Risk and Uncertainty

This appendix deals with the representation of uncertainties and risks in the plan's regional model.[1] It also describes the various studies the Council has performed to understand how the Council's perception of risk and uncertainty bear on its recommendations. A glossary, index, and list of references appear at the end.

This appendix addresses the regional model itself to a limited extent. This appendix identifies a particular range of the model worksheet cells that creates a model "future," the single draw of each source of uncertainty over the study horizon. In the section on "Uncertainties," beginning on page P-18, it describes in detail how the regional portfolio model manifests these modeling futures with Excel® formulas and user-defined functions. The description of the rest of the model, however, appears in Appendix L.

ICON KEY

Key idea

Definition

This appendix provides several tools to help the reader track this discussion. The first tool is the use of icons to flag key definitions and concepts. A table of these icons appears on the left.

The second tool is a set of workbooks containing versions of the regional model, utilities, and a document that describes particular worksheets. The reader can request a copy of these workbooks from the Council or download them from the Council's web site.[2] The first of these files is a compressed file, L24X-DW02-P.zip, containing the workbook files that Appendix L uses. In particular, the file L24DW02-f06-P.xls is a workbook containing a pre-draft plan version of the regional portfolio model. The compressed file also contains examples of utilities and documentation. References to the workbook L24DW02-f06-P.xls appear in curly brackets ("{}"). The second file is L28_P.zip, which contains the workbook L28_P.xls, the regional model that the final plan's preparation used. Note that the treatment of several key sources of uncertainty changed significantly between the draft and final plan. A document in L28_P.zip describes the changes. References to L28_P.xls appear in double curly brackets ("{{}}"). Access to the workbooks should not be necessary for following the discussion in this appendix, however.

References to Council work papers and data sources appear in square brackets ("[]"). The "References" section at the end of the appendix lists these sources. Publicly available sources appear in footnotes. The reader may want to refer to the following Table of Contents for orientation to the remaining appendix.

[1] The reader will find definitions for terms such as "uncertainty," "risk," and "futures" in the glossary. Chapter 6 of the plan also defines and illustrates these terms with examples.
[2] As of this writing, the Olivia_and_Portfolio_Model subdirectory of http://www.nwcouncil.org/dropbox

Table of Contents

Introduction

This appendix begins with a discussion of the Council's approach to decision making under uncertainty. This shapes the means and choice of tools for addressing uncertainty. It also influences the validation of analyses and models. The issue of validation arises not only in the formal validation of the model futures but also extends to basic judgments about assumptions. The Appendix will return many times to the issue of whether the judgments about assumption values are reasonable in the section on "Uncertainties."

Between its discussion of the Council's approach to decision making under uncertainty and the validation of data and models, the appendix introduces the regional model. This serves several purposes. First, the next main section is about the Council's treatment of uncertainties. As mentioned earlier, this appendix identifies a particular range of the model worksheet cells that creates a model "future," the single draw of each source of uncertainty over the study horizon. This introduction identifies that range. The introduction also gives the reader an overview of the philosophy and methods for modeling uncertainty. It describes, for example, the use of Monte Carlo simulation and how the application of this technique facilitates the Council's approach to decision making. Second, this arrangement of topics introduces the means by which the model produces its principal results, the distribution of present value total system costs and associated risk and central tendency measures. This is the topic of the next main section of this appendix, "Risk Measures." Third, discussion of the regional model provides a concrete framework for the discussion of the last section, "Sensitivity Studies." This last section examines not only the purpose and conclusions of the studies, but how Council staff modified the regional model to obtain the results. Finally, the introduction mentions utilities that access regional model output to assist interested parties to perform their own validation of the model's assumptions and results.

Decision Making Under Uncertainty

Strategic decision-making models ***use and manage uncertainty*** differently from many simulation models that incorporate uncertainty. The key difference between the two is the scale of risk and how a decision maker responds to uncertain events.

An example of a simulation that addresses uncertainty, but is ***not*** what we would call a strategic decision analysis, is how many utilities model hydrogeneration. To simulate generation due to hydro streamflow variability, an analyst would create a model using some sample of historical data, say 1939 through 1978 streamflows. The analyst has a great deal of information about the distribution of streamflows. He may be willing to assume that the underlying processes that give rise to the streamflows – and the relationship between generation and streamflows – are stable. Because the variation in hydrogeneration averages out over a sufficient number of years with high probability, the average generation and average system cost are useful statistics, and may be the key outputs of interest.

The decision maker may need to make a choice among different plans to deal with this variation in hydrogeneration, but the tool she uses is essentially sensitivity analysis, albeit sophisticated sensitivity analysis. This kind of analysis is appropriate where the scale of the uncertainty and risk is small enough that the decision maker feels she can live with the outcomes, given the selected plan. In particular, the emphasis is on choosing a plan to which the decision maker feels comfortable committing.

This approach is common to many kinds of analysis. For example, it would be the way an industrial engineer would represent a manufacturing process, if he wanted to maximize productivity. It is the way a civil engineer would model traffic flow, if he were trying to minimize congestion or travel time.

Against these examples, contrast strategic decision analysis. If the scale of change is large, extreme outcomes may be catastrophic. If the outcome would be catastrophic, the decision maker may need to consider individual scenarios. The way each scenario turns out would typically determine how the decision maker would respond to circumstances. Scenario analysis will focus on developing options, deciding what circumstances would trigger the implementation of each option, and evaluating the benefits of using each option. Scenario analysis usually has decision rules or "flags" that tell the decision maker when to change plans or implement options.

An example of strategic decision analysis is planning for a military operation. In the fog of war, leaders must make life or death decisions about tactics and strategy. In addition to the main plan, strategists will develop Plan B, Plan C, and so forth, alternatives to implement if circumstances are not as expected. They create options by deploying resources and small numbers of troops to monitor enemy activity and serve as support if it becomes necessary to adapt to new scenarios.

Note that a general would never consider implementing a fixed strategy, one without options or alternatives, based on average survival. If an option will spare a life, it merits consideration. Whereas the average hydro generation over five or six years is a useful number for certain calculations, such as average power cost, failing to adapt military plans because the expected distribution was acceptable would be ludicrous and tragic. In decision analysis, the tails of the distribution, especially the "bad" tail, assumes greater significance than they do in ordinary simulations. Adaptations that improve the outcomes in the worst of circumstances receive emphasis. Decision making under uncertainty has more to do with making decisions that, while they may not have been optimal in retrospect, did not lead to a catastrophic outcome. This appendix returns to the discussion of managing bad outcomes in the section "Risk Measures."

One of the issues that a decision maker who is making decisions under strategic uncertainty must grapple with is the relative likelihood of each scenario. This issue is central to the question of how much to spend on a given option. If the decision maker believes that scenario A is much more likely than scenario B, which has the same cost, the decision maker might be inclined to spend more to mitigate scenario A. Another difficulty that sometimes arises in scenario analysis is that a decision maker can only evaluate a small number of scenarios. The question arises, "How were these scenarios selected, and how representative are they?"

The next section introduces to a technique, Monte Carlo simulation, which helps address concerns about the likelihood and range of scenarios. The regional model employs Monte Carlo simulation. The regional model, however, also implements planning flexibility. Planning flexibility, described in Appendix L, enables the regional model to evaluate contingency plans and implement those plans as circumstances change during each scenario's study period. Therefore, the regional model performs true strategic decision analysis on a large number of scenarios; effectively, "scenario analysis on steroids."

Another distinction of decision analysis models is how one validates the models. The section that follows the next section discusses those differences.

Monte Carlo Simulation

"Monte Carlo simulation" refers to any method that solves a problem by generating suitable random numbers and observing that fraction of the numbers obeying some property. The method is useful for obtaining numerical solutions to problems that are too complicated to solve analytically.[3] In 1946, S. Ulam became the first mathematician to dignify this approach with a name, in honor of a relative having a propensity to gamble [**1**]. Ulam was involved with the Manhattan project to build the first atomic bomb. Physicists used the technique for evaluating difficult integrals.

The Council applies the Monte Carlo technique to regional resource planning to generate futures. A future consists of a draw of a random vector that represents the following sources of uncertainty:

- Load requirements
- Gas price
- Hydrogeneration
- Electricity price
- Forced outage rates
- Aluminum price
- CO_2 tax
- Production tax credits
- Green tag value

The regional model performs true strategic decision analysis on a large number of scenarios; effectively, "scenario analysis on steroids."

More precisely, the technique produces values for each source of uncertainty in each time period in such a way that the values have the correct correlation with previous and future

[3] The interested reader can consult any of a host of books and Internet resources describing Monte Carlo simulation in general.

values of that source of uncertainty and with the previous and future values of all the other sources of uncertainty. The appendix will explain this technique in detail below.

The principal reason for using Monte Carlo simulation for decision analysis, however, is that it avoids what Richard Bellman referred to as the "curse of dimensionality."[4] To evaluate the outcomes associated with values of uncertainties, an analyst can construct a "decision tree" that associates with each combination of values for the various sources a probability and outcome. The problem, however, is that the "branches" of the decision tree proliferate exponentially with the number of uncertainties addressed. For example, a decision tree with three values of electricity price forecasts ("high," "medium," and "low") would require only three studies. A decision tree large enough to examine three forecasts for each of the nine uncertainties listed above, however, requires $19{,}683 = 3^9$ studies. The regional model uses 750 values for each of 1,045 random variables to represent values in each of the model's 80 periods, which would produce 750^{1045} branches. This number of branches far exceeds the storage capability of any machine imaginable. The regional model, moreover, must perform this calculation roughly a million times to produce a single feasibility space, described below.

Of course, not all of the branches of a decision tree have sufficiently high probability and extreme value that they would contribute much to the solution. It is this observation that leads to Monte Carlo simulation. Monte Carlo simulation chooses random values for each source of uncertainty according to their likelihood.[5] The distribution that results therefore automatically reflects both the likelihood and value of the outcome. Because Monte Carlo simulation is a statistical sampling technique, the criterion for the number of samples is the confidence necessary for statistics of interest, such as the error of the mean or of the mean of a tail. This sample size is typically only weakly sensitive to the number of sources of uncertainty.

The regional model uses Decisioneering Inc.'s Crystal Ball® Excel add-in to perform Monte Carlo simulation. Crystal Ball uses particular terms to refer to the Excel worksheet cells that perform the principal tasks.

Assumption Cells are worksheet cells in a spreadsheet model that contain a value defined by a probability distribution's random variable. These cells are distinguished in the sample workbooks by their distinctive green color. (See, for example, {{R24}}.) This appendix regularly refers to assumption cells in the section "Uncertainties." Crystal Ball reassigns values to each assumption cell at the beginning of each "game" or modeling future.

A **Decision Cell** is a worksheet cell in a spreadsheet model that the user controls. The user controls these indirectly – for example, via an optimizer – or directly. The reader

[4] Bellman, R. (1961), Adaptive Control Processes: A Guided Tour, Princeton University Press.

[5] For a number of good reasons, these values are not truly random in the everyday sense of the word. For example, the random number generator uses a seed value, so that an analyst can reproduce each future exactly for subsequent study. The generator also selects the values to provide a more representative sampling of the underlying distribution, a technique known as Latin Hyper Square or Latin Hyper Cube.

may think of the value of these cells as representing the plan. The optimization program adjusts the decision cells in the regional portfolio model to minimize cost, subject to risk constraints. Appendix L details the function and application of decision cells in the section “Parameters Describing the Plan,” page L-72. These cells are yellow in the regional model. (See, for example, {{R2}}.)

Forecast Cells contain statistical output of the model. The default color for these cells is turquoise. In the regional model, the primary forecast cell is the NPV cost for a plan under a 20-year future, {{CV1045}}. Other forecast cells in the regional model, such as those that regional model macros assign risk values, serve to communicate data back to the OptQuest optimizer.

The assumption and decision cells are, in a sense, the exogenous inputs to the model; the forecast cells report the output. The topic of the next section is the calculation engine that processes the input and produces the output.

Logic Structure of the Portfolio Model

To understand how the regional portfolio model represents uncertainty and generates the system cost values that give rise to risk, it is useful to understand the model itself. The treatment of uncertainties, like load and hydro generation, are to some extent separable from the rest of the model. This section identifies a particular range of the model worksheet cells that creates futures. (See page P-14.) Likewise, the forecast cells that report the final costs and risks inhabit a small range of adjacent cells. The description of the rest of the model appears in Appendix L. The following provides a brief introduction that should be sufficient for understanding that portion of the model that simulates sources of uncertainty.

The Council calls its approach to resource planning “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.[6]

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 a plan 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 P-1 below. Each box in Figure P-1 represents the net present

---

[6] See Appendix L.

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 P-1.

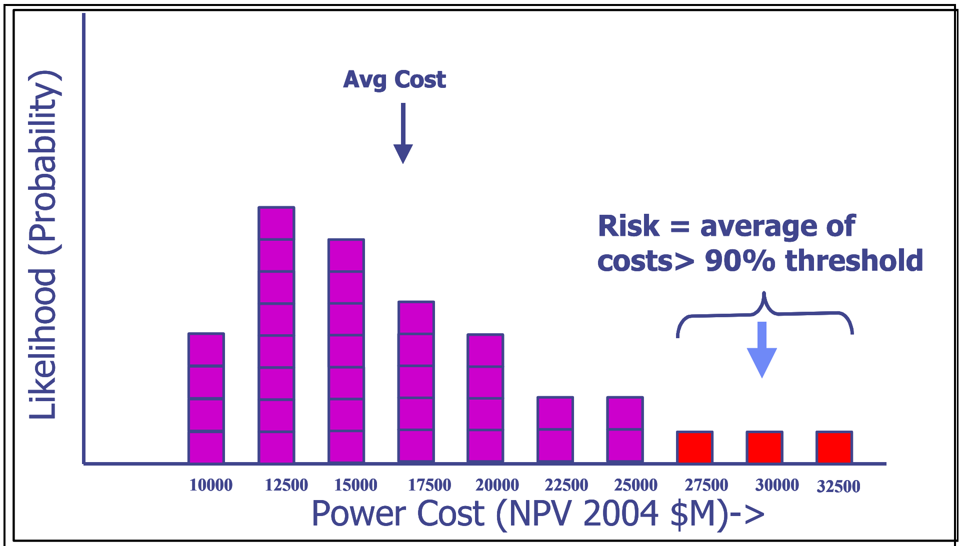

Figure P-1: Plan Cost Distribution, Average Cost and Risk (TailVaR90)

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. A discussion of this choice of risk measure and its comparison with other risk measures appears in section "Risk Measures," below. 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).

Figure P-1 represents the cost distribution associated ***with a single plan***. If the outcomes for different plans are plotted as points, with coordinates given by the expected cost and risk of each plan, one obtains the new distribution illustrated in Figure P-2. 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 versa. The "best" outcome on the efficient frontier depends on the risk that can be accepted.

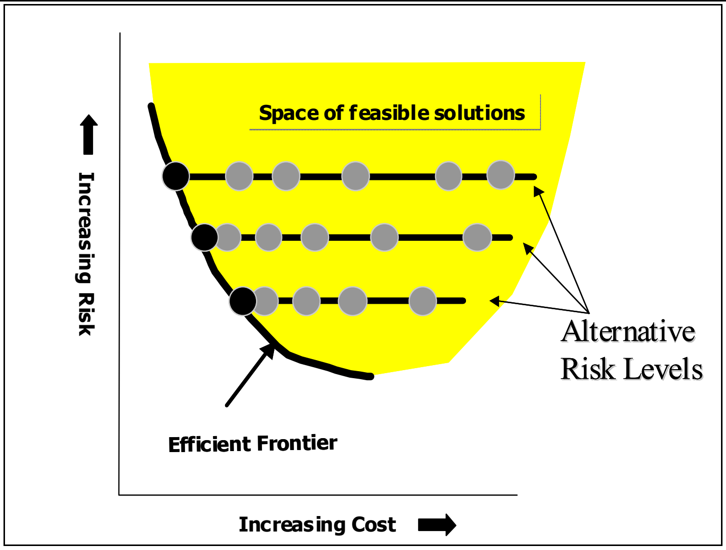

Figure P-2: Feasibility Space

When a user opens the portfolio model workbook, the values they see are values for a particular future and for a particular plan. It is within this future or "game" that the energy and cost calculations take place. How, then, are the futures changed to create a cost distribution for a plan and the plans changed to create the feasibility space?

Figure P-4 illustrates the overall logic structure for the modeling process. The optimization application, the Decisioneering, Inc. OptQuest™ Excel® add-in, controls the outer-most loop. The goal of the outer-most loop is to determine the least-cost plan for each level of risk. It does so by starting with an arbitrary plan, determining its cost and risk, and refining the plan until refinements no longer yield improvements. The program first seeks a plan that satisfies a risk constraint level. Once it has found such a plan, the program then switches mode and seeks plans with equal (or lower) risk but lower cost. The process ends when we have found a least-cost plan for each level of risk. This process is a form of non-linear stochastic optimization.[7]

[7] The interested reader can find a more complete, mathematical description of the optimization logic in reference the following references:
Glover, F., J. P. Kelly, and M. Laguna. "The OptQuest Approach to Crystal Ball Simulation Optimization." Graduate School of Business, University of Colorado (1998). Available at http://www.decisioneering.com/optquest/methodology.html ;
M. Laguna. "Metaheuristic Optimization with Evolver, Genocop, and OptQuest." Graduate School of Business, University of Colorado, 1997. Available at http://www.decisioneering.com/optquest/comparisons.html; and
M. Laguna. "Optimization of Complex Systems with OptQuest." Graduate School of Business,

The optimizer OptQuest controls the Crystal Ball® Excel add-in. OptQuest hands a plan to Crystal Ball, which manifests the plan by setting the values of decision cells in the worksheet. These are the yellow cells in {range R3:CE9}. Crystal Ball then performs the function of the second-outer-most loop, labeled "Monte Carlo Simulation," in Figure P-4. Crystal Ball exposes the selected plan to 750 futures. For each future, Crystal Ball assigns random values to 1,045 assumption cells, the dark green cells throughout the worksheet. (See for example, {R24}.) Crystal Ball then recalculates the workbook. In the portfolio model, however, automatic recalculation is undesirable, as described in Appendix L. The portfolio model therefore substitutes its own calculation scheme. It uses a special Crystal Ball feature that permits users to insert their own macros into the simulation cycle, as shown in Figure P-3. Before Crystal Ball gets results from the worksheet, a macro recalculates energy and cost, period by period, in the strict order illustrated in Figure P-5 and Figure P-6 and as described on page P-13. The reason for performing its own calculations is to assure calculations take place in a strict chronological order, as required by several mechanisms in the model, including the planning flexibility. The values in the Crystal Ball forecast cells then contain final net present value (NPV) costs that Crystal Ball saves until the end of the simulation. Forecast cells are those that have the simulation results and have a bright blue color. The NPV cost, for example, is in {CV1045}.

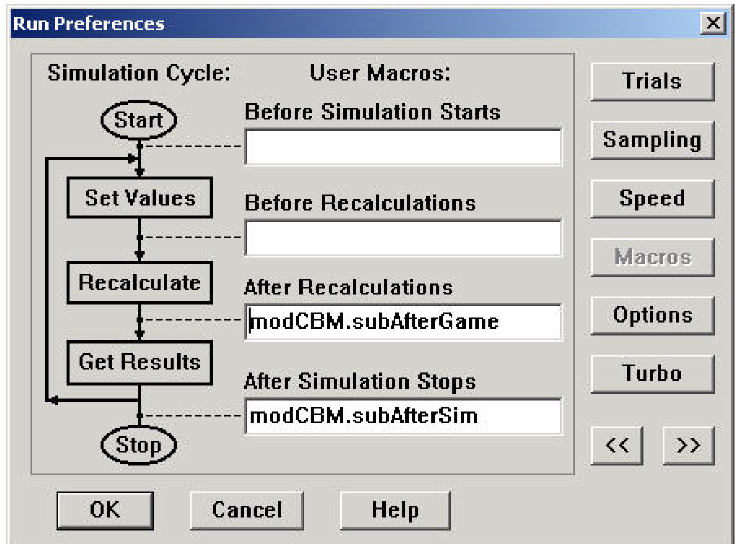

Figure P-3: Crystal Balls Macro Loop

After the simulation for a given plan is complete and Crystal Ball has captured the results for all the games, the last macro in Figure P-3 fires. This macro calculates the custom risk measures and updates their forecast cells. The custom risk measures include, for example, $TailVaR_{90}$, $CVaR_{20000}$, VaR_{90}, and the 90^{th} Quintile. Finally, Crystal Ball returns the distribution of cost associated with futures and the distribution's risk measures to OptQuest.

University of Colorado, 1997. Available at http://www.decisioneering.com/optquest/complexsystems.html

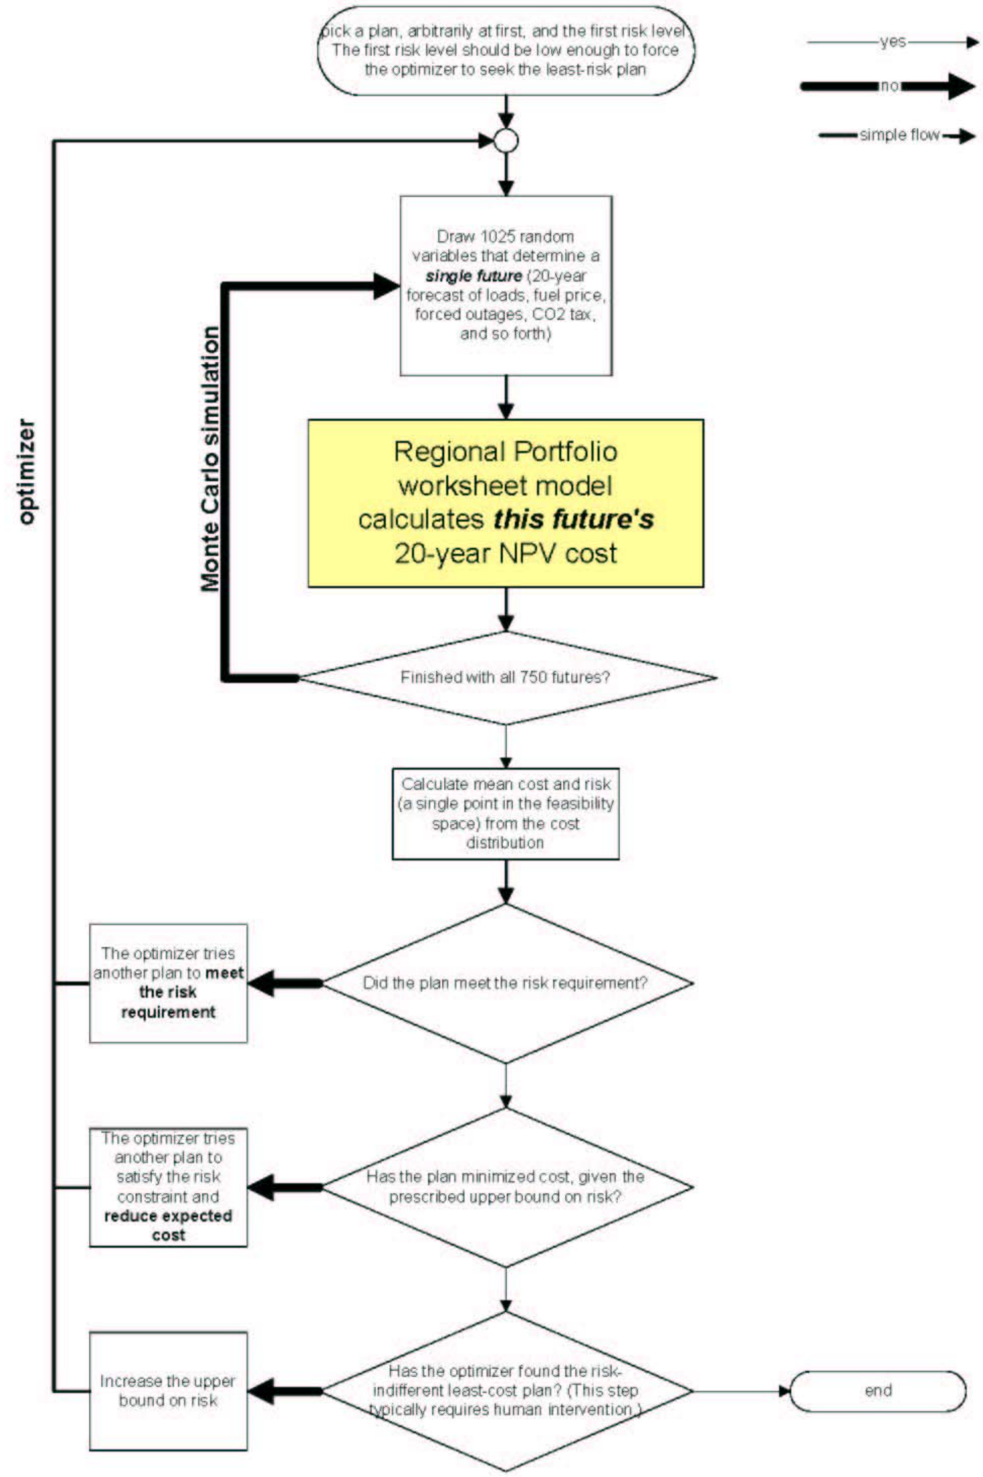

Figure P-4: Logic Flow for Overall Risk Modeling

The portfolio model performs the duties of the innermost task, identified by the shaded box in Figure P-4. Given the values of random variables in assumption cells, the portfolio model constructs the futures, such as paths and jumps for load and gas price,

forced outages for power plants, and aluminum prices over the 20-year study period. It does this only once per game. It then balances energy for each period, on- and off-peak and among areas, by adjusting the electricity price, as illustrated in Figure P-5. The regional portfolio model uses only two transmission zones, however, the region and the "rest of the interconnected system," although it does model some geographic diversity of fuel and electricity price. Only after it iterates to a feasible solution for electricity price in one period does the calculation moves on to the next period. After calculating price, energy, and cost for each period, the model then determines the NPV cost of each portfolio element and sums those to obtain the system NPV. This sum is in a forecast cell.

Some worksheet cells are involved in the energy rebalancing calculation. These cells, many of which contain formulas for electricity prices, must recalculate multiple times for each subperiod. These and other cells that rely on them, such as those that control the long-term interaction of futures, prices, and resources, are the "Twilight Zone" (TLZ) of the regional model. This portion of the worksheet also contains formulas for price elasticity of load and decision criteria.

Figure P-6 illustrates the calculation order described above. The number in the parentheses is the order. The plus sign (+) is a reminder that iterative calculations take place in the area. The workbook calculates the primary uncertainties only once per game, and their cells are near the top of the worksheet {rows 26-201}. (Plant forced outages are the exception. These cells are located elsewhere, as explained below and detailed on page P-81.) The cells associated with the uncertainties are denoted "Futures (1)" in Figure P-6.

Figure P-6 identifies those recalculations that must be made multiple times per subperiod by TLZ {rows 202-321}. NP stands for on-peak {rows 318-682}; FP stands for off-peak {rows 684-1058}. The area at the far right refers to the NPV summary calculations {range CU318:CV1045}.

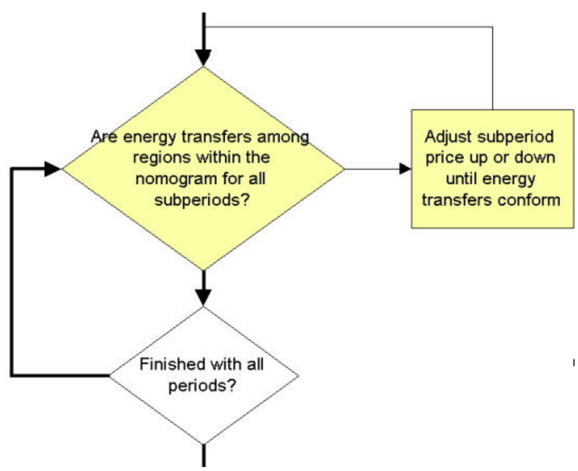

Figure P-5: Logic in the Regional Portfolio Worksheet Model

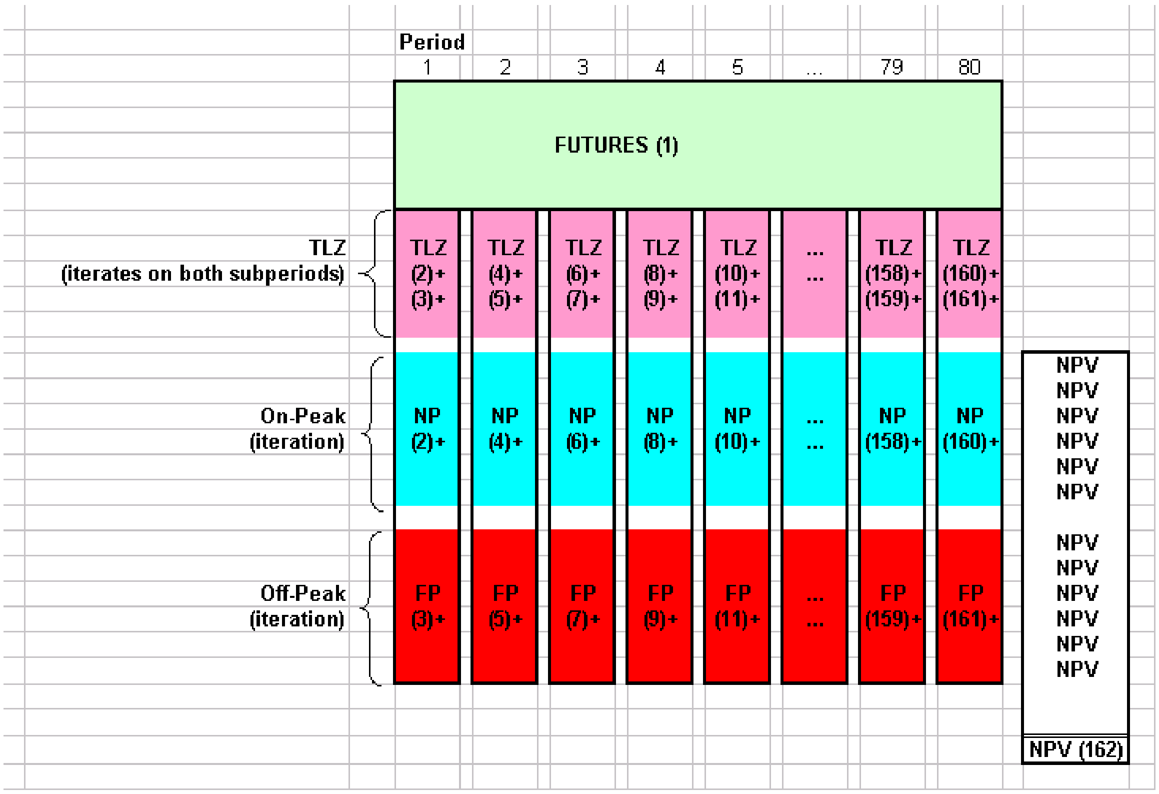

Figure P-6: Portfolio Model Calculation Order

Appendix P documents the uncertainties in the regional portfolio model. This includes the worksheet formulas for describing the uncertainties. Because it would be redundant to cover the same material in Appendix L, Appendix L describes everything *except* the uncertainties.

Figure P-6 permits us to state the scope of this appendix with respect to ranges within of the portfolio model. This **Appendix P**, and in particular the section "Uncertainties," describes the calculations in the area of the worksheet denoted by "FUTURES (1)" and with the dark green assumption cells for plant forced outages associated with each power plant. For ease of reference, the worksheet calculates a future's forced outages in rows associated with the power plants themselves. Consequently, the user will find them not in the range marked "Futures (1)" in Figure P-6, but down in the rows associated with "NP" and "FP" calculations. **Appendix L** discusses all ranges of the regional model except that denoted by "FUTURES (1)".

Model Validation

Given the differences between decision-making models and other simulation models that incorporate uncertainty, it should not be too surprising that how one validates the two differs. This section discusses some of those differences, with attention to the treatment of validation in the regional model.

The example of a simulation model that began this section was a hydrogeneration estimator. To validate the hydrogeneration model, an analyst would make some prediction about how the model would perform with a new set of streamflows. She would be concerned about how well the model reproduced certain patterns of generation. To validate her model, she would then apply the model to a new set of historical streamflows, say 1979 through 1990, and compare the model generation with the actual generation over those years. An analyst would apply a similar process in constructing and validating simulation models for other systems where stochastic processes are important, such as for vehicle traffic flow or industrial manufacturing processes.

With strategic decision-making models, this approach does not work. The past is not a good standard for the future, because we have assumed our modeling futures differ dramatically from one another. It may be appropriate to look at a single future that resembles some past event to see how reasonable the model responds. This, however, is effectively a one-point sample of possible futures. By design, there are many possible futures, and the model should prepare the decision maker for futures that are unanticipated and unfamiliar.

This is not to say that there is no role for more traditional validation. There is a distinction, however, between short-term variation and strategic uncertainty. If we think of an example like electrical load requirements, we recognize there is some short-term variation due to weather and seasonality. We may tend to believe we understand this variation rather well and expect future variation to resemble that which we have seen in the past. This kind of variation lends itself well to statistical analysis of past behavior and patterns.

Once we attempt to forecast load requirements beyond a couple of years, however, we enter the realm of strategic uncertainty. We recognize there are many things that can affect system load requirements. Economic disruptions within and outside the region and technological innovations, for example, can greatly influence energy requirements. We may expect that there is strong chronological correlation in load requirements, i.e., load in a given month will not differ significantly from load in the previous month beyond what we expect from seasonal variation. The underlying tendency or path of system load requirements, however, can move in a host of different directions, so that after just a few years, system load requirements can differ significantly from the expected forecast.

While previous statistical patterns may be helpful in validating the short-term variation behavior of the model, they do not help with strategic uncertainty. Fundamental models, which relate strategic behavior to underlying processes, can be helpful in understanding and reducing strategic uncertainty. Even fundamental models, however, rely on assumptions that are plagued by uncertainty once forecasts extend beyond a few years. Moreover, because of their associated computational burden, it is difficult to incorporate a fundamental model directly into a model for decision making under uncertainty.

Ultimately, the representation and validation of strategic uncertainty is highly subjective. Expert opinion, often formed through careful consideration of many sources of information, including the results of fundamental models, is the arbiter of credibility. When they are available, ranges of expert forecasts can help validate possible futures. The Council attempts to achieve regional model consistency with its forecasts for electricity load requirements and natural gas prices, for example.

This approach is certainly not without its shortcomings. Those who have examined case histories of decision making under uncertainty have noted that experts often overestimate their ability to forecast the future.[8] That is, experts tend to underestimate uncertainty. We do not have to look any further than the load forecasts made by utility experts in the 1970s and 1980s to find examples where each year, the load forecasts fell below the lower jaw of the previous year's set of load forecasts. The Council's own oil price forecasts since the 1980s provide another example where actual prices repeatedly fell outside the range of bounding (high and low) forecasts [**2**]. (See Figure P-7.)

While recognizing these shortcomings, the Council has elected to validate the regional portfolio model using experts' reviews of the futures used in the model. In Appendix L, the reader will find a description of the utility for data extraction and spinner graphs. This utility, and specifically the set of graphs embedded in the principal worksheet, permit anyone to quickly scan through all 750 futures. The user can view the nature and range of futures. For each future, the user can simultaneously view the 20-year projection of electricity prices, loads, natural gas prices, and so forth, for that future. In addition, the user can also witness how power plants are built out under that future and how much energy generation there is by technology for each period under that future. They can view the period costs and net present value cost, and most of the other variables that an analyst would want to see to verify that the model is behaving correctly.

[8] See, for example, John T. Christian, Consulting Engineer, Waban, Massachusetts Geotechnical Engineering Reliability: *How well do we know what we are doing?* The 39th Terzaghi Lecture, Spring 2005 GeoEngineering Seminar Series, Annual GeoEngineering Society Year-End Distinguished Lecture Program and Banquet, University of California at Berkeley

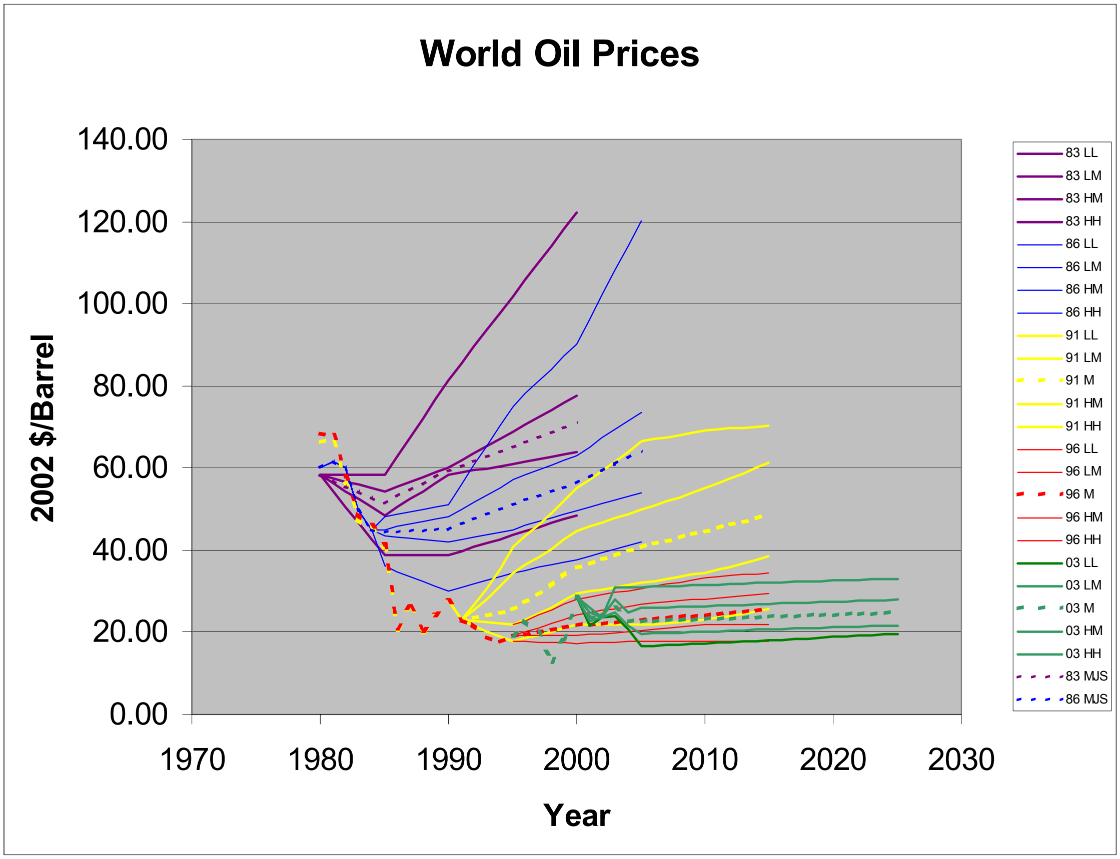

Figure P-7: Council's World Oil Price Predictions

This utility provides the principal means of validation. Rather than attempting to understand statistical distributions for each source of uncertainty in the relationship to other sources of uncertainty, an analyst can witness the final behaviors and see how they stand in relationship to each other. The Council's System Analysis Advisory Committee (SAAC) and the Council have reviewed these futures and found them to be reasonably representative of possible future behaviors.

With this overview of the decision making under uncertainty, this appendix starts the first section, the detailed description of the model's treatment of uncertainties.

Uncertainties

This section consists of two main parts. The first part is an introduction to Stochastic Process Theory implemented in the regional model. There are six main discussions:

- Log normal distributions
- Geometric Brownian motion (GBM)
- GBM with mean reversion
- Simulating Values for Correlated Random Variables
- Principal factor decomposition
- Stochastic Adjustment
- Jumps

The regional model uses each of these techniques to represent the future behavior of sources of uncertainty. The discussion will identify how each technique captures both short-term variation and strategic uncertainty.

The second part of this section steps through each source of uncertainty and describes why that source of uncertainty is modeled the way that it is. The uncertainties include:

- Load requirements
- Gas price
- Hydrogeneration
- Electricity price
- Forced outage rates
- Aluminum price
- CO_2 tax
- Production tax credits
- Green tag value

It explains how each source of uncertainty uses the chosen stochastic process to achieve the desired behavior. It also documents data sources and provides a reference to the sample worksheet to provide a detailed description of how the formulas in the worksheet implement the desired stochastic behavior.

Stochastic Process Theory

Lognormal distributions are a key characteristic of geometric Brownian motion (GBM) and GBM with mean reversion. The regional model uses lognormal distribution in the electricity price, fuel price, load requirements, and aluminum price processes. This discussion therefore starts with a review of the lognormal distribution and then describes the GBM and the GBM with mean reversion processes. Principal factor analysis technique does not rely per se on any of these, and the section will review this technique last.

Lognormal Distribution

It might be useful to understand why the lognormal distribution finds such intensive use in the regional portfolio model and in other simulation and valuation models. There are three reasons the regional model uses lognormal distributions:

1. It solves problems we encounter with simpler distributions,
2. It has an nice intuitive rationale, and
3. It describes much data better than simpler (and sometimes more complex) distributions.

To understand these advantages, we start by examining the problems that a naïve application of simpler distributions might encounter.

If an inexperienced analyst with some background in statistics were to approach the challenge of modeling stochastic prices, he might try to use a simple distribution, such as the normal distribution. However, any unbounded, symmetric distribution, like the normal distribution, must produce negative numbers, as illustrated in Figure P-8. Negative prices, however, are bothersome and may cause some programs to fail in mysterious and unpredictable ways. One fix to this problem is to use an asymmetric, bounded distribution, such as the triangular distribution, to keep prices positive. Of course, the drawback to this approach is that because the distribution has both a lower and upper bound, the analyst must now provide some rationale for choosing the value of the upper price limit.

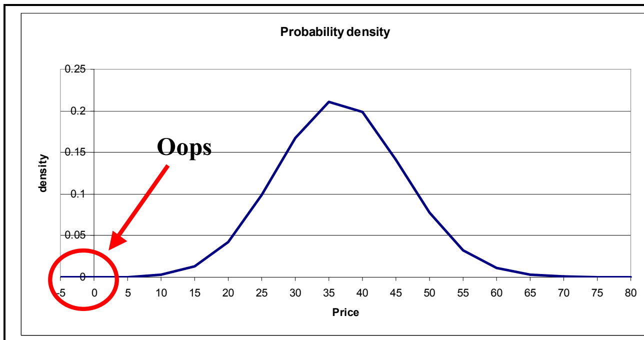

Figure P-8: Symmetric Distribution

The second problem the analyst might encounter would be difficulty in performing meaningful statistics on prices. There are several issues here.

First, prices for commodities typically are not symmetric. Because they are bounded below by zero, but are unbounded above in principle, they can be strongly skewed. This means that simple distributions, like the normal distribution, and statistical tests based on these distributions, do not work. For example, one can not say that 95 percent of the observations lie within two standard deviations of the mean.

Second, prices can drift in ways that mask the information in which an analyst might be interested. To illustrate this, suppose an analyst were interested in estimating the daily variation for natural gas price. Perhaps she is interested in estimating the likely change in natural gas price between today and tomorrow. Because she is interested in the change in daily price, it makes sense to use daily prices for the statistical sample, as opposed to hourly prices or weekly prices. To get a representative sample, she uses the last 100 days of natural gas price history, illustrated in Figure P-9. If she made the mistake of calculating the variation in prices, as measured by their standard deviation, without studying the data beforehand, she would compute the standard deviation to be about \$0.83. The one standard deviation bound around the average price appears in Figure P-11. Clearly, this overestimates the daily price variation. The actual daily price variation is closer to the \$0.14 that Figure P-10 illustrates. If she did discover that price drift was distorting the estimate of price variation, she would need to develop a model of the underlying drift or seasonality to remove that influence.

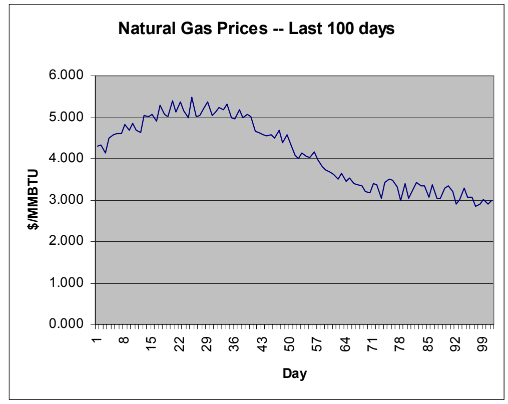

Figure P-9: Price Variation

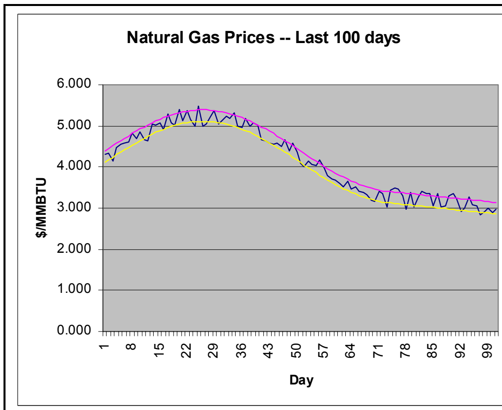

Figure P-10: Daily Price Uncertainty

Third, prices are often the wrong variable to study. Natural gas, for example, is a commodity traded by both hedgers and speculators. Both of these groups, but perhaps especially speculators, buy and sell natural gas to maximize profit. Now, initially it may appear that a \$1.50 price increase of natural gas is equally attractive (or costly) irrespective of whether the underlying price of the gas is \$3.00 or \$4.50. The gross profit would be \$1.50 times the quantity of gas. This ignores the fact, however, that an investor can buy more \$3.00 gas than they can buy \$4.50 gas. That is, what investors are interested in is the return on dollar invested: p_t/p_{t-1}, where p_t is the price today and p_{t-1} was the price yesterday.

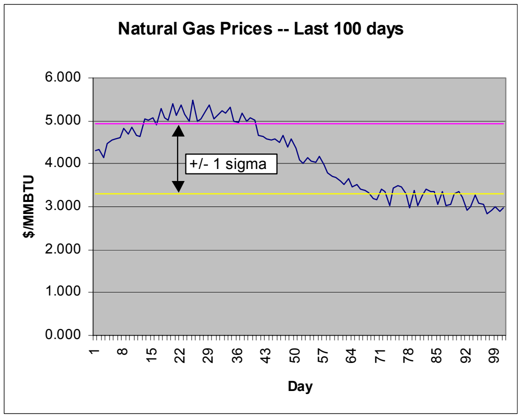

Figure P-11: Standard Deviation of Price Data

The same is true for other commodities and for financial investments. The return on the investment is what matters, not the price. In fact, an analysis of prices for stocks and commodities show that returns, not prices, bump up and down symmetrically in the very short term (hourly or daily) as new information is forthcoming and they are traded. *Symmetry of returns* often explains a large portion of the *asymmetry of prices* described in the first paragraph.

Another advantage of using price returns instead of prices is that the second problem mentioned above disappears. That is, if the analyst uses daily price returns, she will obtain an estimate of daily price variation that more closely resembles that illustrated in Figure P-10.

Using returns, however, seems to give rise to yet another problem: calculating meaningful statistics on price return is tricky. Let us say that our analyst discovers that return on prices has about a normal distribution. It may be easy to calculate the mean and standard deviation of this distribution, but what do these numbers represent? To see the problem with interpreting these statistics, consider the following example. Suppose we have a simple sample of two observations, 50 percent price return and -50 percent price return. That is, on day two, the price increases 50 percent from that on day one; on day three, the price decreases 50 percent from that on day two. The naïve average of these two would be zero price return. In fact, however, we would have

$$1.5 \times 0.5 = .75$$

That is, our final price return would be -25 percent. It is unclear what the average of this distribution means and it is even less clear how to extract meaningful information out of the standard statistics of this distribution.

Consider now taking the log transformation of price return:

$$y_t = \ln(p_t / p_{t-1}), \text{sometimes also denoted } p_t / p_{t-1} \overset{\ln}{\mapsto} y_t \qquad (1)$$

This has the inverse transformation:

$$p_t / p_{t-1} = e^{y_t}, \text{sometimes also denoted } y_t \overset{\mathrm{e}}{\mapsto} p_t / p_{t-1} \qquad (2)$$

The transformed variable y_t has properties that solve the problems this section has raised and has some additional nice properties, as well. First, for small returns the logarithm of returns has a value close to that for the regular return:

$$\text{As } p_t \to p_{t-1}, \text{ then } y_t \to p_t / p_{t-1} - 1$$

The sum and the average of the transformed returns have straightforward and useful interpretations. The sum of the transformed return is the total return over the period:

$$
\begin{aligned}
&\ln(p_2/p_1)+\ln(p_3/p_2)+\cdots+\ln(p_n/p_{n-1})\\
&=\ln(p_2)-\ln(p_1)+\ln(p_3)-\ln(p_2)+\cdots+\ln(p_n)-\ln(p_{n-1})\\
&=\ln(p_n)-\ln(p_1)\\
&=\ln(p_n/p_1)\overset{e}{\mapsto} p_n/p_1,\ \text{the return over the period}
\end{aligned}
$$

The average of the transformed return is the periodic growth rate, also called the geometric mean:

$$
\begin{aligned}
&\frac{1}{n}\left(\ln(p_2/p_1)+\ln(p_3/p_2)+\cdots+\ln(p_n/p_{n-1})\right)\\
&=\frac{1}{n}\ln(p_n/p_1)\\
&=\ln((p_n/p_1)^{\frac{1}{n}})\overset{e}{\mapsto}\sqrt[n]{p_n/p_1},\ \text{the periodic growth rate}
\end{aligned}
$$

The reader will recognize this as the constant rate of growth that, if applied in each period, would increase or decrease the value in the first period to the value in the last period.

If the returns have normal distribution, the prices are said to have lognormal distribution. The lognormal distribution is bounded below by zero and unbounded above, as Figure P-12 illustrates. The population standard deviation of the transformed returns

$$
\sigma_y=\sqrt{\sum_{t=1}^{n}(y_t-\bar{y}_t)^2}=\sqrt{\frac{\sum_{t=1}^{n}y_t^{\,2}}{n}-\bar{y}_t^{\,2}} \qquad (3)
$$

and its inverse transformed value give uncertainty bounds consistent with those illustrated in Figure P-10. Standard quantitative finance texts typically refer the value of σ in Equation (3) (or the corresponding *sample* standard deviation) as the "volatility" of the price sequence. For small values, this volatility approaches the standard deviation of returns.

Standard statistics for the transformed variables are relatively easy to compute and are readily available. For example if μ and σ are the mean and standard deviation of a normally distributed variable, such as the transformed returns y_t,

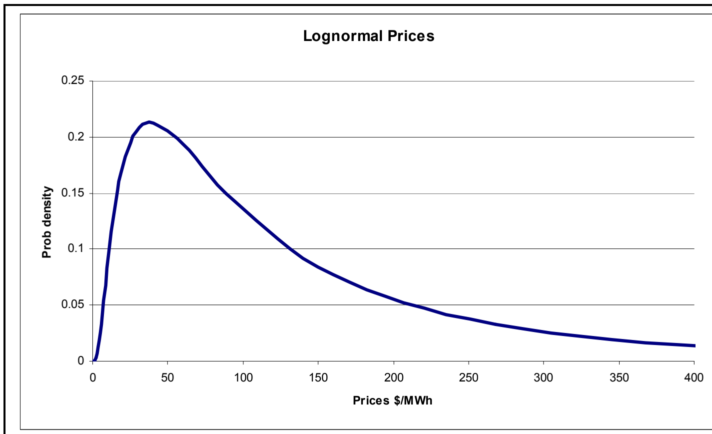

Figure P-12: Lognormal Distribution

$$\text{pdf } f(x) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-(\ln x - \mu)^2 / 2\sigma^2}$$

$$E(x) = e^{\mu + \sigma^2 / 2}$$

$$\text{var}(x) = (e^{\sigma^2} - 1)E^2(x) = (e^{\sigma^2} - 1)e^{2\mu + \sigma^2}$$

where as usual, *E(x)* is the expectation of the lognormally distributed *x* and var*(x)* is the variance of *x*.

Geometric Brownian Motion

The previous section made passing reference to the behavior of prices, bumped around by short-term purchases and sales of the commodity in the market. A standard quantitative representation of this process is Brownian motion. Brownian motion assumes that changes in location (or price) take place in discrete steps. At each step, displacement is determined by a sample from a normal distribution with constant means zero and constant standard deviation sigma.

The standard deviation of the distribution for the sum of these steps is a well-known formula. If there are T steps, the standard deviation is

$$\sqrt{\sigma_1{}^2 + \sigma_2{}^2 + \cdots + \sigma_T{}^2}$$
$$= \sqrt{T\sigma^2} = \sigma\sqrt{T}$$

The standard deviation grows as the square root of the number of steps, as illustrated in Figure P-13.

The previous section explained that the distribution of transformed returns, y_t, is normal for many investments and commodity prices. If the transformed returns follow the kind of process described above, the corresponding prices are said to follow geometric Brownian motion (GBM). At each step, prices have lognormal distribution.

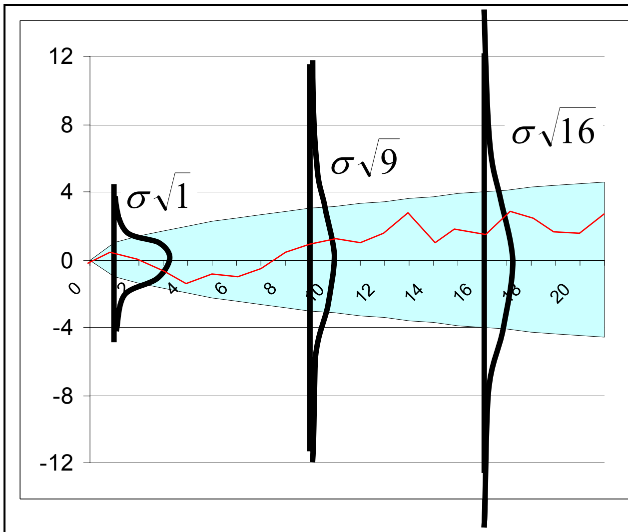

Figure P-13: Uncertainty Growth of GBM Process

GBM with Mean Reversion

Some commodity prices, instead of drifting away from their starting point, instead tend to return to some equilibrium level. This appendix and **Appendix L** describe how fundamental models will produce long-term equilibrium prices that equal long-run marginal costs for new capacity. The long-term equilibrium price represents the level to which prices trend whenever substantial excursions occur. Away from the equilibrium price, long-term supply and demand do not balance, and fundamental economic forces contrive to rebalance them.

There are several price models for a geometric Brownian motion with mean reversion. The regional model uses the following to represent aluminum prices.[9]

$$dp_t = a(b - p_t)dt + \sigma p_t dz$$

where

p_t is stochastic variable in question

dp_t is the change in p_t from the previous step

dz is a drawn from a N(0,1) process

dt is the step size, which has value 1 for discrete processes

a is constant which controls the rate of reversion

b is the equilibrium level

σ is the standard deviation of the log transformed process

The process is identical to an Ito process for a lognormally distributed random variable, but with a drift term that incorporates mean reversion. As prices depart from the equilibrium price b, the term $(b - p_t)$ becomes larger and forces the price back to equilibrium. The strength of the reversion is determined by the constant a. The first-order autocorrelation of price provides an estimate of the value of the constant a. If the constant a has value zero, there is no mean reversion and the price process resembles that of standard GBM. Price will drift away from the starting point with increasing probability. This corresponds to zero autocorrelation. If the constant a is 1.0, the price fluctuates around the equilibrium price and does not drift.

The section "Aluminum Price," beginning on page P-83, describes how this price process represents future aluminum prices. That section includes an explanation of how Excel formulas implement the price process.

There are many other price process models. Some of the more popular models employ jump diffusion and jump diffusion coupled with mean reversion. For the purposes of the regional model, however, these models are excessive. Studies of natural gas and electricity prices suggest that simple geometric Brownian motion does a good job of describing those prices.

Simulating Values for Correlated Random Variables

For each future, the model must generate a large number of correlated values for the stochastic variables. This section describes one standard technique for doing so. The next

---

[9] See, for example, Hull, John C., ***Options, Futures, and Other Derivatives***, 3rd Ed., copyright 1997, Prentice-Hall, Upper Saddle River, NJ., ISBN 0-13-186479-3, page 422

section uses a simplification of this technique to obtain a more economical representation of strongly correlated values.

Suppose that we have a vector $\boldsymbol{\varepsilon}$ of m values ε_j which have some covariance structure Σ. Recall that the covariance matrix is constructed by taking the expectation of the outer product[10] of the vector of deviations from the mean vector **u**:

$$\Sigma = E((\varepsilon - \mathbf{u})(\varepsilon - \mathbf{u})') \quad (5)$$

Because the covariance matrix is a positive definite, symmetric matrix of real numbers, it has representation as the product of its Cholesky factors $\Sigma = TT'$, where T is a lower triangular matrix with zeros in the upper right corner.[11]

Now, take another m-vector η composed of independent variables with zero mean and unit variance. The covariance matrix of the vector η will just be the m x m identity matrix. If we construct the vector $T\eta$, we discover its covariance matrix is

$$\begin{aligned} E(T\eta\eta'T') &= TE(\eta\eta')T' \\ &= TIT' = TT' = \Sigma \qquad (6) \end{aligned}$$

Thus, the vector $T\eta$ has the requisite covariance structure. If we were working with the correlation structure instead of the covariance structure, the conversion is easy. The covariance matrix transforms into the correlation matrix by a simple operation using the diagonal matrix of standard deviations, *D*:

$$\Sigma = DRD \qquad (7)$$

For an example of how to generate correlated values, consider the two-vector $\boldsymbol{\varepsilon}$, where the variables both have zero mean and unit variance. The covariance matrix is the same as the correlation matrix:

$$\begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$$

By the existence of the Cholesky decomposition, there are variables t_{11}, t_{12}, and t_{21}, such that

$$\begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} = \begin{bmatrix} t_{11} & 0 \\ t_{12} & t_{22} \end{bmatrix} \begin{bmatrix} t_{11} & t_{12} \\ 0 & t_{22} \end{bmatrix} = \begin{bmatrix} t_{11}^2 & t_{11}t_{12} \\ t_{11}t_{12} & t_{12}^2 + t_{22}^2 \end{bmatrix}$$

[10] For our purposes, an outer product is the matrix product of a (column) vector right-multiplied by its transpose. This multiplication creates a matrix instead of a scalar, which inner products produce.
[11] See, for example, Burden and Faires, ***Numerical Analysis***, 4th ed., ISBN 0-53491-585-X, Corollary 6.26 and Algorithm 6.6, page 370.

Because the Cholesky matrix is triangular, we can find the values for the entries in the Cholesky matrices by successive substitution:

$$\begin{aligned} t_{11}^{\ 2} &= 1 \\ t_{11}t_{12} &= \rho \\ t_{12}^{\ 2} + t_{22}^{\ 2} &= 1 \end{aligned}$$

so

$$\begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \rho & (1-\rho^2)^{1/2} \end{bmatrix} \begin{bmatrix} 1 & \rho \\ 0 & (1-\rho^2)^{1/2} \end{bmatrix}$$

which means

$$\varepsilon = \begin{bmatrix} 1 & 0 \\ \rho & (1-\rho^2)^{1/2} \end{bmatrix} \eta \qquad (8)$$

Of course, this technique applies to vectors of arbitrary dimension. Note, however, the number of non-zero entries in T increases as $(m^2+m)/2$, as do the number of multiplications and additions, roughly, to create a sample vector. When m is large, the computation burden can increase dramatically. For this reason, practitioners have developed various numerical efficiencies to reduce the computation burden. One of these efficiencies is the topic of the next section.

Principal Factor Decomposition

Principal factor analysis is a general statistical technique for capturing complex statistical behavior with a small number of random variables. In the regional model, principal factor analysis simplifies the representation of strategic uncertainties that have strong chronological correlation, i.e. follow some underlying path over time.

Natural gas price has such a strategic uncertainty, as well as short-term variation due to weather effects and regional economics. For example, consider the price path illustrated by the dotted line in Figure P-14.[12] One way to model the path is by adding up several simpler paths, each of which is a draw from a separate statistical population of similar, simple paths. The advantage of this approach is that the resulting sum will look like a path, i.e., the entries will be strongly correlated, and it gives rise to a great number of possible such paths. This section explains how to perform the construction.

[12] Figure P-14 illustrates the ranges of natural gas prices that the Council adopted for the plan. The middle, solid line is the median price forecast; there is equal probability that annual prices will lie above and below this line.

Before the reader attempts to work his or her way through this section, which is among the more mathematically challenging, they should be aware of its purpose. The regional model implements an adaptation of the concepts presented here. While these concepts have rigorous application to statistical problems with abundant and representative data, the application in the regional model is more art than science. While this is consistent with the spirit of validation articulated on page P-16, it means that understanding the mathematics is not essential to grasping the basic technique of adding up constituents "sub-paths" point-wise. This section merely provides the basis for the technique, to assure the reader that it is neither arbitrary nor original.

Before tackling the construction of paths for future prices under strategic uncertainty, we begin with a simpler construction, one for which data exists and that may be more familiar to some readers. Suppose that, instead of representing strategic natural gas price uncertainty, Figure P-14 represented possible *forward or futures* prices for natural gas. Suppose further, that our objective were to estimate *tomorrow's* forward curve for natural gas, that is, tomorrow's prices for future delivery of natural gas in each year through 2024. There is data about the variation in the forward curves for natural gas price, in principle, because each day traders buy and sell gas forward. Every day, for example, traders buy and sell 2006 gas, and it is possible to get statistics about how that price varies. Others statistics of interest that we can obtain is how the price of 2006 gas price correlates with that of 2005, 2007, and all other years.

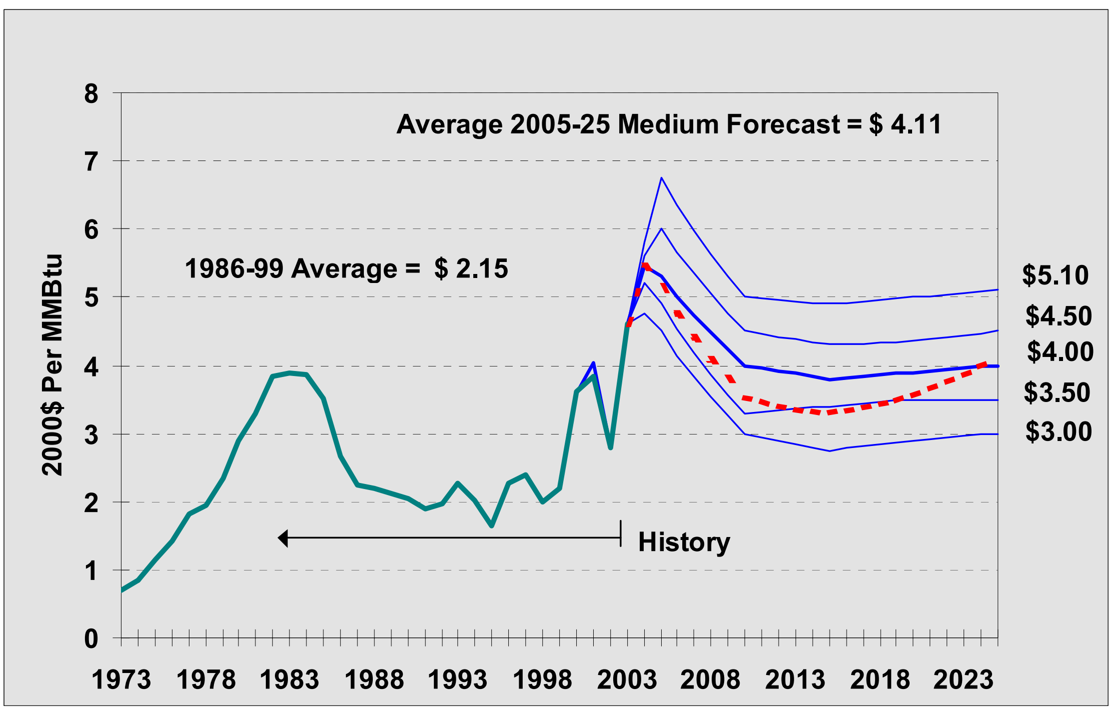

Figure P-14: Possible Natural Gas Prices

For this purpose, the medium forecast of natural gas prices in Figure P-14 will play the role of today's forward curve. The higher and lower price forecasts will represent the typical daily variation in the forward curve. (We will not use the higher and lower price forecasts directly in this example, so we do not need to be precise in how we think about them or their magnitude.)

The dotted line in Figure P-14 will play the role of one possible forward curve that may materialize tomorrow. We want to be able to generate many such forward curves, say, because we are valuing a portfolio of natural gas forward positions and want to understand how much variation and risk there may be in holding that portfolio overnight.

Recall from the discussion of "Lognormal Distribution," beginning on page P-19, that it is convenient, for all the reasons discussed in that section, to use transformed price returns. We will do that, but the approach will look different from the discussion in that section. Specifically, in that section, the price returns represented prices from successive periods. The section "Geometric Brownian Motion" described paths that result when these transformed returns stem from independent, uniform "innovations."[13] In fact, we are not going to make any such assumptions about how prices in 2006 relate to those in 2005 or 2007. We may have information that a large supply of natural gas is coming on-line in 2006, for example, so in a sense the 2006 product is distinct from those in 2005 and 2007. Instead, for each year's price, we represent its covariance with any other year using principal factor analysis, and the only innovation we are interested in is the one-step change between today and tomorrow. (Remember, we are simulating tomorrow's forward curve.) If we do this for each forward year, we get a new curve.

We start by taking the logarithm of the price for each year j, j=1 to m, of today's forward curve. The prices and transformed prices appear in equations (1). Denote this transformed price by $\ln(\boldsymbol{p}_{j,0})$. Denote the corresponding transformed price for *tomorrow* by $\ln(\boldsymbol{p}_{j,1})$. The innovations ε_j are drawn from the distribution of the transformed returns $\ln(\boldsymbol{p}_{j,t+1}/\boldsymbol{p}_{j,t})$ obtained from historical data for that forward year. A given draw then gives us the means of estimating a possible price for tomorrow's forward curve:

$$\begin{aligned} &\ln(p_{j,1}) = \ln(p_{j,0}) + \varepsilon_j, \text{ where} \qquad (9) \\ &\varepsilon_j \approx \ln(p_{j,t+1}/p_{j,t}) \end{aligned}$$

The second line merely says that the innovations are distributed like the transformed daily price returns for year j.

The previous section provides a technique for simulating this vector of innovations. We can construct the covariance matrix from historical data, find the Cholesky decomposition, and use a higher-dimensional version of equation (8) to produce the samples. If natural gas prices behave as many commodity prices do, the innovations will

[13] By innovation, we mean small, random shocks. These are generated by drawing a value from a random variable.

be roughly normally distributed, so the vector η in equation (8) will be drawn from a normal distribution.

When practitioners applied these techniques to very large vectors, however, they discovered that these calculations could become burdensome. The computations increase roughly as the number of non-zero elements, $(m^2+m)/2$, in the Cholesky factor. They discovered that, by using principal factor analysis, they could substantially reduce that computational burden, especially when the entries in the vector of prices were strongly correlated.

Principle factor analysis is based on the fact that any symmetric matrix, such as any covariance matrix, has a "spectral decomposition"

$$\mathbf{A} = \lambda_1 \mathbf{e}_1 \mathbf{e'}_1 + \lambda_2 \mathbf{e}_2 \mathbf{e'}_2 + \cdots + \lambda_m \mathbf{e}_m \mathbf{e'}_m \qquad (10)$$

where

$\mathbf{A}$ is a $(m \times m)$ symmetric matrix

λ_i is the ith eigenvalue

$\mathbf{e}_i$ is the ith normalized eigenvector $(m \times 1)$

$\mathbf{e'}_i$ is the transpose of $\mathbf{e}_i$

If there are strong correlations among entries of the random vector, several of the eigenvalues tend to be much larger than the rest. The eigenvectors are principal patterns of correlated variation in entries and these give rise to the paths to which this section has referred. If the terms in equation (10) are sorted with respect to magnitude of their eigenvalues (they will all be positive), we can represent the covariance matrix as the sum of two matrices, one associated with the first *k* dominant eigenvalues and the second associated with the remaining eigenvalues. Because these two terms are also symmetric matrices, they both have Cholesky terms:

$$\Sigma = \mathrm{LL'} + \mathrm{SS'}$$

where

L is $m \times k$

S is $m \times (m-k)$

The S matrix should be nearly diagonal and if we replace it by a diagonal matrix, we obtain an equation for creating the innovations that corresponds to equation (8):

$$\mathbf{X} - \boldsymbol{\mu} = \mathbf{L}\mathrm{f} + \boldsymbol{\varepsilon} \qquad (11)$$

where

$\mathbf{X}$ is the k - vector of random variables

$\boldsymbol{\mu}$ is their k - vector of means

$\mathbf{L}$ is the $(k \times m)$

matrix of k - eigenvectors

f is an m - vector of independent random variables

$\boldsymbol{\varepsilon}$ is a k - vector of independent random "specific factors"

The entries in the m-vector f may be taken to be distributed N(0,1); the specific factors may also be taken as independent, normally distributed with mean zero, but the variance of each is determined by the residual variance necessary to match that of **X-μ**. Efficiencies arise when *m* is much less than *k*.

In Figure P-14, the possible forward curve is the weighted sum of the following three eigenvectors:

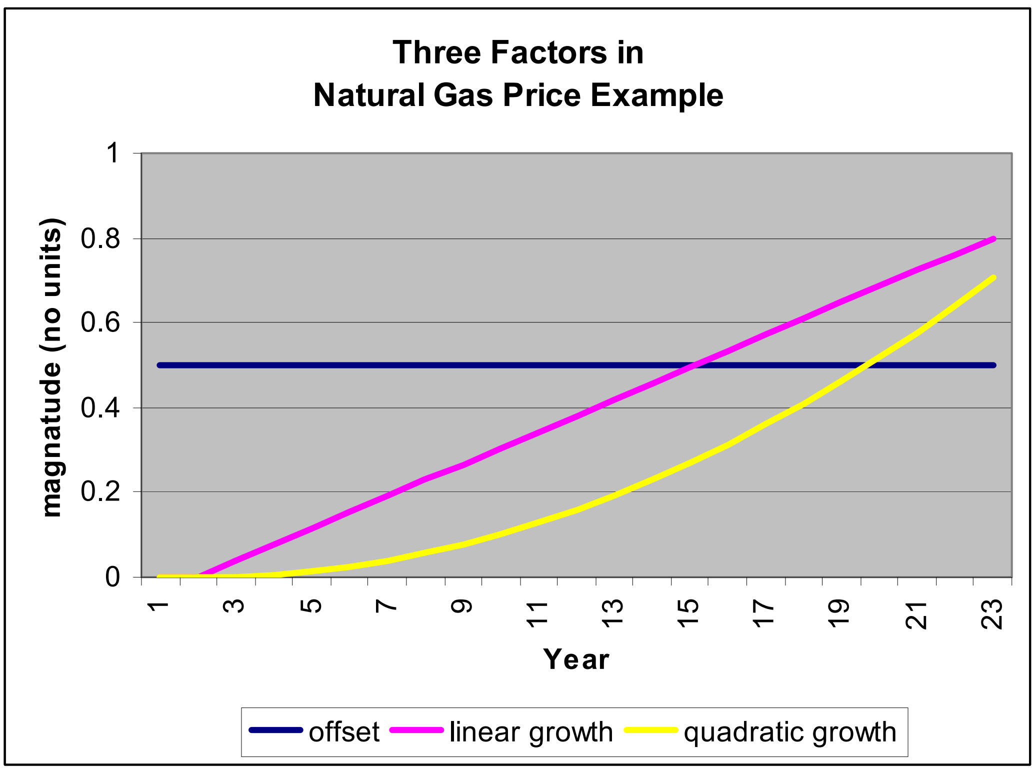

Figure P-15: Constituent Eigenvectors

For the possible (dotted line) forward curve in Figure P-14, the offset, linear growth, and quadratic growth eigenvectors, of "sub-paths," are weighted by 0.00, -0.75, and 0.90, respectively. These sub-paths are then added to the transformed returns, as in equation

(9), and transformed back to prices using the standard exponential transformation described on page P-21. Figure P-16 illustrates the steps.

Slightly different weightings provide dramatically different paths. For example, the

P0	ln(P0)	offset	linear	quadratic	sum (e)	ln(P0)+e	P1= exp(ln(P0)+e
4.62	1.53	0.00	0.00	0.00	0.00	1.53	4.62
5.45	1.70	0.00	0.00	0.00	0.00	1.70	5.47
5.30	1.67	0.00	-0.03	0.00	-0.03	1.64	5.16
5.01	1.61	0.00	-0.06	0.01	-0.05	1.56	4.76
4.74	1.56	0.00	-0.09	0.01	-0.08	1.48	4.39
4.48	1.50	0.00	-0.11	0.02	-0.09	1.41	4.10
4.23	1.44	0.00	-0.14	0.04	-0.10	1.34	3.82
4.00	1.39	0.00	-0.17	0.05	-0.12	1.27	3.56
3.96	1.38	0.00	-0.20	0.07	-0.13	1.25	3.49
3.92	1.37	0.00	-0.23	0.09	-0.14	1.23	3.42
3.88	1.36	0.00	-0.26	0.12	-0.14	1.22	3.39
3.84	1.35	0.00	-0.29	0.14	-0.15	1.20	3.32
3.80	1.34	0.00	-0.31	0.17	-0.14	1.20	3.32
3.82	1.34	0.00	-0.34	0.21	-0.13	1.21	3.35
3.84	1.35	0.00	-0.37	0.24	-0.13	1.22	3.39
3.86	1.35	0.00	-0.40	0.28	-0.12	1.23	3.42
3.88	1.36	0.00	-0.43	0.32	-0.11	1.25	3.49
3.90	1.36	0.00	-0.46	0.37	-0.09	1.27	3.56
3.92	1.37	0.00	-0.49	0.42	-0.07	1.30	3.67
3.94	1.37	0.00	-0.51	0.47	-0.04	1.33	3.78
3.96	1.38	0.00	-0.54	0.52	-0.02	1.36	3.90
3.98	1.38	0.00	-0.57	0.58	0.01	1.39	4.01
4.00	1.39	0.00	-0.60	0.64	0.04	1.43	4.18

Figure P-16: Steps in the Calculation

weighting (0,1.25, -1.2) gives rise to the path illustrated in Figure P-17. The weighting (-1.4, 1.25, -1.2) generates the curve in Figure P-18.

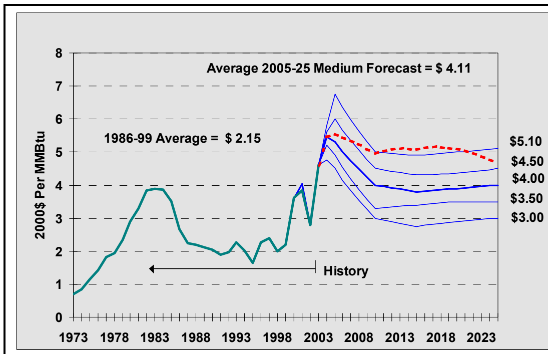

Figure P-17: Path with (0,1.25, -1.2) Weighting

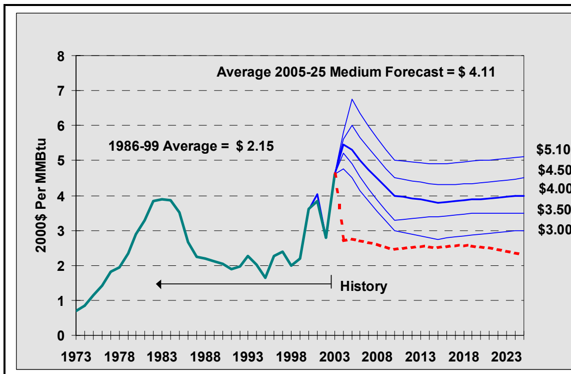

Figure P-18: Path with (-1.4, 1.25, -1.2) Weighting

Returning to the original challenge of creating new paths for future prices and loads, it would be natural to attempt construction of future paths based on historical data. It turns out, though, that those kinds of patterns generally did not garner credibility with experts. They usually failed to capture the experts' scale of uncertainty. Effectively, the curve weighting and parameters were calibrated to the experts' expectations. This is in keeping with the spirit of strategic decision analysis articulated on page P-16, however, which recognizes the subjective nature of characterizing complex and unpredictable behaviors.

Three factors, like those illustrated in Figure P-15, appear to be sufficient to capture the kind of underlying path behaviors that experts wished to see. Of course, these paths do not suffice to produce all of the kinds of necessary behavior. There is short-term (period) variation, as we might expect to see with weather differences. Prices and requirements possess short-term correlations, within the modeling period, and these require attention. There are also jumps that reflect excursions from long-term supply and demand equilibrium or other economic disruption. The construction of the jumps is the subject of the next section.

The example of natural gas price simulation in the workbook L24DW02-f06-P.xls provides a good example of how the regional model treats the factors. This takes place in rows {56 to 62}. As shown in Figure P-19, the value for the period 7 in {X62} is a sum of three products. The first product is the weighting for the linear growth {$S$56}, times the random number in {$R$56}, times the value of the factor in {$W$57}. The random

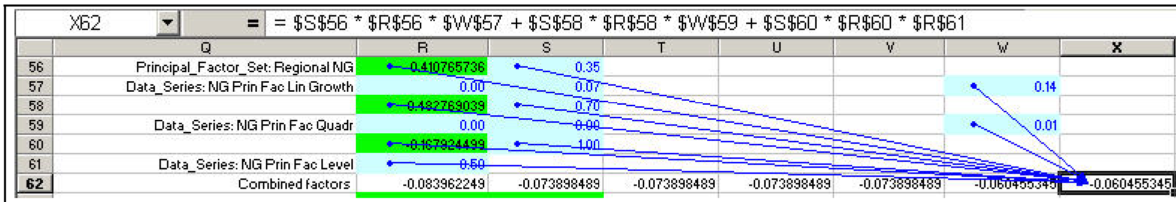

Figure P-19: Adding the Principal Factors

number plays the role of an entry of η in equation (8) or of **f** in equation (11). The distribution of the random number will depend on the simulated uncertainty. The value of the linear factor does *not* increase smoothly over the 80 periods, from 0.0 to 1.33. Instead, because the Olivia model[14] that created this workbook used annual values, the values only change once each four columns, and the logic points back to the last data value for the factor.

The remaining two terms in the sum {X62} add the quadratic and offset factors. Because the offset factor does not change across periods, the formulas in row {62} all point to the offset factor value in cell {R61}.

This is not the last step in creating the behavior for natural gas price. Other influences, such as jumps, add to the combined factor, and the worksheet applies the necessary

[14] Olivia is a Council application that creates Excel worksheet portfolio models. Appendix L describes Olivia.

inverse transformation to the sum. The next two sections discuss specific factors and jumps. The subsequent section describes the stochastic adjustment, and the section following that one shows the final inverse transformation.

Specific Factors

Specific Factors arise in equation (11) as a means to capturing variance not accounted for by the principal factors. They are "specific" in the sense that they describe only the remaining variance for a stochastic vector's entries.

In the regional model, specific factors are typically describing seasonal variation, which can be greater at certain times of the year. For example, loads tend to have greater uncertainty during the winter and summer, so the model adds independent variance to those seasons. Figure P-20 shows the crystal ball dialog box that specifies the distribution of the random variable in cell {AQ 124}. This is a normal distribution with mean zero and a standard deviation of five percent. As described in the section

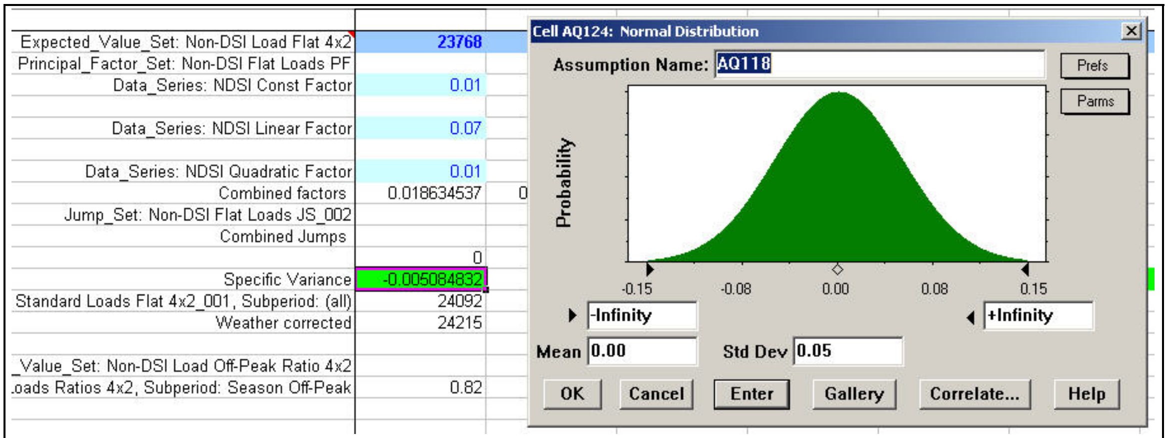

Figure P-20: Assumption Cell

"lognormal distribution," this small standard deviation will correspond to roughly a five percent change in the standard deviation of quarterly loads.

Jumps

Excursions occur in prices and loads for several reasons, including a disequilibrium in long-term supply and demand. Gas and electricity prices, as we have seen in the last few years, can depart significantly from their equilibrium values when capacity shortages occur. It typically takes a year or two for new capacity to come on-line. Load excursions will occur due to business cycles or large economic displacements. It is important to have this kind of behavior in the regional model because large and sudden changes, which can last a significant time, are key sources of uncertainty and risk. These changes, moreover, may stem from activities and prices outside the region and may therefore be uncorrelated with local events.

One of the shortcomings of the principal factor approach to simulating price paths is that it does not accommodate very well excursions that begin at random times and last for a random number of periods. Rather that forcing the principle factor metaphor, the regional model represents these excursions with a different, simpler technique.

In the regional model, jumps can begin at random times and have random magnitude and duration. There is logic to model the "recovery" from excursions and to constrain when jumps can take place.

Figure P-21, which shows the wholesale electricity price[15] in row {102} of our sample workbook L24DW02-f06.xls, illustrates a typical jump with recovery. The first jump, illustrated by the heavy line, and the subsequent recovery have an obvious impact on the electricity price, illustrated by the light line. In addition, a second jump begins in the 79th period and lasts the remaining two periods of the study.

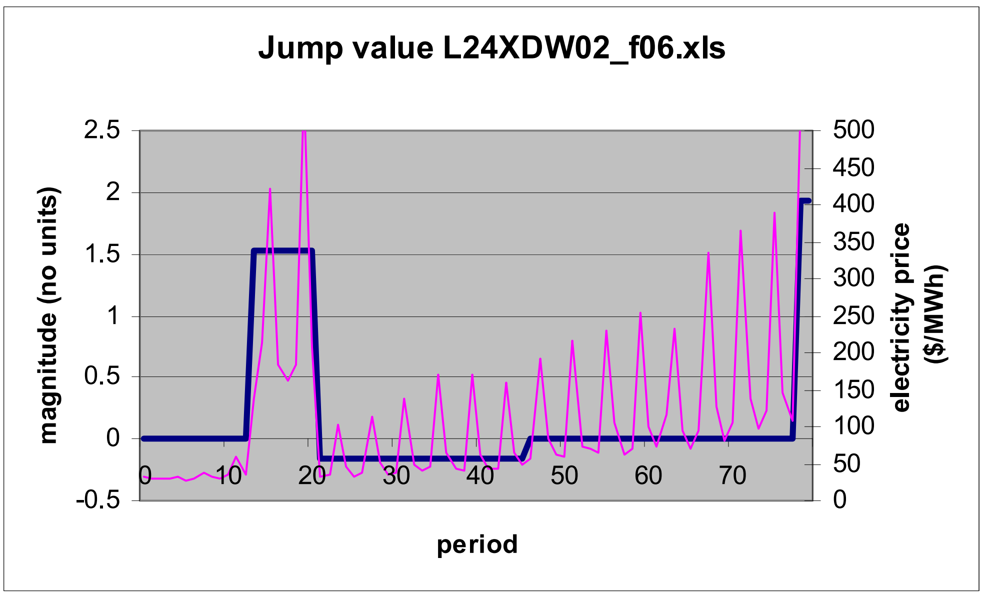

Figure P-21: Typical Jump with Recovery

The worksheet logic that produces the jump pattern appears in rows {99} through {102}. In principle, there can be as many jumps as the user desires. For this two-jump system, we first have the following Crystal Ball assumption cell values:

	R	S	T
99	13.74923	1.534261	10.59427
100	32.05555	1.935131	8.386783

[15] This is the flat market price before any resource response. Resource responsive price modeling is the subject of Appendix L. "Flat" market prices are average prices, where the average is with respect to on- and off-peak hours in whatever period is under discussion.

where R, S, and T are the wait, size, and duration of the jump, respectively. The values for the wait and duration of a jump specify the number periods that must pass before a jump can begin and end, respectively. For proportional jumps, the model ignores the last parameter, because the size of the jump determines its duration. This particular example uses proportional jumps.

Then the formulas in the row {101} calculate intermediate values, which specify the periods in which events occur.

R101	=$R$99	wait_1	start time of jump 1
S101	=R101+ IF($S$99= 0,0,12/$S$99)	wait_1+ 12/size_1	end time of jump 1
T101	=$S$99	size_1	size_log xfr jump 1
U101	=S101	end time of jump 1	start time of recovery 1
V101	=U101+ S101*EXP(T101)	end time of jump 1 + duration recovery 1	end time of recovery 1
W101	=-T101/10	-size_1/10	size_log xfr recovery 1
X101	=V101+ $R$100	end time of recovery 1+ wait_2	start time of jump 2
Y101	=X101+ IF($S$100= 0,0,12/$S$100)	wait_2 + 12/size_2	end time of jump 2
Z101	=$S$100	size_2	size_log xfr jump 2
AA101	=Y101	end time of jump 2	start time of recovery 2
AB101	=AA101+ Y101*EXP(Z101)	end time of jump 2 + duration recovery 2	end time of recovery 2
AC101	=-Z101/10	-size_2/10	size_log xfr recovery 2

Figure P-22: Intermediate Jump Calculations

The first six columns, {R101 through W101}, calculate parameters for the first jump; those in the next six columns pertain to the second jump. Note that the formulas for the first jump are almost identical to those for the second. If the user specified additional jumps, there would be six additional columns in that row for each additional jump.

The size and duration of the jump recovery are proportional to the inverse of the size of the jump. The scaling factors of 10 and 12 in columns {S, W, Y, and AC} control the sizes. The size of these factors produce "realistic" behavior, i.e., behavior that conformed to expectations about the future. Originally, the size of the duration and jump assured that the price or load adjustment, after appropriate inverse transformation

$$y_t \overset{e}{\mapsto} p_t / p_{t-1}$$

would average to 1.0. The intent was to create prices that averaged out to the long-term equilibrium value over time. This approach, however, produced recoveries that were much too large and lasted too long. The Council therefore abandoned recoveries of that size and duration. Part of the rationale in moving away from an adjustment that averaged to 1.0 was some disbelief that there was justification for prices returning to a fixed long-term price. Equilibrium prices, after all, can change as underlying economics change.

Row {102} interprets the values in row {101} based on the period number in row {46} and produces the final jumps:

```
R102 = IF(AND(R$46>$R101,R$46<=$S101),$T101,0)+          jump_1
       IF(AND(R$46>$U101,R$46<=$V101),$W101,0)+          recovery_1
       IF(AND(R$46>$X101,R$46<=$Y101),$Z101,0)+          jump_2
       IF(AND(R$46>$AA101,R$46<=$AB101),$AC101,0)        recovery_2

S102 identical, except S$46 instead of R$46
T102 identical, except T$46 instead of R$46
```

Row {102} contains the values that must then undergo inverse transformation. This final transformation is the subject of the next section.

CO_2 and emission taxes exhibit a special kind of jump behavior not shared by loads and prices. There is only one jump, but its value can change in particular periods. When the Council queried experts about the likelihood of carbon tax legislation, the experts agreed that any changes would probably occur with a change in the federal administration. Therefore, emission taxes can arise only in the year of a presidential campaign (2008, 2012, etc.). These are step functions of uncertain size and timing, otherwise. Any jump remains in place through the end of the study. The section on CO_2 tax uncertainty further describes this behavior.

Stochastic Adjustment

Prices in the model derive from the Council's assumptions for long-term equilibrium prices.[16] For reasons discussed in Chapter 6, these equilibrium prices can be associated with the median price because there is equal probability of being above and below the median price. Some users may prefer, however, for the long-term equilibrium prices to match the price distribution's ***mean***. Because prices in the regional model use a lognormal distribution, however, the mean price is *higher* than the median price.

To accommodate this situation, the model can apply a "stochastic adjustment" to the benchmark price. This adjustment, a number between zero and one, is chosen so that the distributions mean price matches the benchmark price. An example of a stochastic adjustment for on peak wholesale electricity market prices appears in the second row of Figure P-23.

	Series: Market Prices Independent Term_005				
Expected_Value_Set: Market Exp Price On-Peak 4x2	32.29	33.04	32.99	32.33	32.66
Stochastic_Adjust_Set: Stoch Adj On-Peak 4x2	0.87	0.73	0.78	0.76	0.85
Principal_Factor_Set: Reg Mkt Prc	-0.02037443	1.00			
Data_Series: Mkt Prin Fac Level	0.50				
	0.007267999	1.00			
Data_Series: Mkt Prin Fac Lin Growth	0.00	0.07			
Combined factors	-0.010187215	-0.009678455	-0.009678455	-0.009678455	-0.009678455
Jump_Set: Elec Mkt_002	8.770426174	0.072691876	8.899814829		
	16.07130502	0.100080134	11.46780741		
Combined Jumps	8.770426174	173.850772	0.072691876	173.850772	220.5905142
	0	0	0	0	0

Figure P-23: Stochastic Adjustment

[16] Because the median and the mean both described the final distribution of prices after any adjustment, we refer to the starting place as the "benchmark price." The benchmark price is typically the long-term equilibrium price.

Each period typically requires a separate stochastic adjustment. Appendix L describes a utility, the macro subTarget, which automates the process for finding values for the stochastic adjustment.

Combinations of Principal Factors, Specific Factors, and Jumps

The preceding sections describe how the model represents stochastic behavior using combinations of principal factors, specific factors, and jumps. It is easiest, however, to model these elements with simple symmetric or unbounded distributions. The inverse lognormal transformation then guarantees physical values that have positive value and behavior that is more realistic.

In the example of natural gas price from the sample workbook, the model combines these influences in row {68}. For example, the formula in column {R} is

= R53*R54*EXP(R66+R62+R67)

The first two terms are the benchmark price and the stochastic adjustment factor, respectively. The remaining three terms R66, R62, and R67, are the jump, principal factor, and specific factor contributions, which must be inverse transformed according to equation (2). Because the inverse transformation produces the ratio of the new value to today's value or, in the case of strategic uncertainty, the value of the benchmark, the worksheet must multiply it by the benchmark price (modified by any stochastic adjustment) to obtain the new price.

This concludes the discussion of the regional model's representation of stochastic processes. The appendix now turns to how the model applies these principles to the specific sources of uncertainty that are of interest.

Load

Electricity requirement, or load, in the regional model has characteristics that depend on the timescale. On an hourly basis, loads have distinct on- and off-peak variation. Hourly electricity prices typically move with this load. However, the period duration in the regional model is three months. When we consider load requirements averaged over three months there are

- strong chronological correlation,
- seasonal shapes,
- excursions due to changing economic circumstances, and
- long-term elasticity to electricity prices.

The long-term correlation with electricity prices differs in magnitude and direction from the short-term correlation. That is, loads generally correlate positively to electricity prices in the short term but negatively in the long term.

This appendix has described the techniques the regional model uses for capturing this broad spectrum of behaviors. This section details the specific formulas and data that implement those techniques.

Electric load serves a number of purposes in the regional model. Its main role is its contribution to energy balance, and thus costs, in the regional model. Two other roles that it serves, however, are as a term in the reserve margin calculation and as an influence on medium-term electricity prices.

Appendix L describes how the regional model uses energy requirements to determine energy balance, costs, and reserve margin. This Appendix P will trace back the logic and data from the point where Appendix L begins the discussion. This will be a "bottoms-up" description. The description proceeds from the final values used in Appendix L to the constituent components from which they are constructed. Because the influence of load on medium-term electricity prices is an issue of modeling uncertainty, either that of load or of electricity price, that entire discussion appears in this appendix, in the section "Electricity Price."

Energy Balance and Cost

The discussion of energy load in Appendix L begins with the average megawatts for the period, on peak and off peak. The specific worksheet cells in the sample worksheet that provide on-peak and off-peak load in column {AQ} are {AQ 183} and {AQ 236}, respectively. The formulas in these two cells are similar. The on-peak calculation in {AQ 183} is

=AQ$125*AQ$133

Tracing back from these cells, the reader will find that AQ$125 is the period estimate for the flat load in that period. (By flat load, we mean the average load across all – on peak and off-peak – hours.) The value in cell {AQ 133} is a constant factor for converting monthly flat average megawatts to average megawatts over the on peak hours.

The source of these conversion constants is reference [**3**]. The process used to arrive at them is as follows:

1. From Northwest Power Pool energy and peak loads for 2000 through 2002, calculate a monthly load factor;

2. Estimate the on-peak and off-peak energy using the number of corresponding hours in each month and the simple load duration curve model illustrated in Figure P-24;

3. Estimate the monthly and quarterly multiplication factors; and

4. Recognize that the quarterly factors, illustrated in Figure P-25, change little and are effectively constants.

The preceding section, Stochastic Process Theory, describes how the model represents uncertainties with principal factors, jumps, and specific factors. As shown in Figure P-26, the period estimate for flat load that appears in cell AQ 125 is the product of the benchmark level load requirement {AQ$113} times the inverse transformation (equation 2) for specific variance {AQ$124}, jump {AQ$123}, and combined factor terms {AQ$120}:

=AQ$113*EXP(AQ$123+AQ$120+AQ$124)   (12)

The specific factor contribution ({AQ124}) is nonzero, roughly five percent, only for winter and summer seasons. Council staff [4] concluded that this was an appropriate amount of seasonal variation of loads due to weather uncertainty.

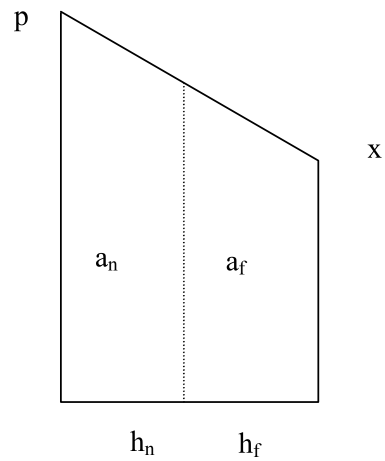

Figure P-24: Load Duration Curve to Estimate On- and Off-Peak Energies

$$\text{area or energy } a = a_n + a_f$$
$$\text{hours } h = h_n + h_f$$
$$\text{min load } x = 2(a/h) - p$$
$$\text{average on-peak energy } a_n/h_n = p + (h_n/h)*(x-p)/2$$
$$\text{average off-peak energy } a_f/h_f = (a-a_n)/h_f$$

The jump contribution in cell {AQ 123} represents longer-term excursions in load requirements due to a host of influences, including general economic activity. Because business cycles tend to last several years, the regional model uses only a single jump. The logic for the jump is a variation of the example that the previous section illustrated. In particular, the duration of the jump is specified rather than being a function of the size of the jump, and the recovery is specialized.

The wait, size, and duration for jumps are all random variables. The specification for the wait, size, and duration appear in Figure P-27.

Spring		1.14	0.82
Summer		1.10	0.87
Fall		1.14	0.82
Winter		1.18	0.78
	average	1.14	0.82

Figure P-25: Quarterly Multiplication Factors (Unitless)

AQ125	=AQ$113*EXP(AQ$123+AQ$120+AQ$124)		
	Q	AQ	AR
112			
113	Expected_Value_Set: Non-DSI Load Flat 4x2	23768	20287
114	Principal_Factor_Set: Non-DSI Flat Loads PF		
115	Data_Series: NDSI Const Factor	0.01	
116			
117	Data_Series: NDSI Linear Factor	0.07	
118			
119	Data_Series: NDSI Quadratic Factor	0.01	
120	Combined factors	0.018634537	0.018634537
121	Jump_Set: Non-DSI Flat Loads JS_002		
122	Combined Jumps		
123		0	0
124	Specific Variance	0.005084832	0
125	Standard Loads Flat 4x2_001, Subperiod: (all)	24092	20669
126	Weather corrected	24215	20669

Figure P-26: Average Megawatt Requirements Calculation

The jump recovery is such that, after transformation, the jump area equals the recovery area. The size of the jump before transformation equals the size of the recovery before transformation. This is an arbitrary choice, to make the calculation simple. To make the areas after transformation the same, the duration of the recovery is a function of the jump duration and size. To make the areas the same, we have D_1(exp(J)-1)=D_2(1-exp(-J)), where D_1 is the duration of the original jump and D_2 is the duration of the "recovery" jump. This gives us D_2=D_1*exp(J). Having equal areas means the load excursions average out over a sufficiently long (D_1+D_2) period.

The combined principal factors have the weightings, distributions, and eigenvalues illustrated in Figure P-29.

Random Variables

	Type	Cell	Distribution	Parameters	
Jump 1	wait	{{R121}}	**uniform**	min 0	max 85
	size	{{S121}}	**uniform**	min -0.10	max 0.80
	duration	{{T121}}	**uniform**	min 8	max 20
Principal Factors	offset	{{R114}}	**normal**	mean 0	stdev 1
	linear	{{R116}}	**normal**	mean 0	stdev 1
	quadratic	{{R118}}	**normal**	mean 0	stdev 1
Specific Variance		{{row 124}}	**normal**	mean 0	stdev 0.05

Figure P-27: Assumption Cell Values for Load

The Council selected these to provide realistic behavior [**4**]. The validation for this behavior is the topic of this section's "Comparison with the Council's Load Forecast."

Finally, the baseline load forecast {row 113} corresponds to the Council's weather-adjusted, non-DSI load forecast [**5**] and reflects the following assumptions:

- Nine percent losses for distribution and transmission
- Existing conservation through hydro-year 2003

	R	S	T
121	48.54747	-0.03909	16.17535

			interpretation
R122	=$R$121	wait_1	start time of jump 1
S122	=R122+ $T$121	wait_1+ duration_1	end time of jump 1
T122	=$S$121	size_1	size_log xfr jump 1
U122	=S122	end time of jump 1	start time of recovery 1
V122	=U122+ S122*EXP(T122)	end time of jump 1 + duration recovery 1	end time of recovery 1
W122	=-T122	-size_1	size_log xfr recovery 1

R123 = IF(AND(R$46>$R122,R$46<=$S122),$T122,0)+ jump_1
IF(AND(R$46>$U122,R$46<=$V122),$W122,0) recovery_1

S123 identical, except S$46 instead of R$46
T123 identical, except T$46 instead of R$46

Figure P-28: Jump Data and Formulas for Load

- Frozen efficiency for hydro year 2004 and beyond [17]
- Monthly distribution of annual energies, and the aggregation of those monthly energies into quarterly energies

This baseline forecast serves as the median of the distribution of energy requirements. The model has all future conservation in the conservation supply curves described in Appendix L. The only exceptions are conservation implemented before 2003 and conservation due to building codes and appliance standards implemented before 2003.

Energy Reserve Margin

Appendix L describes how the model uses weather-adjusted energy load requirement in each period to determine the energy reserve margin. The energy reserve margin plays a prominent role in the decision criterion to proceed with construction of new power plants.

The load estimate in cell {AP289} is the hydro year's average, weather-corrected non-DSI load (the range {AL126: AO126}), plus the DSI load in the final period.

=-AVERAGE(AL126:AO126)-AO327

Pricipal Factors

	offset	linear	quadratic
	Weight		
	0.000	0.300	0.051
Dec of Cal			
Year	Value		
2003	0.01	0.01	0.00
2004	0.01	0.02	0.00
2005	0.01	0.03	0.00
2006	0.01	0.04	0.00
2007	0.01	0.05	0.01
2008	0.01	0.06	0.01
2009	0.01	0.07	0.01
2010	0.01	0.08	0.02
2011	0.01	0.09	0.02
2012	0.01	0.10	0.03
2013	0.01	0.11	0.04
2014	0.01	0.12	0.04
2015	0.01	0.13	0.05
2016	0.01	0.14	0.06
2017	0.01	0.15	0.07
2018	0.01	0.16	0.08
2019	0.01	0.17	0.09
2020	0.01	0.18	0.10
2021	0.01	0.19	0.11

Figure P-29: Principal Factors for Load

[17] The frozen efficiency load forecasts assume no new conservation of any kind, although it does incorporate any *prior* conservation and the effect of *existing* codes and standards on future requirements. Instead, conservation supply curves in the model reflect future conservation measures and new codes and standards.

The model's weather corrected load is simply the load, less the stochastic part that represents weather variation in the winter and summer. Specifically, if the user examines cell {AO 126}, the last cell in the average computed in the previous equation, they will find formula

=AO$113*EXP(AO$123+AO$120)

This of course matches to equation (12), less the term that corresponds to specific variance for weather.

Hourly Behavior

The regional model captures hourly prices and requirements through descriptive statistics. In particular, the transformed hourly variation in load given by equation (3) and its correlation with hourly electricity price determine revenues to meet load. Appendix L describes the calculation in its discussion of Single-Period load behavior. The intra-period hourly load variation is 25 percent, as specified in cell {R 185}. The hourly correlation with other variables appears in this section's, "Hourly Correlation" discussion.

Comparison with the Council's Load Forecast

Statute requires that the Council's Northwest Regional Conservation and Electric Power Plan have a 20-year forecast of electricity demand.[18] This forecast of electricity demand serves as the basis for other, alternative forecasts that are necessary for specific purposes, such as a source of input data for the Aurora™ model. The alternative forecasts use assumptions that differ from those for the primary forecast. For example, an alternative forecast may use different assumptions about energy losses or about the representation of conservation. To compare the regional model's load forecast to the primary forecast, this section determines what adjustments to the primary forecast would make the two forecasts comparable. The section then compares the modified primary forecast and the loads from regional model futures.

The regional model uses a non-DSI forecast. The model simulates the behavior of DSI load separately, using electricity and aluminum prices in the model. (See Appendix L for a description of DSI modeling in that appendix's "Multiple Period" section of Principles.) The non-DSI load forecast appearing in the plan (Appendix A) is of sales (MWa) by calendar year, including conservation expected to arise from a forecast of retail electricity rates but excluding conservation due to codes and standards implemented since the Council's 4th Plan. The basis of electricity rate forecast is an earlier calculation of long-term equilibrium wholesale prices. The annual loads appear in Table P-1, which details the values in Appendix A, Table A-2.

[18] Public Law 96-501, Sec. 4(e)(3)(D)

YEAR	Non-DSI Sales (Price Effects) Low	Medlo	Medium	Medhi	High
2004			18072		
2005	17191	17824	18433	19020	20221
2006	17200	17955	18663	19360	20727
2007	17214	18098	18906	19721	21257
2008	17228	18239	19145	20093	21814
2009	17257	18398	19405	20479	22397
2010	17297	18570	19688	20879	23007
2011	17320	18729	19959	21275	23598
2012	17353	18906	20251	21696	24214
2013	17366	19067	20521	22106	24843
2014	17430	19274	20830	22547	25501
2015	17489	19482	21147	23000	26187
2016	17522	19672	21456	23449	26906
2017	17554	19864	21770	23907	27645
2018	17586	20058	22089	24375	28407
2019	17619	20254	22413	24853	29190
2020	17652	20453	22742	25341	29997
2021	17686	20653	23076	25839	30827
2022	17719	20855	23415	26347	31681
2023	17753	21059	23760	26866	32560
2024	17787	21265	24109	27396	33466
2025	17822	21474	24464	27937	34397

Table P-1: Council's Non-DSI Calendar-Year Sales Forecast

Some background about the Council’s load-forecasting methods will be helpful to following the development of forecast adjustments. Electricity prices, building codes, and appliance standards determine the implementation rate for conservation and consequently, energy requirements. Because Council policy can affect codes and standards directly and electricity prices indirectly, it is useful to separate these influences.

One way to approach this decomposition is to start with a "frozen efficiency" load forecast. The frozen efficiency load forecast reflects the amount of energy requirement that would arise only from current appliance standards and codes. Next, one would attempt to estimate how much conservation would arise in the future from the price effect of retail electricity rates. That is, ratepayers should pursue some conservation because it costs less than the electricity it displaces. The Council refers to load forecast net of this reduction as the “price-effects” forecast.

The Council has demonstrated, however, that additional benefit accrues to ratepayers from conservation beyond that which ratepayers would pursue to offset anticipated electricity purchases. Specifically, additional conservation can reduce fuel cost and defer the utility's capacity expansion. Electric power rates may go up or down because of this conservation, but this additional conservation would minimize ratepayers’ total power costs. To induce this additional conservation, however, the region typically must pursue additional codes and standards or other conservation measures. The Council refers to the forecast that arises by virtue of this additional conservation as a "sales" forecast, that is, the actual sale of electricity to consumers after the effects of codes and standards, energy conservation, utility program savings, and consumers’ own response to prices.

The regional model, on the other hand, represents conservation using supply curves, which include new utility programs, appliance codes and standards, and price effects. Consequently, the regional model needs the frozen efficiency load forecast. If price effects or program savings were subtracted from the load, the model would be double counting their effect.

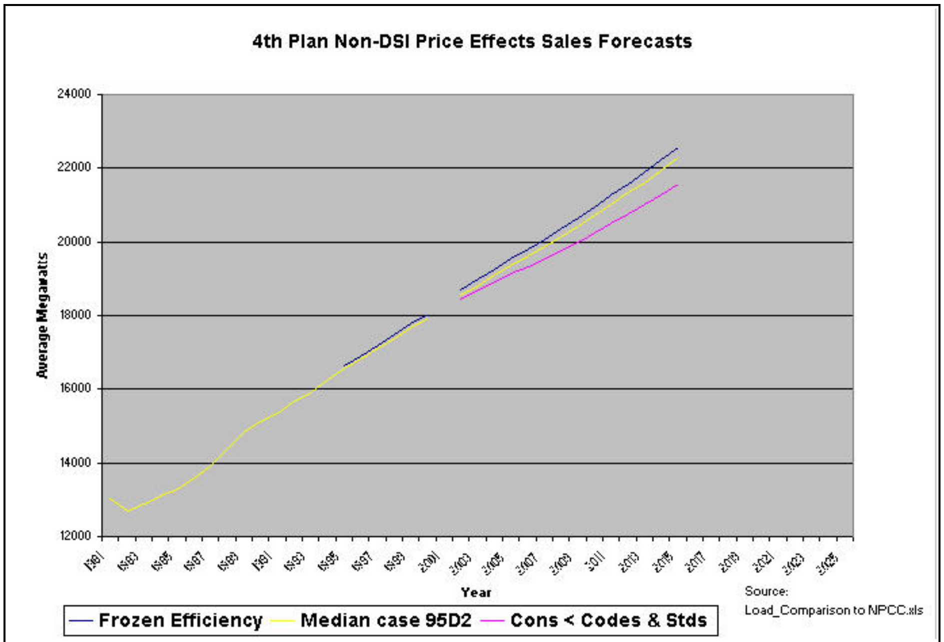

Figure P-30: Comparison of 4th Plan Load Forecasts

As mentioned in Appendix A, the load forecast of the Fifth Plan builds directly on work of the Fourth Plan. Figure P-30 illustrates the relationship in the Fourth Plan between the

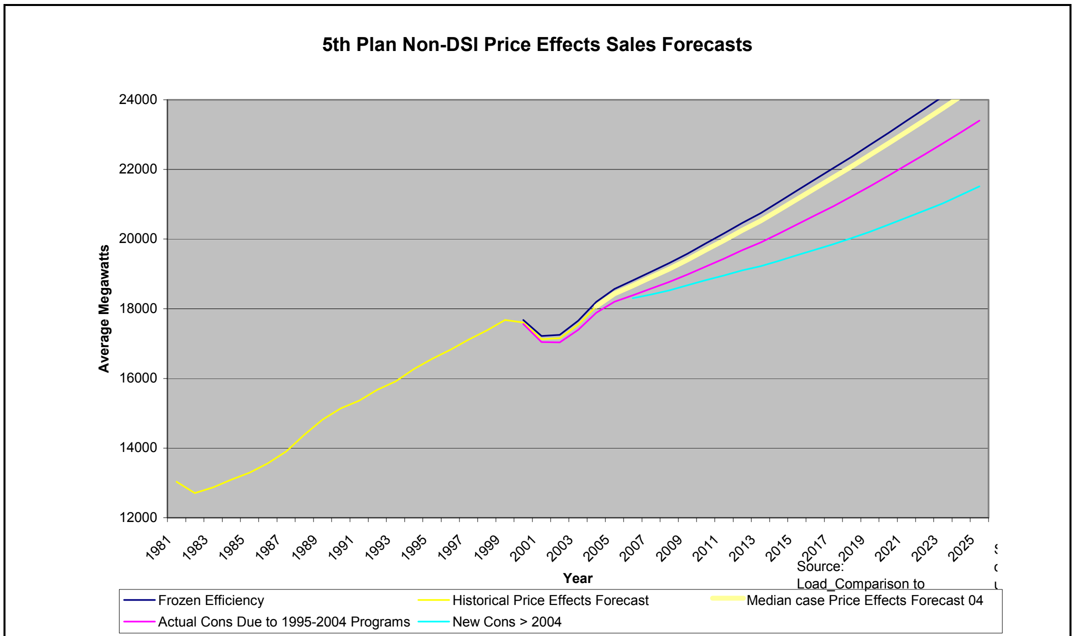

Figure P-31: Comparison of 5th Plan Load Forecasts

frozen-efficiency, price-effects, and sales load forecasts. To prepare the load forecast for the Fifth Power Plan, the Council used a revised price-effects forecast (Table A-2). The revised price-effects forecast builds on the price-effects forecast in the Fourth Plan, incorporating history over the last five years. In particular, the revised price-effects forecast does not reflect the conservation arising from codes and standards enacted since the Fourth Plan. Figure P-31, which has five loads, illustrates the resulting situation. Before 2004, it shows an estimate of the price-effects forecast due to actual history. This price-effects forecast is continued after 2004 as the "median case price-effects forecast 04". We know, however, that codes and standards since 1995 have in fact reduced loads, and ***this reduced forecast is our best estimate of where a price-effects forecast might wind up if the Council had updated the analysis for the Fifth Plan***. (The effect on loads of any new conservation, subsequent to the this ***Fifth Plan*** is captured by the line "new conservation > 2004.") Similarly, our best estimate of where the "frozen efficiency" load forecast would lie relative to the price-effects forecast comes from using the increment between the “price-effects” forecast and the "frozen efficiency" forecast in the last plan. In summary, therefore, the "frozen efficiency" load forecast used in the regional model starts with a revised price-effects forecast anchored in 1995 but reflecting economic history since then, reduces this forecast by the effect of conservation due to codes and standards implemented since the Fourth Plan, and adds the increment for frozen efficiency increment developed in the Fourth Plan. The frozen efficiency adders appear in Table P-2, and the estimated Code and Standards Savings since the Fourth Plan are in Table P-3.

	Frozen Efficiency Adders (From 95D4)				
YEAR	Low	Medlo	Medium	Medhi	High
2004	66	70	78	87	105
2005	60	64	74	86	109
2006	53	57	68	83	111
2007	48	53	66	83	116
2008	46	51	67	86	125
2009	46	51	69	91	137
2010	45	51	71	97	149
2011	46	52	74	103	163
2012	49	56	80	114	184
2013	57	67	92	131	210
2014	65	76	105	151	238
2015	72	85	116	167	265
2016	72	85	116	167	265
2017	72	85	116	167	265
2018	72	85	116	167	265
2019	72	85	116	167	265
2020	72	85	116	167	265
2021	72	85	116	167	265
2022	72	85	116	167	265
2023	72	85	116	167	265
2024	72	85	116	167	265
2025	72	85	116	167	265

Table P-2: Frozen Efficiency Adders

Finally, the revised forecast must capture losses due to distribution and transmission. An energy loss, which amounts to 9 percent, will increase the end use forecast measured at the customers’ electric power meters. The power plants in the regional model, of course, must meet both end use and losses of energy.

	Conservation Captured Since the 4th Plan		
YEAR	Residential	Commercial	Total
2004	174	14	187
2005	212	18	231
2006	254	23	276
2007	298	27	325
2008	343	31	373
2009	387	35	422
2010	433	39	472
2011	478	43	521
2012	524	47	571
2013	571	50	621
2014	618	54	672
2015	664	58	722
2016	711	62	773
2017	758	66	824
2018	794	70	863
2019	830	74	903
2020	852	78	929
2021	875	82	956
2022	898	86	984
2023	922	90	1012
2024	946	94	1040
2025	966	98	1064

Table P-3: Conservation Since the 4th Plan

The data presented in tables and graphs to this point reflect calendar year averages. Because the regional model uses hydro quarters, we must make the conversion to hydro year averages. The final formula for combining these effects is in Figure P-32. The table of the resulting values, by hydro year, appears in Table P-4.[19] The load forecast in this table serves as the basis for comparison between the Council's primary forecast and the regional model loads.

One subtlety of the formula in Figure P-32 is that we have implicitly assumed transmission and distribution losses are included in the frozen efficiency adders and the codes and standards savings. In any case, the adjustment for losses due to these effects is very small.

$$L_{HY,T} = \frac{8}{12}\left\{(1+\lambda)\cdot L_{CY,T} + PE_{CY,T} - C_{CY,T}\right\} + \frac{4}{12}\left\{(1+\lambda)\cdot L_{CY,T-1} + PE_{CY,T-1} - C_{CY,T-1}\right\}$$

where

$L_{HY,T}$ is load (MWa) for hydro year T

λ is loss factor (0.09)

$L_{CY,T}$ is load (MWa) for calendar year T

$PE_{CY,T}$ is price effect for calendar year T

$C_{CY,T}$ is conservation in calendar year T arising from programs implemented since the 4th Plan

Figure P-32: Calculation of Adjusted Primary Forecast

The regional model uses futures containing chronological loads that can vary quite

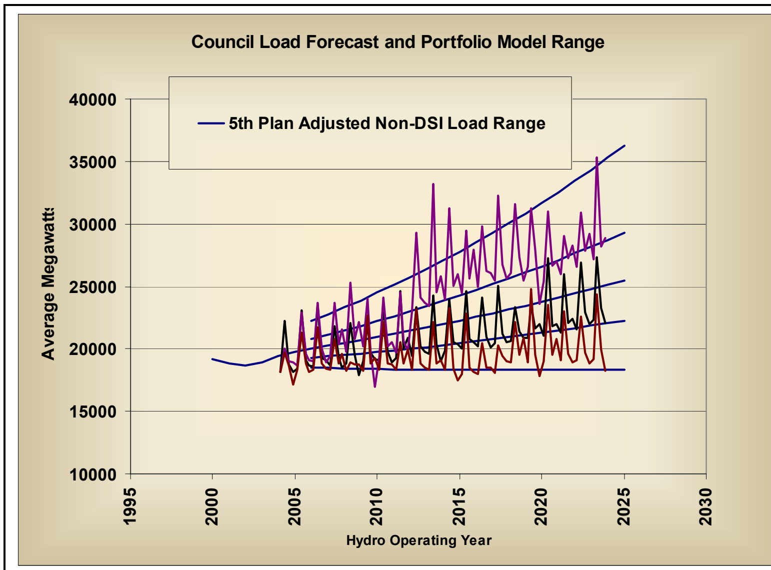

Figure P-33: Regional Model Loads from Futures

[19] The hydro year September 2006 through August 2007 is defined to be hydro year 2007.

dramatically. Jumps and excursions due to business cycles and weather are evident in individual futures, as illustrated in Figure P-33. This figure compares three randomly chosen futures from the 750 futures to the five load forecasts presented in Table P-4. Figure P-33 also has the disadvantage of comparing quarterly energy load values against annual averages. Even with only three futures, the figure is rather difficult to sort out. Two refinements to this graph that help make the data from the regional model more accessible are the presentation of the load data across all futures statistically and the averaging of the quarterly data into annual values.

HYDRO	Olivia Input Loads				
YEAR	Low	Medlo	Medium	Medhi	High
2004			19398		
2005			19800		
2006	18516	19298	20045	20778	22234
2007	18472	19393	20249	21112	22755
2008	18431	19492	20458	21462	23308
2009	18403	19604	20682	21828	23892
2010	18388	19733	20931	22211	24505
2011	18366	19858	21179	22597	25116
2012	18346	19994	21441	23002	25742
2013	18320	20129	21699	23414	26393
2014	18324	20293	21980	23847	27072
2015	18343	20473	22279	24298	27782
2016	18335	20634	22567	24739	28507
2017	18314	20787	22852	25180	29250
2018	18302	20951	23151	25639	30025
2019	18295	21121	23459	26113	30828
2020	18297	21303	23782	26608	31665
2021	18304	21491	24115	27118	32532
2022	18311	21681	24453	27638	33425
2023	18318	21872	24796	28170	34344
2024	18324	22065	25144	28713	35290
2025	18334	22264	25502	29271	36269

Table P-4: Hydro-Year Forecast for Regional Model

Figure P-35 compares the 0, 10th, 50th, 90th, and 100th percentiles against the forecast from Table P-4. The data falls somewhat outside of the jaws of the revised, primary forecast, as we would expect. The quarterly values have greater variation largely due to seasonal variation, and the Council believes there is some very small probability that annual average load will fall outside of the jaws.

Figure P-34 addresses the second problem, replacing quarterly values with annual averages. Now it is evident, for example, that the median forecast (50th percentile) lies directly on top of the adjusted Council "Medium" forecast.

In Figure P-34, there appears to be greater uncertainty associated with the futures in the early part of the study rather than near the end of the study. Indeed, if these forecasts are truly comparable we would expect the 0 percent and 100th percentiles to lie outside of the jaws.

One of the things going on here is the difference in assumption about electricity prices between the Council's primary forecast and the regional model. The Council's primary forecast, again, stems from a 1995 load forecast, which assumes much smaller variation in electricity price. The regional model sees electricity prices that are orders of magnitude larger, in particular. The regional portfolio model incorporates electricity price elasticity of loads. This elasticity will cause the variation in load excursions to diminish on average, especially in outlying years where greater electricity price variation occurs.

Another influence is the limited samples of futures. The regional model data presented in Figure P-34 are directly from the model's Monte Carlo simulations. As the sample size

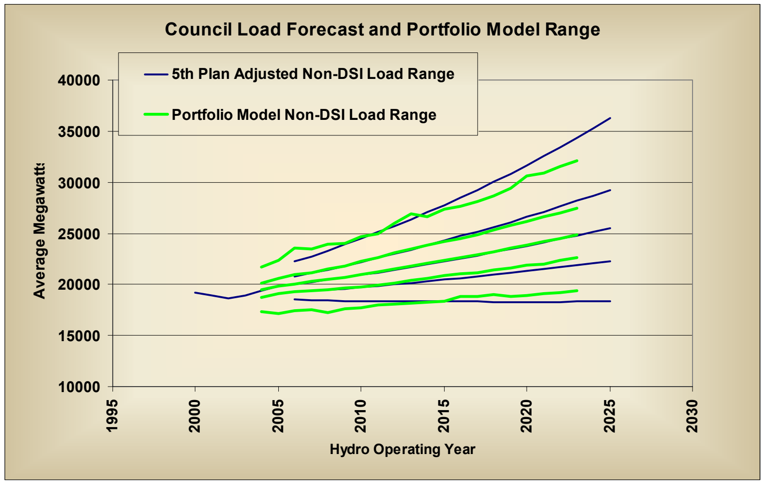

Figure P-34: Annual Averages of Deciles for Regional Model Loads

increases beyond the 750 samples reflected here, the zero percent and 100 percent deciles would grow apart. The maximum range of excursion in the Monte Carlo simulation is sensitive to the number of samples in the simulation.

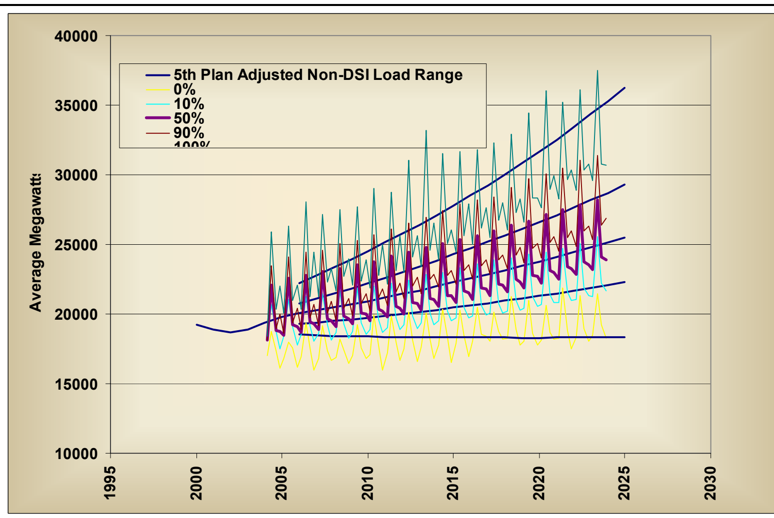

Figure P-35: Deciles for Quarterly Load

Because the regional model simulates hydro quarters, conversion of energy to that period is necessary. The basis of conversion (Ref [**6**]) is averages of monthly load allocation factors from Ref [**7**], which is integral to the study for the Council's primary load forecast.

Gas Price

Like electricity requirement load, natural gas price has characteristics that depend on the time scale. Although natural gas price does not vary a great deal across the day, there can be substantial variation within the month. The kinds of behavior that natural gas price demonstrates include:

- Chronological correlation, stronger than that for electricity prices perhaps due to the storage capability of natural gas
- Seasonal shapes
- Excursions due to disequilibrium of long-term supply and demand
- Daily variation within the month and hydro quarter
- Basis differential, in particular between regions separated by the Cascade mountain range
- Relatively small hourly price variation, because of storage capability within natural gas transmission lines. This eliminates the requirement for modeling on- and off-peak price differences

Natural gas prices also exhibit correlation with other variables. Natural gas prices correlate with loads and with electricity prices because weather affects all of these. Moreover, natural gas-fired generation is a marginal resource for power generation and consequently affects electricity price. Finally, higher electricity load generally places higher demand on natural gas markets. The model must capture both the long-term and short-term correlation among these variables.

Natural gas prices serve several functions or roles within the regional model. Short-term prices determine economic dispatch of gas-fired thermal generation. Forward gas prices feed decision criteria for the construction of new capacity. This section discusses the simulation of each of these uses.

As noted above, gas prices also influence longer-term electricity price. This influence of natural gas price on electricity prices appears in the discussion of electricity price uncertainty (See the section “Electricity Price”). Short-term correlation is outlined at the end of this chapter.

Worksheet Function and Formulas

Appendix L identifies how the regional model uses natural gas prices for the dispatch of gas-fired thermal generation and for the decision criteria for construction of new power plants. Appendix L traces natural gas price back to specific workbook cells. The description of natural gas prices in this Appendix P begins with those cells and continues the description back to the “building blocks” of these prices.

East of Cascade's gas prices {AQ180} are derived from those for west of Cascade's {AQ 68}. The worksheet range {A176: U176} provides the seasonal basis differential. The source of these basis differential values is [**8**]. The formulas in {Row 178} limit the lowest price in the East to $.20 per million BTU. This constraint assures Eastside prices remain positive irrespective of what Westside prices may do.

The formulas in {Row 180} add the values in {Row 179} to those in {Row 178}, but the values in {Row 179} are zero. This is a vestige of earlier logic, which attempted to add a contribution for fixed costs differentially to the east-of-Cascades natural gas prices. Council staff later decided that a fixed-cost adder would be inappropriate.

A lognormal process creates west-of-Cascade's natural gas prices, using combined factors, specific variation, and two jumps. Figure P-36 identifies the random variables for the natural gas price representation. The character of the jumps differs from that for the load's representation. The Council deemed the original size and duration of the jumps too large to be realistic. The Council substituted the representation in Figure P-37.

	R	S
63	14.7678	0.0886
64	12.1453	0.1863

Cell	Formula		interpretation
R65	=$R$63	wait_1	start time of jump 1
S65	=R65+ IF($S$63= 0,0,3/$S$63)	wait_1+ 3/size_1	end time of jump 1
T65	=$S$63	size_1	size_log xfr jump 1
U65	=S65	end time of jump 1	start time of recovery 1
V65	=U65+ S65*EXP(T65)	end time of jump 1 + duration recovery 1	end time of recovery 1
W65	=-T65/10	-size_1/10	size_log xfr recovery 1
X65	=V65+ $R$64	end time of recovery 1+ wait_2	start time of jump 2
Y65	=X65+ IF($S$64= 0,0,3/$S$64)	wait_2 + 3/size_2	end time of jump 2
Z65	=$S$64	size_2	size_log xfr jump 2
AA65	=Y65	end time of jump 2	start time of recovery 2
AB65	=AA65+ Y65*EXP(Z65)	end time of jump 2 + duration recovery 2	end time of recovery 2
AC65	=-Z65/10	-size_2/10	size_log xfr recovery 2
R66	= IF(AND(R$46>$R65,R$46<=$S65),$T65,0)+	jump_1	source: L28_P.xls
	IF(AND(R$46>$U65,R$46<=$V65),$W65,0)+	recovery_1	
	IF(AND(R$46>$X65,R$46<=$Y65),$Z65,0)+	jump_2	
	IF(AND(R$46>$AA65,R$46<=$AB65),$AC65,0)	recovery_2	

Figure P-36: Jump Data and Formulas for Natural Gas Price

The specific variance contributes to shoulder months in the spring and the fall. In contrast with several other stochastic variables, there seems to be much greater uncertainty in the price of natural gas during these off-peak seasons (see Reference [**9**]). This is perhaps due, in part, to the storage capability for natural gas and the buying that takes place in anticipation of the heating season and occasional surpluses resulting from warm winters. The values for the specific variances appear in Figure P-37.

Random Variables

	Type	Cell	Distribution	Parameters		
Jump 1	wait	{{R63}}	**uniform**	min 0	max 30	
	size	{{S63}}	**uniform**	min 0	max 0.70	
Jump 2	wait	{{R64}}	**uniform**	min 4	max 20	
	size	{{S64}}	**uniform**	min 0	max 0.70	
Principal Factors	offset	{{R56}}	**triangle**	min -1	mode 0	max 1
	linear	{{R58}}	**triangle**	min -1	mode 0.1	max 1
	quadratic	{{R60}}	**triangle**	min -1	mode 0	max 1
Specific Variance		{{row 67}}	**normal**	mean 0	stdev 0.30	

source: L28_P.xls

Figure P-37: Assumption Cells for Natural Gas Price

The principal factors appear Figure P-38. These were chosen largely to create realistic behavior. Some comparative statistics appear in the section "Comparison with the Council's Gas Price Forecast."

Principal Factors

Dec of Cal Year	offset	linear	quadratic
Weight			
	0.35	0.70	1.00
Value			
2003	0.50	0.07	0.00
2004	0.50	0.14	0.01
2005	0.50	0.21	0.02
2006	0.50	0.28	0.03
2007	0.50	0.35	0.05
2008	0.50	0.42	0.07
2009	0.50	0.49	0.10
2010	0.50	0.56	0.13
2011	0.50	0.63	0.16
2012	0.50	0.70	0.20
2013	0.50	0.77	0.24
2014	0.50	0.84	0.29
2015	0.50	0.91	0.34
2016	0.50	0.98	0.39
2017	0.50	1.05	0.45
2018	0.50	1.12	0.51
2019	0.50	1.19	0.58
2020	0.50	1.26	0.65
2021	0.50	1.33	0.72

source: L28_P.xls

Figure P-38: Principal Factors for Natural Gas Prices

The influence of principal factors, specific variance, and jumps combine just as they did for the construction of load futures. The cell {AQ68} contains the formula that combines these:

= AQ53*AQ54*EXP(AQ66+AQ62+AQ67)

where {AQ53} contains the benchmark (Council "medium" forecast, Reference [8]) value for natural gas in this period, {AQ54} is a special "stochastic adjustment," {AQ66} contains the sum of the jumps, {AQ62} is the sum of the factors, and {AQ67} is the contribution from the specific variance (seasonal uncertainty). The stochastic adjustments in row {66} are multipliers that would guarantee that the average, rather than the median, of the prices in that period match the benchmark. Early in the Council's studies, the Council identified their "medium" forecast with the average of the futures prices. Subsequently, the Council decided that the Council's medium forecast is a median forecast and the stochastic adjustment became 1.0 (no effect). That is, the Council constructs its forecast so that there is equal likelihood of the long-term equilibrium price being on either side of the forecast.

Forward Prices for Decision Criteria

Forward prices for natural gas play a key role in decisions about whether to construct new gas-fired power plants. The price of its fuel largely determines the value of the gas-fired power plant, and if future natural gas prices are low, the power plant will have greater value.

Some decision makers believe forward prices for natural gas are the best predictor of future spot price. The relationship between forward prices and current and future spot prices has been the subject of financial research for over 70 years. Arbitrage between forward and *current* spot price is possible for financial instruments and for commodities that can be stored. There is, therefore, a strict relationship between current spot prices and forward prices for these products. Natural gas, however, can only be stored in significant volumes up to about six months. Beyond that period, arbitrage opportunities are rare or nonexistent. For electricity, of course, the opportunities are even scarcer.

The relationship between forward prices and future spot prices is even weaker. The argument is often that the forward price incorporates all information about future spot price. This ignores, however, the question of whether the forward price in fact does a good job of predicting future spot price. A substantial body of research has demonstrated that long-term forward prices are a poor predictor of future spot prices for commodities that cannot be stored.[20] (In this context, "long-term" would be any period significantly longer than that which the commodity is stored.) Moreover, such an assessment ignores the influence of scarcity or abundance on the attitudes of hedgers or speculators, and these can bias the price up or down when there is uncertainty, even when all market participants share the same view of *expected* future spot price.[21]

Even if long-term forward prices are no better than throwing darts for predicting future prices, however, this does not mean that forward prices are irrelevant to the value of a power plant. On the contrary, an appropriate use of forward contracts is for hedging. If decision makers purchase the natural gas forward and sell the output of the power plant forward before proceeding with construction, changes in the values for forward contracts for natural gas and electricity will offset any change in the value of the plant due to fuel and output price variation. This provides a means for managing such risk associated with the "merchant" (un-hedged) portion of the plant. That is, an owner can make the merchant portion as small as desirable by hedging the rest of the plant. For various reasons, this hedging is likely to "lock in" as loss for the owners. However, decision makers view this loss as the cost of reducing risk, much like an insurance premium.

Forward prices continuously change, and this is an important source of uncertainty. One challenge for the portfolio model is to continuously forecast changing forward prices for natural gas and electricity. The question is, what is a reasonable basis for making such a forecast? Experience shows that forward prices tend to track current spot prices. Figure P-39 (Reference [**10**]) illustrates the relationship over time between current spot prices and a contract for delivery of natural gas in July 2003. The same kind of relationship exists for electricity. FERC analysis of electricity prices[22] in fact explicitly supports the position that spot prices move forward prices. (See discussion of the role of electricity spot prices in forecasting electricity forward prices on page P-72.)

---

[20] See, for example, Frank K. Reilly, ***Investment Analysis and Portfolio Management***, 2nd ed., The Dryden Press, Chicago 1979. See especially Chapter 24, "Commodity Futures," which discusses research for shell eggs, cattle, and other perishable commodities. For a more recent examination of electricity prices, see Longstaff and Wang, "Electricity Forward Prices: A High-Frequency Empirical Analysis," Anderson Graduate School of Management, UCLA, 2002.

[21] John C. Hull, ***Options, Futures, and Other Derivatives***, 4th ed., Prentice Hall 2000. See section 3.12, "Futures Prices and the Expected Future Spot Price."

[22] U.S. Dept. of Energy, Federal Energy Regulatory Commission, ***Final Report On Price Manipulation In Western Markets, Fact-Finding Investigation Of Potential Manipulation Of Electric And Natural Gas Prices***, Docket No. PA02-2-000, March 2003. PDF version.

This observation led the Council to adopt averages of current spot prices for natural gas over the prior 18 months as a simulated forecast of natural gas forward prices. In the cell {{AQ 249}}, the model averages the prior six periods (18 months) to estimate the

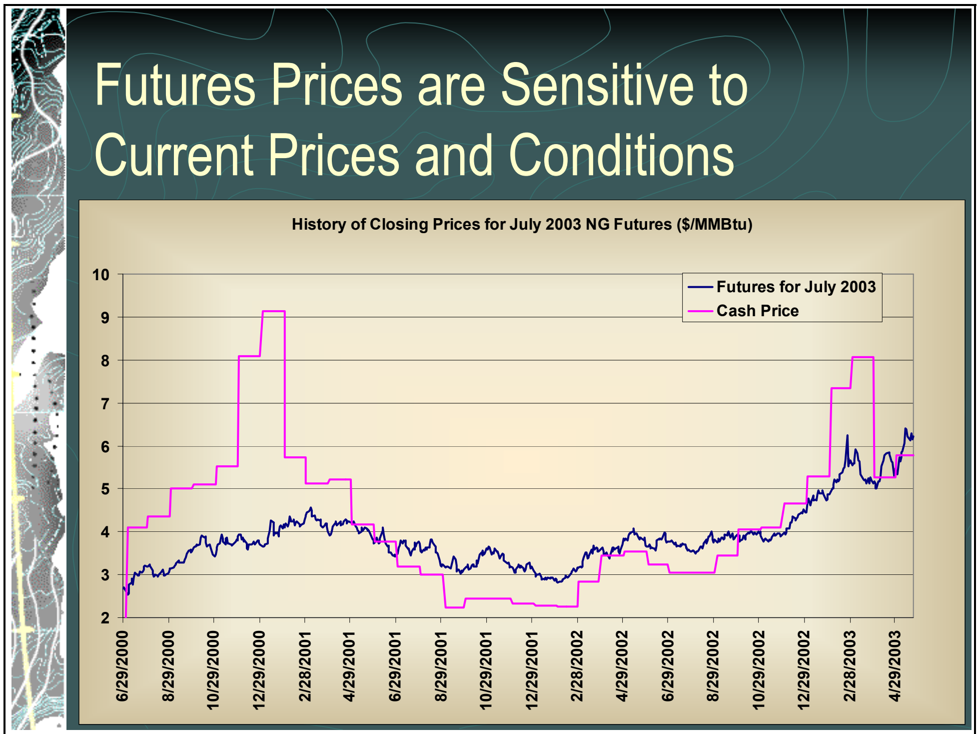

Figure P-39: Relationship Between Futures and Spot Prices

corresponding forward price for the decision criteria.

Hourly Behavior

Hourly volatility of natural gas prices within the period is taken as 10 percent, as indicated in the cell {{R55}}. Hourly price data for gas is not available to the Council, but casual exchanges with traders suggest this figure is representative. This appendix discusses the correlation of natural gas prices to other variables at the end of this chapter.

Comparison with the Council's Gas Price Forecast

In addition to preparing a long-term load forecast for the region, the Council prepares and updates long-term natural gas price forecasts. A comparison of the regional model's gas prices to the Council's forecast is more direct than the comparison to loads provided in the previous section.

Figure P-40 illustrates the quarterly natural gas price averages for four randomly chosen futures. Also shown, with the shaded area, is the range (high, median, and low) associated with the Council's natural gas price forecast. The quarterly averages fall well outside the range. In most of the futures, for example, there is at least one quarter when the natural gas price exceeds $10/MMBTU, well above the Council's "high" forecast. Some of the same caveats used in the comparison of the regional portfolio model's futures to the Council's load forecast apply here. The Council's forecast is a long-term equilibrium price forecast and does not capture excursions due to, for example, two- or three-year disruptions in supply and demand balance. Also, the Council's forecast is of annual averages, and quarterly averages will be more volatile.

By looking at statistical averages of the quarterly values for the regional portfolio model's natural gas price futures (Figure P-41), a more representative picture emerges. Quarterly averages for gas price can run from as low as $0.90 per MMBTU (2004 $) to as high as $28.24 per MMBTU, although those extremes are unlikely. The seasonal variation in price is not as extreme as that for load, so calculating annual averages for comparison with the Council's forecast is not essential. By carefully examining the deciles for quarterly gas price averages, it appears (Figure P-42) that there is about a 20

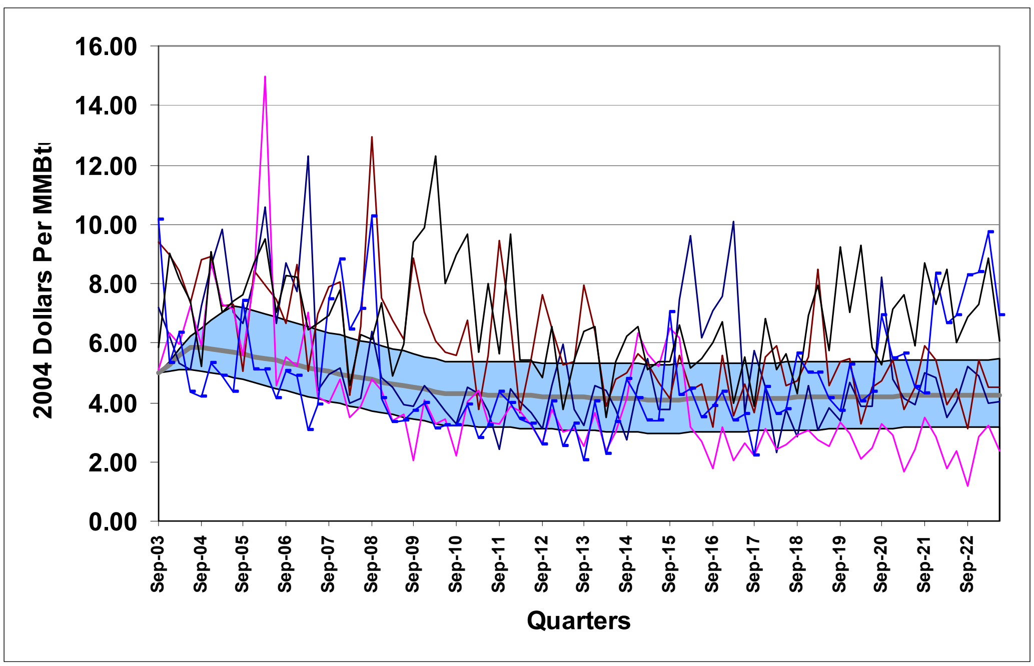

Figure P-40: Quarterly Natural Gas Prices for Selected Futures

percent chance of finding quarterly averages above the Council's high natural gas price forecast and a 20 percent chance of finding quarterly averages below the Council's low price forecast. The median of the price futures falls on top of the Council's median price forecast. This is all desirable behavior for these forecasts.

The results of the comparison of the regional model's natural gas price futures with the Council's forecasts are favorable. The only improvement on the regional model's representation that is evident after the fact is that, as has been the case in the past, the

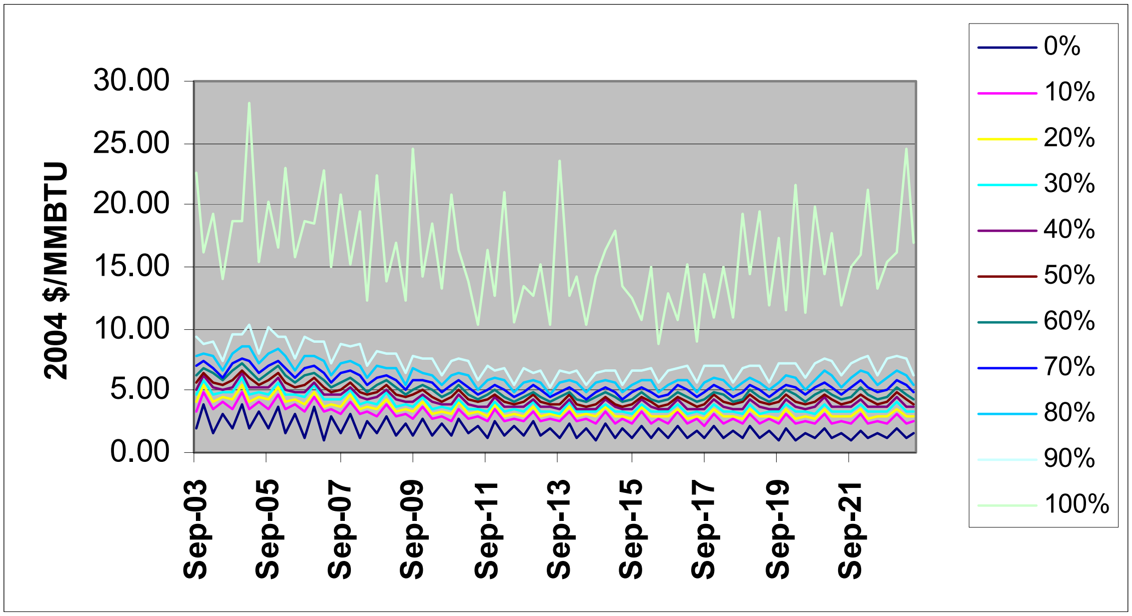

Figure P-41: Deciles for Quarterly Natural Gas Price Futures

Council's price forecasts may underestimate uncertainty (see Figure P-7). This may be a difficult situation to improve. The intuition of experts determines the range of uncertainty; without behavior that is consistent with experts' intuition, the results of the model do not have credibility. Perhaps the best outcome will be one where *low probability ranges* are as wide as feasible.

Hydro

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.

The regional model assumes that most hydrogeneration is insensitive to price. Hydrogeneration already occurs primarily on peak for both economic and reliability purposes, as much as non-economic constraints permit. The regional model captures differences in on- and off-peak generation, as described below. Nevertheless, there often remains a relatively small amount of energy that operators can shift among months for commercial reasons, without adversely affecting the refill probabilities of the system. Appendix L describes how the regional model captures that behavior using reversible supply curves. (See the Appendix L section "Price-Responsive Hydro.") The scope of hydrogeneration modeling that this Appendix P discusses is the energy that is not responsive to price.

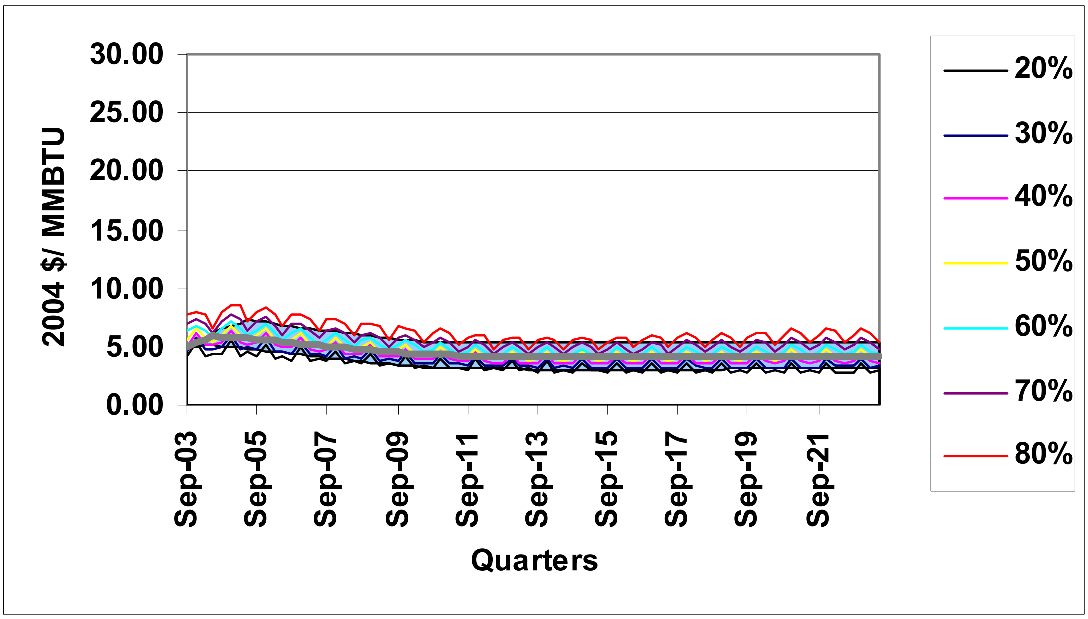

Figure P-42: Correspondence of Council's Natural Gas Price Range to Futures' Deciles

Data Sources and Representation

The source of all data for the price-invariant hydrogeneration is a BPAREGU.OUT file [**11**]. The Council's GENESYS model, specifically the HYDREG subroutine, produces this file.[23] HYDREG is the monthly hydro regulator for Genesys, the same hydro regulator that BPA, the Northwest Power Pool, and Canada use for determining rights under the Pacific Northwest Coordination Agreement (PNCA). HYDREG produces monthly generation for each hydro generation project in the region for each of 50 years (hydro years 1929-1978) of stream flow conditions. Figure P-43 illustrates the output of HYDREG for a single month in 2001 under a single (1929) stream flow condition. As

[23] Genesys is available for download from the Council's website. Contact John Fazio or Michael Schilmoeller, Council staff (503-222-5161), for directions on acquiring, installing, and using the model.

explained below, HYDREG models more facilities than appear in Figure P-43, and a complete list of such facilities appears in Figure P-44 and Figure P-46.

The regional energy value reported at the top of Figure P-43, under the heading "FINAL," primarily determines the energy used in the regional portfolio model. However, not all facilities in Figure P-43 contribute to the "FINAL" value. There are three reasons why energy is not included. First, the facility may have no generation. An example is Columbia Falls *gage* ("COLFLS") in Montana, which is a constraint on the hydro regulator. Gages always have zero energy under the column "AVMW" in Figure P-43. In Figure P-44 and Figure P-46, these have the word "gage" included in their names.

The second reason a facility may not contribute to the "FINAL" energy is that the facility may be located in Canada. Their operation is critical to the regulator, but the unit obviously does not directly contribute to regional energy. Any dams located in Canada have an asterisk in Figure P-43. In Figure P-44 and Figure P-46, these have the expression "(CAN)" included in their names, and their location is CAN. The capacity, ownership, and regulation status of Canadian facilities does not appear in the latter figures.

The third reason a facility would not contribute to the "FINAL" energy is that the PNCA does not incorporate its generation. Three Idaho facilities, Brown Lee, Oxbow, and Hell's Canyon, are part of the region and are regulated, but are not under the PNCA. The names of these three facilities have an asterisk in Figure P-43, as well.

Another class of regional plants that *contribute* to the region's energy supply but *do not* contribute to the "FINAL" energy is unregulated or "independent" plants. These are run-of-river plants and dams with capacity that is so small that HYDREG ignores their regulation. The names of these plants do not appear in Figure P-43 but are in Figure P-44 and Figure P-46, along with ownership and location information. They appear with regulation status "unreg." The total generation for the independents, however, does appear under the heading INDP at the top of Figure P-43.

HYDREG knows whether the hydro generator is east or west of the Cascades, and it produces a separate subtotal for each area. A special Council application [**12**] parses the BPAREGU.OUT file and creates a simple table of regional hydro generation (average MW) for both the Eastside and the Westside of the Cascades, by month and by hydro condition. Because the regional portfolio model needs all regional generation, the parsing application uses the "FINAL" energy from the BPAREGU.OUT file, adds in the unregulated generation from the "INDP" field, and adds the generation of Brown Lee, Oxbow, and Hell's Canyon.

One subtlety to preparing the hydro generation data lies in extracting on- and off-peak power from the monthly average energies that HYDREG produces. For the regional model, the on-peak period is 6 a.m. to 10 p.m. Monday through Saturday. The remaining hours are off peak. (Western power operations professionals refer to this subperiod definition of 16 on-peak hours on the six days of the week as the 6x16 or "six by sixteen" standard.) Although HYDREG does not provide subperiod values for systems hydrogeneration, extensive studies of sustained peaking capability for the system provide some guidance.

For their Fourth Power Plan, the Council commissioned Dr. Mike McCoy to make estimates of two-, four-, and 10-hour sustained peaking capability for the hydroelectric system.[24] An analysis of the conclusions from this study suggests that the peaking capability in average megawatts decreases roughly linearly with the number of hours of sustained capability [**13**]. With this assumption, the following equation relates on- and off-peak generation capability, using the 6x16 on-peak standard, to the average energy and 10-hour sustained peaking capability.

Let
E_p denote the on-peak (6days x 16hours) power (MW)
$\hat{E}_p$ denote the sustained 10-hour on-peak (5 weekdays) power (MW)
$\overline{E}$ denote the flat or average power over the entire week (MW)
X denote the average power (MW) over hours that do not contribute to 10-hour sustained peak
then

$$\overline{E} = \frac{5 \times 10 \times \hat{E}_p + (7 \times 24 - 5 \times 10)X}{7 \times 24}$$

$$E_p = \frac{(\text{saturday peak}) + (\text{weekday sus peak}) + (\text{weekday non-sus peak})}{\text{total peak hours}}$$

$$= \frac{16 \times X + 5 \times 10 \times \hat{E}_p + 5 \times 6 \times X}{(16 + 5 \times 10 + 5 \times 6)}$$

So solving for X gives us

$$X = \frac{7 \times 24 \times \overline{E} - 5 \times 10 \times \hat{E}_p}{(7 \times 24 - 5 \times 10)} \text{ so}$$

$$E_p = \frac{5 \times 10 \times \hat{E}_p + (16 + 5 \times 6) \times X}{(16 + 5 \times 10 + 5 \times 6)}$$

which gives us on-peak power in terms of average power and 10-hour, sustained peak power.

[24] Northwest Power Planning Council, "A Trapezoidal Approximation to the Pacific Northwest Hydropower System's Extended Hourly Peaking Capability Using Linear Programming," Appendix H 2, *Fourth Northwest Power Plan*.]

```
              BPA REGULATOR OUTPUT FOR SEPTEMBER (PERIOD  1)     WATER YEAR 1929 STUDY YR 2001 GAME   1
          2000 BIOLOGICAL OPINION - FSH027C Updated Spi

       DESIRED   FINAL     URC     ECC     PDP   XTRA1   XTRA2    INDP    PUMP       Draft Mode        ENER
EAST              7582.   6327.   7482.   7582.   7582.  14079.    267.             User Draft Point  9.00
WEST              1189.   1126.   1184.   1189.   1189.   1385.    400.
TOTAL    11615.   8771.   7453.   8665.   8771.   8771.  15464.    667.    137.

 PLANT   NO.  NAT Q  Q OUT   QMIN  FORCE  BYPAS  OTHER  OVERG I HKSM  AVMW   DRAFT ENDSTO  ELEV    URC    ECC    AER   CON VIOL
 -----  ---- ------  -----  -----  -----  -----  -----  ----- ------ -----  ------ ------ ------ ------ ------ ------  --- -- --
CUSH 1  2208    113   1390    100      0      0      0      0  49.08    26    38.3  149.1  718.5  171.3  161.5  149.1  FC
CUSH 2  2206    114   1391      0      0      0    100      0  30.59    43
ALDER   2190    450    683    300      0      0      0      0  20.25    14     7.0   74.4 1202.3   81.4   79.7   74.4  FC
LAGRND  2188    450    683      0    683      0      0      0   0.00    21
WHITE   2160    606    237    100      0      0    130      0  32.00     3   -11.1   20.9  540.9   23.5   21.4   20.9  FC
ROSS    2070   1047   2288    788      0      0      0      0  84.31    70    37.2  474.0 1592.5  530.5  482.3  473.9  QH  FC
DIABLO  2067   1613   2854      0      0      0      0      0  53.56    74
GORGE   2065   1759   3000   1500      0      0      0      0  27.06    87
U BAKR  2028   1030   1265      0      0      0      0      0  39.41    26     7.0   72.7  707.3  111.2   74.2   72.7  FC
L BAKR  2025   1220   1495     80      0      0      0      0  19.16    29     1.2   70.6  437.5   71.8   71.8   70.6  FC
MICA  * 1890  19310  19310  10000      0      0      0      0  84.62   834     0.0 5825.1 2470.1 5825.1 5825.0 5825.1  FC
REVELS* 1870  25280  25280      0      0      0      0      0  84.62   806     0.0  557.0 1875.6  557.0  557.0  557.0  FC
ARROW * 1831  36329  40642   5000      0      0      0      0  84.62     0   129.4 3450.2 1442.0 3579.6 3233.6 3450.2  FC
LIBBY   1760   5072   9656   4000      0      0    200      0 107.95   221   137.5 1923.8 2432.5 2510.5 1731.6 1923.8  PD  FG
BONFER  1740   6694  11279      0  11279      0      0      0  84.62     0
DUNCAN* 1681   2140   2973    100      0      0      0      0  84.62     0    25.0  680.8 1889.2  705.8  678.8  680.8  FC
CORA L* 1665  12270  13971   5000      0      0      0      0  84.62    21  -111.5  396.9 1745.3  396.9  396.9  396.9  FC
CANAL * 1664  12270   8971      0      0      0      0      0  84.62   177
UP BON* 1663  12270   5000      0      0      0      0      0  84.62    21
LO BON* 1660  12270   5000      0      0      0      0      0  84.62    23
S SLOC* 1658  12270   5000      0      0      0      0      0  84.62    25
BRILL * 1652  12197  13898      0      0      0      0      0  84.62    96
H HORS  1530    646   1419   1419      0      0      0      0 180.93    49    23.2 1290.0 3538.0 1549.0 1259.7 1290.0  QP  QL SL
COLFLS  1520   2727   3500   3500   3500      0      0      0 146.48     0
KERR    1510   3847   5083   3200      0      0      0      0 146.48    72    13.9  600.8 2892.8  614.7  575.4  600.8  FC
THOM F  1490   8840  10076   6000      0      0      0      0 132.37    45
NOXON   1480   6862   8098   3727      0      0      0      0 128.93    94     0.0  108.5 2329.0  116.3  108.5  108.5  FC
CAB G   1475   8136   9371   5000      0      0      0      0 117.27    65
PRST L* 1470    118      1      0      0      0      0      0 110.37     0    -3.5   25.0    2.1   35.5   26.0   25.0  FC
ALBENI  1465   9656  14665   4000      0      0     50      0 110.37    30   116.7  465.7 2060.0  582.4  465.7  465.7  FC
BOX C   1460   9806  14815      0      0      0      0      0 108.30    41
BOUND   1450   9939  14948      0      0      0      0      0 105.62   308
7-MILE* 1442  10206  15215      0      0      0      0      0  84.62   242
WANETA* 1440  10206  15215      0      0      0      0      0  84.62   244
CDA LK* 1341    704   1634    300      0      0      0      0 122.77     0    27.9   84.6 2126.6  112.5   86.9   84.6  FC
POST F  1340    704   1634    300      0      0      0      0 122.77     6
UP FLS  1332   1350   2280      0      0      0      0      0 118.92    10
MON ST  1330   1350   2280      0      0      0      0      0 115.92    13
NINE M  1315   1753   2683      0      0      0      0      0 105.45    11
LONG L  1305   2156   3092      0      0      0      0      0 101.70    36     0.2   50.1 1535.0   52.5   50.2   50.1  FC
L FALL  1302   2156   3092      0      0      0      0      0  89.97    16
COULEE  1280  56077  64261  50000      0      0      0      0  84.62  1574  -113.3 2329.7 1283.0 2614.3 2368.4 2329.7  PD
CH JOE  1270  56117  64301      0      0      0    500      0  60.12   822     0.0    0.0  953.8    0.0    0.0    0.0
WELLS   1220  58698  66882      0      0      0   1200      0  47.23   337
CHELAN  1210    647   1637     50      0      0      0      0  68.79    43    29.7  308.5 1098.0  341.5  308.3  308.5  FC
R RECH  1200  59404  68578      0      0      0      0      0  42.75   457
ROCK I  1170  61975  71149      0      0      0      0      0  36.08   209
WANAP   1165  62061  71235      0      0      0   2200      0  33.25   410
PRIEST  1160  62340  71514  36000      0      0   2200      0  27.60   413
BRNLEE*  767  14452  14452   5000      0      0      0      0  50.27   252     0.0  293.8 2045.0  491.7  411.2  293.8  PD
OXBOW *  765  14452  14452      0      0      0    100      0  50.27   112
HELL C*  762  14497  14497      0      0      0      0      0  50.27   215
DWRSHK   535   1060   1300   1300      0      0    100      0  93.13    51     7.2  388.6 1518.9  902.6  378.1  388.6  QL  SA
LR.GRN   520  22361  22600  11500      0      0    670      0  50.27   154     0.0  225.0  733.0  245.8   78.1  225.0  FC
L GOOS   518  22361  21783  11500      0      0    630      0  43.27   147   -24.5  285.0  638.0  285.0  128.6  285.0  UR  FC
LR MON   504  21657  20758  11500      0      0    750      0  36.30   142    -9.6  190.1  540.0  190.1   83.2  190.1  FC
ICE H    502  21647  20367   7500      0      0    740      0  29.20   137   -11.4  204.8  440.0  204.8   90.8  204.8  FC
MCNARY   488  79759  87654  50000      0      0   4000      0  22.23   458     0.0    0.0  338.7    0.0    0.0    0.0
J DAY    440  80738  88631  50000      0      0    800      0  16.76   664    -0.1  127.8  262.5  269.7  127.8  127.8  PD  FC
RND B    390   3154   3302   2800      0      0    200      0  47.09    81     4.4  131.9 1941.7  138.3  135.5  131.9  FC
PELTON   388   3354   3502   3000      0      0      0      0  21.10    33
REREG    387   3354   3502      0      0      0      0      0  11.74     8
DALLES   365  84499  92540  50000      0      0   4300      0   9.20   540
BONN     320  87725  95766      0      0      0   8400      0   4.86   424     0.0    0.0   74.1   -1.0    0.0    0.0  PL  UR
TMTHY    117     87    177     10      0      0      0      0  86.88     0     2.7   28.4 3186.1   31.1   29.2   28.4  FC
OK GRV   115    344    434      0      0      0      0      0  86.88    27
NFORK    111    822    912      0      0      0      0      0  23.89     9
FRDAY    110    822    912      0      0      0      0      0  13.99     8
R MILL   108    822    912      0      0      0      0      0   5.10     5
SWFT 1    82    715    555      1      0      0      0      0  69.69    16    -4.8  219.4  997.2  225.4  225.4  219.4  FC
SWFT 2    80    715    555      0      0      0      0      0  40.59     5     0.0    0.0  603.0    0.0    0.0    0.0  FC
YALE      78    837    561      0      0      0      0      0  32.34    10    -3.5   93.1  488.7   95.6   95.6   93.1  FC
MERWIN    76    933   1475    800      0      0      0      0  13.77    20    24.6   62.4  224.4   92.1   63.2   62.4  FC
PCKW L*   63     56     56     10      0      0      0      0  40.54     0     0.0    0.0 2850.5    1.4    0.0    0.0  FC
MOSSYR    48   1194   3957   2858      0      0      0      0  40.54   103    82.9  566.3  763.0  654.3  654.3  566.3  FC
MAYFLD    42   1336   4099   3000      0      0      0      0  14.63    57
-------------------------------------------------------------------------------------------------------------------------------
                    RESULTS FOR SEPTEMBER  (PERIOD  1)     WATER YEAR 1929          INTERLACE PERIOD    3

PHASE                DESIRED  ACTUAL = (REGUL  + INDEP   - PUMP)     SPILL   MICA+RVL  DRAFT PT   DRAFT
------------------      MW       MW       MW       MW       MW        MW       MW                 MW-MO
PROPORTIONAL DRAFT  12145.0   9300.8   8770.8    667.0    137.0       0.           6.000000
MICA NON-TRTY (BC)   1046.1      0.0      0.0                         0.       0.  MAX STORE    0  MAX RETURN    0
MICA NON-TRTY (US)   1020.8      0.0      0.0                         0.       0.  MAX STORE    0  MAX RETURN    0
ALLOCATE SPILL                   0.0      0.0                         0.
FINAL OPERATION     12145.0   9300.8   8770.8    667.0    137.0       0.                     7419.
ENERGY CONTENT (MW-MO) RELATIVE TO TARGET  54683. 33 PLANTS ABOVE BY  54683.     0 BELOW BY      0.
ENERGY CONTENT (MW-MO) RELATIVE TO ECC   1807.  TO PDP(AER)      0.
THE DALLES FLOW AT URC=      0.  ECC=      0.  TREATY=  92540.  FINAL=  92540.
```

Figure P-43: Sample from a BPA HYDSIM Regulator BPARegu.out file

Name	Cap (MW)	ownership	regulated	location
Albeni Falls	43	Fed	Reg	OR/WA
Alder	50	Non-Fed	Reg	OR/WA
American Falls	92	Non-Fed	Unreg	ID
Anderson Ranch	40	Fed	Unreg	ID
Arrow (CAN)				CAN
Big Cliff	18	Fed	Unreg	OR/WA
Big Creek (Flathead Irr Prj, MT)	1	Non-Fed	Unreg	MT
Black Canyon	10	Fed	Unreg	ID
Bliss	75	Non-Fed	Unreg	ID
Boise Diversion (USBR)	2	Fed	Unreg	ID
Bonners Ferry gage				ID
Bonneville	1093	Fed	Reg	OR/WA
Boundary	951	Non-Fed	Reg	OR/WA
Box Canyon (PEND)	60	Non-Fed	Reg	ID
Brill (CAN)				CAN
Brownlee	585	Non-Fed	Reg	ID
Bull Run (PGE)	21	Non-Fed	Unreg	OR/WA
C.J. Strike	83	Non-Fed	Unreg	ID
Cabinet Gorge	222	Non-Fed	Reg	ID
Calispel Creek	1	Non-Fed	Unreg	ID
Canal (CAN)				CAN
Carmen Smith	90	Non-Fed	Unreg	OR/WA
Cascade (IDPC)	12	Non-Fed	Unreg	ID
Cedar Falls (SCL)	20	Non-Fed	Unreg	OR/WA
Chandler	12	Fed	Unreg	OR/WA
Chelan	48	Non-Fed	Reg	OR/WA
Chief Joseph	2457	Fed	Reg	OR/WA
City of Idaho Falls	42	Fed	Unreg	ID
Clear Lake (IDPC)	3	Non-Fed	Unreg	ID
Clearwater 1,Clearwater 2	41	Non-Fed	Unreg	OR/WA
Coeur D'Alene Lake gage				ID
Columbia Falls gage				MT
Condit	10	Non-Fed	Unreg	OR/WA
Copco 1	20	Non-Fed	Unreg	OR/WA
Copco 2	27	Non-Fed	Unreg	OR/WA
Corra Linn (CAN)				CAN
Cougar	25	Fed	Unreg	OR/WA
Cowlitz Falls (Lewis Co PUD)	70	Non-Fed	Unreg	ID
Cushman 1	43	Non-Fed	Reg	OR/WA
Cushman 2	81	Non-Fed	Reg	OR/WA
Dalles	1807	Fed	Reg	OR/WA
Detroit	100	Fed	Unreg	OR/WA
Dexter	15	Fed	Unreg	OR/WA
Diablo	123	Non-Fed	Reg	OR/WA
Duncan (CAN)				CAN
Dworshak	400	Fed	Reg	ID
Electron	26	Non-Fed	Unreg	OR/WA
Faraday	35	Non-Fed	Reg	OR/WA
Fish Creek	11	Non-Fed	Unreg	OR/WA
Foster	20	Fed	Unreg	OR/WA
Gorge (SCL)	207	Non-Fed	Reg	OR/WA
Grand Coulee	6494	Fed	Reg	OR/WA
Green Peter	80	Fed	Unreg	OR/WA
Green Springs	16	Non-Fed	Unreg	OR/WA
Hells Canyon	392	Non-Fed	Reg	ID
Henry M Jackson (Snohomish PUD)	112	Non-Fed	Unreg	OR/WA
Hills Creek	30	Fed	Unreg	OR/WA
Hungry Horse	428	Fed	Reg	MT
Ice Harbor	603	Fed	Reg	OR/WA
Iron Gate	18	Non-Fed	Unreg	OR/WA
Island Park Hydroelectric Proj	5	Fed	Unreg	ID
John C Boyle	80	Non-Fed	Unreg	OR/WA
John Day	2160	Fed	Reg	OR/WA
Kerr	168	Non-Fed	Reg	MT
La Grande	64	Non-Fed	Reg	OR/WA
Leaburg	14	Non-Fed	Unreg	OR/WA
Lemolo units 1& 2	62	Non-Fed	Unreg	OR/WA
Libby - USCEPD	525	Fed	Reg	MT

Figure P-44: Facilities Contributing to Hydrogeneration (1/2)

Off-peak (168-6x16) hours are a subset of the hours to which *X* pertains. Therefore, the off-peak power is exactly *X*. The sustained peaking information is from reference [**14**], which provides relationships between 2-, 4-, and 10-hour sustained peak capacity as a function of system energy for each month.

The special Council application [**15**] that parses the BPAREGU.OUT file uses the appropriate number of on- and off-peak hours for each month to estimate average on- and off-peak power (MW). For the regional model, another Council application reduces these data to hydro year quarters [**16**].

Worksheet Function and Formulas

Turning to the worksheet function that provides this data to the regional model, we note that several versions of the function exist and are available to the public. One of these, for example, is an Excel add-in that provides monthly energies in both megawatt-hours and average energy, on peak and off peak, as well as sustained, 10-hour peak generation for the region, for each stream flow condition, and separately for east and west of the Cascades. The version used in the regional portfolio model, however, is not an Excel add-in, but instead a VBA function that reads a worksheet ("For AddIn ver 7") of data.[25] This section returns shortly to the description of this function.

The regional model uses hydrogeneration for three purposes: meeting energy requirements, influencing electricity price, and for planning long-term resource requirements. The influence on electricity price is discussed in the following section, "Electricity Price." For planning long-term resource requirements, the model uses critical hydrogeneration levels, which the model assumes remain constant. Consequently, this section outlines only the use of hydrogeneration for meeting energy requirements.

The discussion of hydrogeneration in Appendix L refers to the on-peak average MWh hydrogeneration in a specific, but representative cell, {AQ 36} in the example workbook L24DW02-f06-P.xls. (This is identical to cell {{AQ 36}} in L28_P.xls.) The on-peak calculation in {{AQ 36}} is

```
=(AQ33-300*AP$21/79)*1152
```

This differs from the formula in {AQ 36}, "=AQ33*1152," in the draft plan workbook. Between the draft and final plan, the Council added a loss of hydroelectric availability over the twenty years of the study. The beginning of this section describes the reasons for this loss. The loss is deterministic and increases linearly with time to 300MWa by the end of the study. Incorporating that loss is what the additional term -300*AP$21/79 achieves.

[25] The use of Excel add-ins complicates the use of distributed computing with Decisioneering, Inc.'s CB Turbo®, described in Appendix L. Each machine would have to be equipped with a copy of the add-in, so changing any logic in the add-in becomes burdensome.

The cell {{AQ 33}} references the VBA function that provides average MW for the period:

=vfuncHydro4x2W($R$24:$CS$24,1)

VBA function vfuncHydro4x2W takes as its first argument a range containing cells that assume random values – one for each hydro year – between 0.0 and 50.0. In the preceding example of {{AQ 33}}, the range is $R$24:$CS$24. These random numbers determine the stream flow condition for the hydro year (September through August of the following year). We return to this determination in a moment.

After the range, the function takes an integer that specifies the subregion for which hydrogeneration is requested. A zero designates hydrogeneration for east of the Cascades; the one in {{AQ 33}} designates hydrogeneration west of the Cascades.

The function returns a range two rows high and 80 columns wide, in the case of the regional model. The range contains cells with the hydrogeneration (MWa) for that subregion, for each period (column). The first row contains on-peak hydrogeneration; the second row contains off-peak hydrogeneration.

It may be helpful to examine the VBA function vfuncHydro4x2W from a couple of perspectives. The definition of the vfuncHydro4x2 function is as follows:

```
Function vfuncHydro4x2(ByRef rYears As Range, ByVal lLoc As Long, Optional
        ByVal, lStartPeriod As Long = 0) As Variant

Takes:
rYears - Range, pointing to a vector of single [0.00-50.00] representing the years 1929-
        1978, sorted ascending by annual energy. For example, the user can have Excel
        pass 50 * rand() as sYear to this function to get draws of hydro condition.
        Ascending order permits user to correlate annual energy with other variable.  To
        access a particular year, use the sfuncYear() function, below.

lLoc - 0, East only
       1, West only
       2, East+West Generation

lStartPeriod - Optional'
       0, (default), Range of returned energies starts with Sep - Nov
       1, Dec - Feb
       2, Mar - May
       3, Jun - Aug
================================================================
Returns:
       A variant containing an array of period Hydrogeneration (MWa) for east-side or
       west-side generation, or both.  The value of each element of the array corresponds
       to the value of the hydro year choice, for the appropriate region and subperiod
```

For a different perspective on what this function is doing, consider the auditing references in Figure P-45. The average MW of generation in cell {{U26}} is one entry of a range, {{R26:CS27}}, which the function is returning. The value of {{U26}} is the on-peak hydrogeneration east of the Cascades for a particular hydro year. For which hydro year does the function return generation? The function is returning the fourth quarter for the first hydro year, so it uses the random number at the beginning of the hydro year, cell {{R24}} from the input range {{$R$24:$CS$24}}.

To what historical hydro year do the values correspond? In Figure P-45, the random number in cell {{R24}} has the value 49.38926508. There are 50 years of hydrogeneration data. The generation returned is for the year, according to the rank by annual hydrogeneration energy, from lowest to highest. For example, the random number 49.38926508 lies in the last bin, (49,50], so the year with the highest annual hydrogeneration would be returned, in this case hydro year 1973-1974. If the random number had been 0.5 on the other hand (or any number less than 1.0), the function would return the driest year on record, 1931, as determined by total annual generation.

A separate function simplifies the process of getting data for a particular hydro year. The regional model does not use the function sfuncYear, but the Council would make it available to any party on request. It returns a real number corresponding to each hydro year that the vfuncHydro4x2 function returns. Its definition follows.

Function sfuncYear(ByVal lYear As Long, ByVal lType As Long) As Single
Takes a calendar year, e.g., 1937, and returns a real single with a value in the middle of the correct "bin" for that year, for use as input to vfuncHydroGen. For example, 1937 is the second lowest year for Eastside Hydro, in terms of annual energy and is therefore the second entry in vfuncHydroGen(*,0). Then sfuncYear(1937,0) = 1.5 (The first bin is [0,1), the second is [1,2), etc.

lYear - calendar year, as long
lType - 0, East Generation only
1, West Generation only
2, East+West Generation

This concludes the description of the model worksheet VBA function. This section next considers the assumed hourly behavior of hydrogeneration.

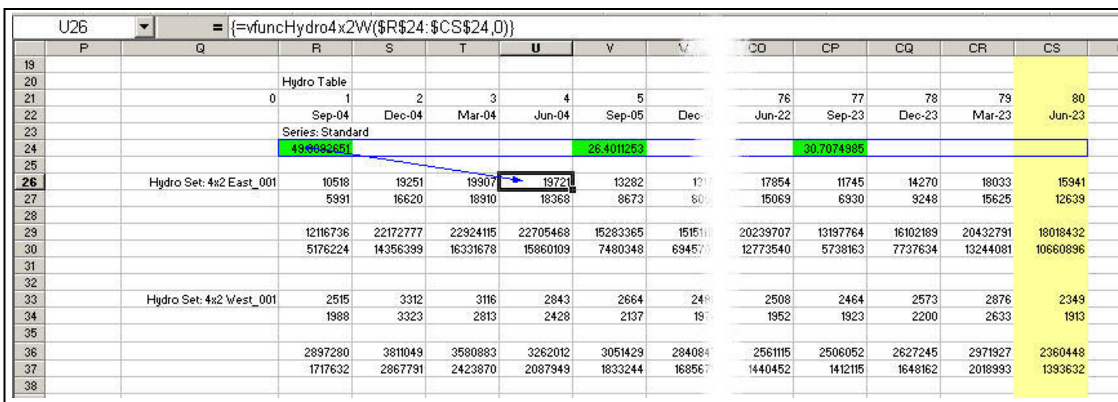

Figure P-45: Hydro Function References

Name	Cap (MW)	ownership	regulated	location
Little Falls (WWPC)	32	Non-Fed	Reg	OR/WA
Little Goose	810	Fed	Reg	OR/WA
Long Lake	70	Non-Fed	Reg	OR/WA
Lookout Point	120	Fed	Unreg	OR/WA
Lost Creek	49	Fed	Unreg	OR/WA
Lower Baker	64	Non-Fed	Reg	OR/WA
Lower Bonnington (CAN)				CAN
Lower Granite	810	Fed	Reg	OR/WA
Lower Malad	14	Non-Fed	Unreg	ID
Lower Monumental	810	Fed	Reg	OR/WA
Lower Salmon	60	Non-Fed	Unreg	ID
Mayfield	162	Non-Fed	Reg	OR/WA
McNary	980	Fed	Reg	OR/WA
Merwin	136	Non-Fed	Reg	OR/WA
Mica (CAN)				CAN
Mill Creek	1	Fed	Unreg	OR/WA
Milner (IDPC)	59	Non-Fed	Unreg	ID
Minidoka	8	Fed	Unreg	ID
Monroe Street	15	Non-Fed	Reg	OR/WA
Mossyrock	300	Non-Fed	Reg	OR/WA
Nine Mile	26	Non-Fed	Reg	OR/WA
North Fork	38	Non-Fed	Reg	OR/WA
Noxon Rapids	467	Non-Fed	Reg	MT
Oak Grove	51	Non-Fed	Reg	OR/WA
Oxbow (IDPC)	190	Non-Fed	Reg	ID
Packwood	30	Non-Fed	Unreg	OR/WA
Packwood Lake gage				OR/WA
Palisades (USBRCO)	177	Fed	Unreg	ID
Pelton	97	Non-Fed	Reg	OR/WA
Pelton Re-Regulation	18	Non-Fed	Reg	OR/WA
Post Falls	15	Non-Fed	Reg	OR/WA
Priest Lake gage				OR/WA
Priest Rapids	923	Non-Fed	Reg	OR/WA
Prospect units 1-4	44	Non-Fed	Unreg	OR/WA
Revelstoke (CAN)				CAN
River Mill	19	Non-Fed	Reg	OR/WA
Rock Island Powerhouse	624	Non-Fed	Reg	OR/WA
Rocky Reach	1280	Non-Fed	Reg	OR/WA
Ross Dam	360	Non-Fed	Reg	OR/WA
Round Butte	247	Non-Fed	Reg	OR/WA
Roza	13	Fed	Unreg	ID
Seven Mile (CAN)				CAN
Shoshone Falls	13	Non-Fed	Unreg	ID
Slide Creek	18	Non-Fed	Unreg	OR/WA
Smith Creek (EWEB)	38	Non-Fed	Unreg	OR/WA
Snoqualmie	42	Non-Fed	Unreg	OR/WA
Soda Springs	11	Non-Fed	Unreg	OR/WA
South Slocan (CAN)				CAN
Stone Creek (EWEB)	12	Non-Fed	Unreg	OR/WA
Strawberry Creek (Lower Valley P&L)	2	Non-Fed	Unreg	ID
Swan Falls	25	Non-Fed	Unreg	ID
Swift 1	204	Non-Fed	Reg	OR/WA
Swift 2	70	Non-Fed	Reg	OR/WA
T.W. Sullivan	15	Non-Fed	Unreg	OR/WA
Thompson Falls (MPC)	93	Non-Fed	Reg	MT
Thousand Springs	9	Non-Fed	Unreg	ID
Timothy Lake gage				OR/WA
Toketee Falls	43	Non-Fed	Unreg	OR/WA
Trail Bridge (EWEB)	10	Non-Fed	Unreg	OR/WA
Twin Falls (IDPC)	44	Non-Fed	Unreg	ID
Upper Baker	105	Non-Fed	Reg	OR/WA
Upper Bonnington (CAN)				CAN
Upper Falls (WWP)	10	Non-Fed	Reg	OR/WA
Upper Malad	8	Non-Fed	Unreg	ID
Upper Salmon Falls	35	Non-Fed	Unreg	ID
Wanapum	1038	Non-Fed	Reg	OR/WA
Waneta (CAN)				CAN
Wells (DOPD)	774	Non-Fed	Reg	OR/WA
White River (PSPL)	70	Non-Fed	Reg	OR/WA
Yale	108	Non-Fed	Reg	OR/WA
Yelm (Centralia)	10	Non-Fed	Unreg	OR/WA

Figure P-46: Facilities Contributing to Hydrogeneration (2/2)

Hourly Behavior

Recall that there are two types of hydrogeneration in the regional model, the type that this section discusses, which does not respond to electricity market prices, and the market-price responsive type. Appendix L has a description of how the regional model captures the latter. (See pages L-48 and L-106.)

At the hourly time step, there is certainly a difference for non-price responsive hydrogeneration on- and off-peak. Because the function vfuncHydro4x2W already accounts for these differences through separate returned values, however, the question of any remaining variation means variation *within* the respective subperiods. If there is any such residual variation in hydrogeneration, the model assumes it is small and uncorrelated with electricity price. The hydrogeneration valuation calculations in the model therefore implicitly assume a zero correlation between hourly hydrogeneration and hourly electricity market price. (See page L-50.)

Electricity Price

Many forecasters use long-term equilibrium price models to estimate future electric power prices. These models result in annual average electricity prices that equal the fully allocated cost of the plant used for expanding system capacity, which in the West is typically a combined-cycle combustion turbine (CCCT). 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, the regional model must incorporate correlation of electricity prices with hydropower availability, loads, and natural gas prices. Correlation between electricity prices and load on the time scale of the hydro quarter should have the opposite sign of the correlation on the time scale of years. That is, demand elasticity of loads needs attention.

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, e.g., 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.

This section begins with an overview of the construction of electricity prices in the regional model. It describes how the model accommodates the requirements just mentioned. The treatment addresses price averages at the time scale of the hydro quarter-year. The model uses electricity prices for energy requirement valuation, as input to various decision criteria, and for producing load elasticity, and the section explores those in turn. The section then traces the formulas in the sample Excel workbook portfolio model from the point where the discussion of Appendix L, "The Portfolio Model" leaves off. Finally, it elaborates on some of the hourly price behavior, which typically is different from that at the time scale of the hydro quarter.

Background

At its December 19, 2002, meeting, the Council's System Analysis Advisory Committee (SAAC) discussed the influence that various sources of uncertainty have on each other. Figure P-47 resembles the Influence Diagram that the SAAC used. Most of the influences are predictable. As hydro generation increases, for example, electricity prices should decrease. In the short term, increases in load, natural gas prices, and forced outages should push up the price of electricity.

There are hosts of factors besides regional hydro generation, load requirements, natural gas prices, and forced outage rates, however, that influence regional electricity prices. (For brevity, we will refer to regional hydro generation, load requirements, natural gas prices, and forced outage rates as the "local variables" in the following.) First, the values of local variables do not capture the corresponding influences from outside the region. For example, economic recession and load reduction in California or the Pacific Southwest would probably have the effect of depressing electricity prices in the Pacific Northwest. Second, there are certainly factors that influence electricity price besides the four just identified. Over the long term, technology innovation could easily trump the influence of these four. Unanticipated changes in legislation or the regulation of electricity could influence the availability of supply both within the region and outside the region. Changes in supply availability from outside what we traditionally think of as the region is another factor. Examples of these influences are regional Independent Power Producers (IPP) and California's initiative to implement a strong reserve margin. While it might be possible to model these individual factors explicitly, a surrogate for these effects is an unanticipated excursion in electricity price that is independent of the local variables. That is, such excursions are the primary means by which supply outside the traditional region's system influences regional costs.

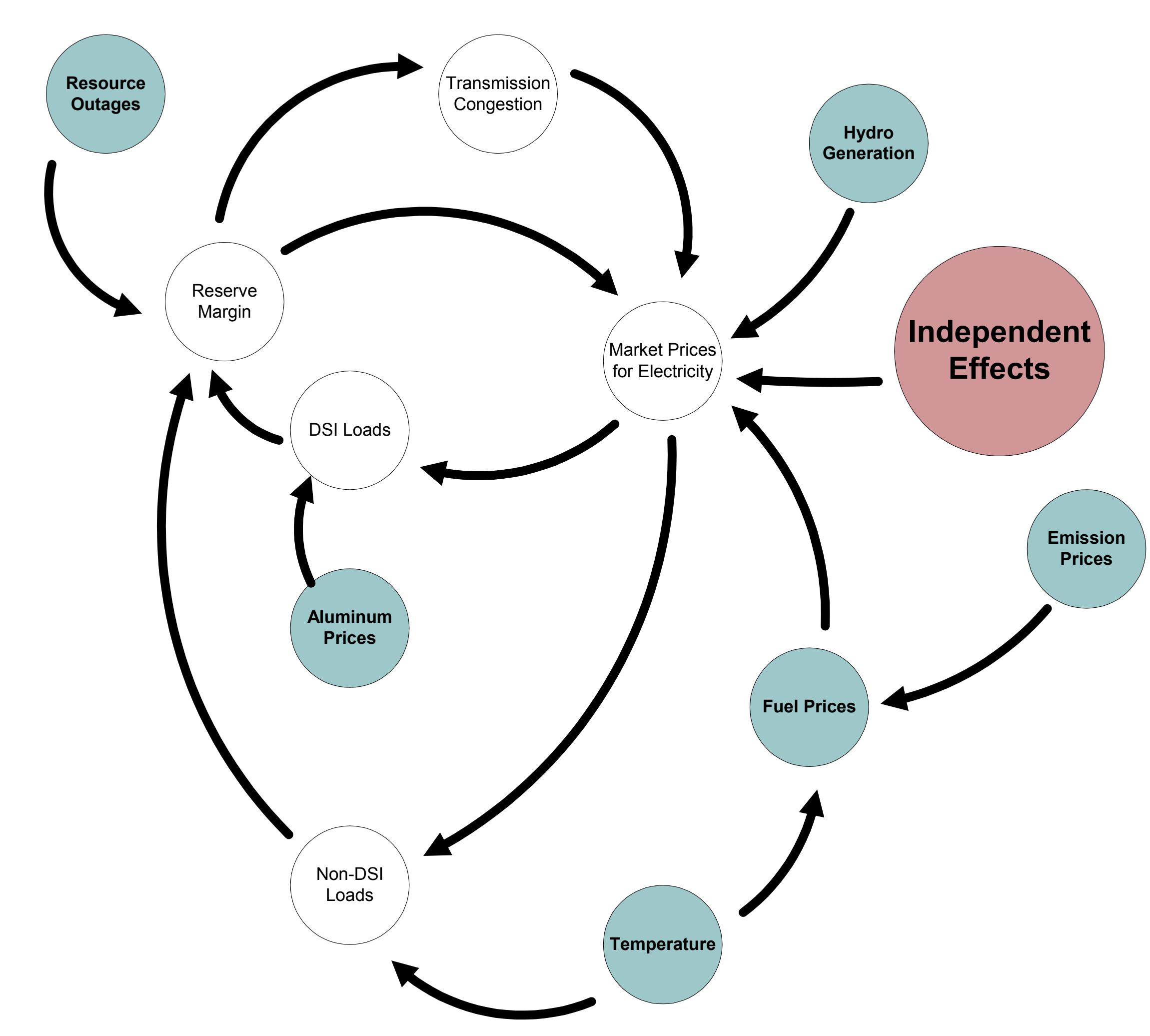

Figure P-47: Independent Effects

The Council used Bench Mark Heuristics (BMH) to study the statistical behavior of electricity prices, transmission, load requirements, natural gas prices, hydro generation, and a host of other related data [**17**]. BMH studied each of the factors individually, and created a detailed regression model for each, using an ARMA process to simulate the error term. BMH then modeled the relationship between local electricity prices and local loads, natural gas prices, and hydro generation, seasonal, and weekday factors. Based on the best explanatory model BMH produced, local variables explain only about 43 percent of the change in daily electricity prices [**18**]. When markets are in transition, the influence of these local variables is even smaller. There is a significant amount of variation in electricity price behavior that local variables do not explain. Figure P-47 illustrates the influence of such Independent Effects with a conspicuous bubble.

Both local and independent effects, of course, work together to produce the final electricity prices. For modeling purposes, however, we conceive of these influences as follows. If in every period, loads and other local variables had "normal" values, what remained would be a path of electricity prices that must be the result of the independent effects. (The influence of independent effects, of course, could differ from "normal" conditions for all the reasons articulated in the previous paragraph.) To construct an electricity price series, therefore, it is valid to reverse this process. That is, it should be reasonable to apply the influence of loads, hydro generation, and a natural gas price to values representing the Independent Effect to obtain the resulting electricity price.

Unfortunately, we are not quite finished, because we may still need to adjust for any energy imbalance. The section "The Influence of Resource-Load Imbalances," beginning on page P-71, discusses this adjustment issue.

The process just described is the one that the regional model uses to produce an electricity price series. The next discussion focuses on the construction of the prices associated with Independent Effects. The subsequent discussion outlines the incorporation of influences for local hydro, load requirements, and natural gas prices. Forced outages influence prices to the extent that they affect energy imbalance.

Principal Factors

Dec of Cal Year	offset	linear
Weight	1.000	1.000
Value		
2003	0.50	0.07
2004	0.50	0.14
2005	0.50	0.21
2006	0.50	0.28
2007	0.50	0.35
2008	0.50	0.42
2009	0.50	0.49
2010	0.50	0.56
2011	0.50	0.63
2012	0.50	0.70
2013	0.50	0.77
2014	0.50	0.84
2015	0.50	0.91
2016	0.50	0.98
2017	0.50	1.05
2018	0.50	1.12
2019	0.50	1.19
2020	0.50	1.26
2021	0.50	1.33

source: L28_P.xls

Figure P-48: Principal Factors for the Independent Component of Electricity Price

	R	S	T
99	29.35806	0.409331	11.08000719
100	22.56537	0.26246	10.60643151

Cell	Formula		interpretation
R101	=$R$99	wait_1	start time of jump 1
S101	=R101+ IF($S$99= 0,0,12/$S$99)	wait_1+ 12/size_1	end time of jump 1
T101	=$S$99	size_1	size_log xfr jump 1
U101	=S101	end time of jump 1	start time of recovery 1
V101	=U101+ S101*EXP(T101)	end time of jump 1 + duration recovery 1	end time of recovery 1
W101	=-T101/10	-size_1/10	size_log xfr recovery 1
X101	=V101+ $R$100	end time of recovery 1+ wait_2	start time of jump 2
Y101	=X101+ IF($S$100= 0,0,12/$S$100)	wait_2 + 12/size_2	end time of jump 2
Z101	=$S$100	size_2	size_log xfr jump 2
AA101	=Y101	end time of jump 2	start time of recovery 2
AB101	=AA101+ Y101*EXP(Z101)	end time of jump 2 + duration recovery 2	end time of recovery 2
AC101	=-Z101/10	-size_2/10	size_log xfr recovery 2

R102 = IF(AND(R$46>$R101,R$46<=$S101),$T101,0)+ — jump_1
IF(AND(R$46>$U101,R$46<=$V101),$W101,0)+ — recovery_1
IF(AND(R$46>$X101,R$46<=$Y101),$Z101,0)+ — jump_2
IF(AND(R$46>$AA101,R$46<=$AB101),$AC101,0) — recovery_2

source: L28_P.xls

S102 identical, except S$46 instead of R$46
T102 identical, except T$46 instead of R$46

Figure P-49: Jump for Independent Component of Electricity Price

The Independent Term for Electricity Price

The model constructs the Independent Effect for electricity price in a manner very similar to the way it constructs natural gas prices and loads. See the section, "Stochastic Process Theory," above for details. Underlying strategic paths for average price[26] are the sum of principal factors, jumps, and optionally a stochastic adjustment. (The final regional model does not make use of the stochastic adjustment.) The model applies this path separately to on- and off-peak prices from the Council's long-term, electricity equilibrium price forecast to obtain corresponding prices for the regional model.

The principal factors appear in Figure P-48. The model permits up to two jumps, and the values and formulas for those jumps appear in Figure P-49. Both principal factors and jumps, in turn, rely on stochastic variables in assumption cells, the data for which appear in Figure P-50. The values for all of these objects ultimately originate from SAAC and Council staff judgments about what seem to be realistic and feasible futures. (See the section "Model Validation.")

Random Variables

	Type	Cell	Distribution	Parameters		
Jump 1	wait	{{R99}}	**uniform**	min 0	max 80	
	size	{{S99}}	**uniform**	min 0	max 2.5	
	duration	**<-------------------- not used -------------------->**				
Jump 2	wait	{{R100}}	**uniform**	min 16	max 36	
	size	{{S100}}	**uniform**	min 0	max 2.5	
	duration	**<-------------------- not used -------------------->**				
Principal Factors	offset	{{R94}}	**triangle**	min -1	mode 0	max 1
	linear	{{R96}}	**triangle**	min -0.83	mode -0.33	max 1.17

source: L28_P.xls

Figure P-50: Assumption Cells for Independent Component of Electricity Price

The Influence of Loads, Natural Gas Price, and Hydro Generation

The BMH study [17] provides the foundation for estimating the influence of loads, hydro generation, and natural gas price, on Mid-C electricity price. This study identified a regression equation for electricity price against these other influences. The equation, of course, is only accurate for the specific series of electricity prices and values of local variables assumed in the study. One difficulty with this approach, however, is that we

[26] Here average price refers to period (hydro quarter) average, across on- and off-peak hours. This is synonymous with "flat" market prices, where the average is with respect to on- and off-peak hours in whatever period is under discussion.

assume electricity prices to some extent independent from these other factors. The sensitivity to each of the influences, however, is implicit in the regression equation. By taking the difference between regression equations corresponding to two Independent sets of Independent variables we obtain a difference between two electricity price series. If we interpret this has the difference in electricity price due to changes in assumptions about the independent variables, we obtain the result we need.

The BMH model is of the form:

$$\ln(P_e(t)) = \alpha_0 + \alpha_1 \ln(P_g(t)) + \alpha_2 L(t) + \alpha_3 H(t) + \varepsilon(t)$$

where

$P_e(t)$ is electric price (\$/MWh) over interval t

$P_g(t)$ is gas price (\$/MMBTU) over interval t

$L(t)$ is peak load (MW) over interval t

$H(t)$ is hydrogeneration (MWa) over interval t

$\varepsilon(t)$ is an error term with a specified ARMA structure, having zero mean

α_i are constants determined by a statistical estimation technique, such the effect of weekday

Coefficient		Value
α_1	ln(Sumas price $/MMBTU)	4.40E-01
α_2	Max Load (MW)	4.38E-05
α_3	Hydro (MWa)	-1.34E-05

Figure P-51: Electricity Price Sensitivity Coefficients

Given three specific series $P^*_g(t)$, $L^*(t)$, and $H^*(t)$, this model predicts a specific $P^*_e(t)$. Given a distinct, arbitrary series $P_g(t)$, $L(t)$, and $H(t)$ and the associated, predicted $P_e(t)$, we have the following description of differences in electric price, given differences in the independent variables.

$$\begin{aligned}\ln(P_e(t)) - \ln(P^*{}_e(t)) &= \alpha_0 + \alpha_1 \ln(P_g(t)) + \alpha_2 L(t) + \alpha_3 H(t) + \varepsilon(t) - [\alpha_0 + \alpha_1 \ln(P^*{}_g(t)) + \\ &\qquad \alpha_2 L^*(t) + \alpha_3 H^*(t) + \varepsilon^*(t)] \\ &= \alpha_1\left[\ln(P_g(t)) - \ln(P^*{}_g(t))\right] + \alpha_2\left[L(t) - L^*(t)\right] + \alpha_3\left[H(t) - H^*(t)\right] + \varepsilon'(t)\end{aligned}$$

where

ε' is an error term with the same properties as ε and ε^*

We note several things. First, we have lost the constant coefficient, alpha zero. Second, the price of electricity does not appear on the right-hand side of this equation. The sensitivity of electric price to our independent variables does not depend on the absolute electric price.

Now, handed another series $Q_e(t)$ that shares the same sensitivity as $P_e(t)$ to our independent variables, we would predict $\ln(Q_e(t))$-$\ln(Q^*_e(t))$ would be described by the right-hand side of the preceding equation, where $Q^*_e(t)$ represents the value of $Q_e(t)$ when the perturbations of the independent variables are all zero.

The last step, then, is to take $Q^*_e(t)$, $P^*_g(t)$, $L^*(t)$, and $H^*(t)$ as the expected values of the electricity price, gas price, loads, and hydrogeneration values the regional model begins with, before accounting for the effect of the last three variables on the first. This gives us a means of forecasting electricity price $Q_e(t)$ given our assumed expected values for the four variables and excursions in the three independent variables. By taking the exponent of both sides,

$$\ln(P_e(t)) - \ln(P^*{}_e(t)) = \alpha_1 \ln(P_g(t)) + \alpha_2 L(t) + \alpha_3 H(t) - \alpha_1 \ln(P^*{}_g(t))$$
$$- \alpha_2 L^*(t) - \alpha_3 H^*(t) + \varepsilon'(t)$$

implies

$$P_e(t) = P^*{}_e(t) \cdot \frac{1}{c} \cdot P_g(t)^{\alpha_1} \exp\{\alpha_3 H(t) + \alpha_2 L(t)\} \qquad (13)$$

where

$$c = P^*{}_g(t)^{\alpha_1} \exp\{\alpha_3 H^*(t) + \alpha_2 L^*(t)\}$$

Note in particular that equation (13) consists of the product of three terms, the unadjusted electricity price, a term of the form

$$P_g(t)^{\alpha_1} \exp\{\alpha_3 H(t) + \alpha_2 L(t)\}$$

and a term that corresponds to the reciprocal of this expression, albeit with different values for certain variables. The section returns to the use of this expression later, at the discussion of "Worksheet Function and Formulas," below.

The Influence of Resource-Load Imbalances

After taking into account local influences, such as natural gas price, the resulting electricity price may prove to be infeasible from the standpoint of energy balance. The portfolio model assumes that dispatchable resources respond to market prices for electricity.[27] When a power system is unconstrained by transmission or other

[27] Strictly speaking, the assumption is that dispatchable resources respond to some explicit, widely visible signal of generation value. In the world before price deregulation, the measure of merit was "system lambda," which indicated the variable cost of generation on the system. Regulators among others sometimes refer to this concept as the "avoided cost." Economists refer to this kind of value as a "shadow price." It simply represents a means for assigning value to alternative means to meeting system

import/export limitations, one typically does not have to worry about whether a given market price is somehow infeasible. Higher prices simply mean more generators will run.

If a lot of new generation capacity arrives in the region, the region produces more energy at the same wholesale electricity market price level. Now, if loads are unchanged and exports are constraining, prices must fall to balance demand. Electricity prices are neither completely independent nor completely dependent of other variables. If the price is high, the resulting generation, after exports, may be surplus to requirements. Energy must be conserved, however: energy consumed must equal energy produced. In this example, the price must fall until the situation becomes feasible. The situation will be feasible when generation equals loads plus exports. Similarly, if the price is high, the resulting generation, after imports, may be inadequate for our requirements. The price must rise.

The Resource-Responsive Price (RRP) algorithm in the regional model finds a price that balances the system's energy. It does this by iteratively adjusting the price. Appendix L, in the section "RRP Algorithm," beginning on page L-51, describes this process in detail. Although this adjustment is made infrequently, keep in mind that it may be necessary and is part of the model logic. The RRP adjustment is also the principal means by which the model captures the influence of surplus and deficit resources and of forced outages.

The Application to Decision Criteria

The regional model makes extensive use of spot electricity prices for estimating forward electricity prices and future spot prices. The philosophical basis for this choice is the observation that forward prices and estimates of future spot prices generally track existing spot prices, as discussed in the section "Gas Price" and illustrated in Figure P-39. For forward electricity prices, the argument received fortification in March 2003, when FERC staff released their final analysis of "Price Manipulation in Western Markets," which features a section on "The Influence of Electricity Spot Prices on Electricity Forward Prices."[22] After examining prior analyses and studying the relationship between the prices, the report concludes "the forward power contracts negotiated during the period 2000-2001 in western United States were influenced by then-current spot prices, presumably because spot power prices influenced buyers' and sellers' expectations of spot prices in the future."

Because the horizon that a planner must consider depends in a sensitive fashion on the particular decision, technology, or power plant type she is considering, the role of electricity prices in each decision criterion differs. For this reason, Appendix L addresses their role in each specific criterion. (See section "Decision Criteria," beginning on page L-80 of Appendix L.)

requirements or the requirements of others. In describing the portfolio model, all of the arguments work if one substitutes these identical concepts for that of deregulated market price for electricity.

In all cases, an average of current electricity prices over some brief history determines the influence on the decision criterion. When this section turns to "Worksheet Function and Formulas," it will identify the specific average and describe its formula.

The Application to Load Elasticity

Load elasticity played an important role in the history of the Council. Arguably, it was a failure to recognize load elasticity that was responsible for some of the region's planning failures in the 1970s and was therefore the impetus for creating the Council.

Despite the prominence of the issue of load elasticity, the first versions of the regional model did not attempt to address it. The primary reason for this is that the effect of load elasticity is small relative to the load uncertainty that the model already incorporated. That is, because the regional model must already address futures where loads are much lower than could be accounting for price elasticity alone, it would seem unnecessary to include this smaller influence.

At the SAAC meetings where the Council staff presented the representation of load behavior, however, several of the participants felt uncomfortable that there was no separate accounting for this effect. Ultimately, the Council staff agreed that if for no other reason than to simplify the communication around treatment of load, it would be easier to include price elasticity explicitly.

Dr. Terry Morlan, who has prepared prior Council load forecasts, provided the basic characterization of price elasticity [**19**]. As we use the expression here, price elasticity of load is the change in load induced by a change in price over some specified time period.

$$\varepsilon_{l,p} = \frac{\Delta L}{L} \Big/ \frac{\Delta P}{P} = \frac{\Delta L}{L}\frac{P}{\Delta P}$$

where L and P are the load and price, respectively, at the beginning of the period. His sources indicate that the price elasticity over five years, which has a value of about -0.1, is less than that over 20 years, which he estimated at closer to -0.4. He said these factors would correspond to non-DSI retail rates, not wholesale price, which typically contribute about half to rate change. For a single year, and using wholesale prices, -0.02 max would probably be a better figure for non-DSI loads. To understand the impact of this selection of values, examples may be helpful. A doubling in prices, say from \$30/MWh to \$60/MWh, well in line with changes the region has seen in the last couple of year, predict almost a 20 percent reduction in loads over 10 years, about 3,600 MW. A one-year shock like the 2000-2001 energy crisis, where annual prices approached \$300/MWh would result in a similar change.

While at first glance, these seem comparable to changes the region has witnessed, in fact most of the change in loads corresponding to the 2000-2001 energy crisis is attributable to DSI load changes. (The regional model captures DSI loads separately. See the section on the principles of DSI modeling under the section "Multiple Periods" of Appendix L.) This level of elasticity, therefore created unrealistic behavior – over-response of non-DSI load – in the regional model.

Another difficulty with modeling this level of elasticity in the regional model was that it seemed to create model instability. Feedback from load to price can create an undampened oscillation. High price can lower requirements load via elasticity, and low loads can depress electricity prices via the model's resource-responsive price (RRP) algorithm. One way to avoid this behavior is to use small elasticities, but without extensive study, it is not clear what the upper limit on the magnitude of the elasticities needs to be.

In the end, the model did incorporate load price elasticity, but the model caps their influence, and their magnitude is one-tenth of the original values. This section will return to formulas that implement the elasticity in the next discussion. The issue of how best to represent price elasticity, however, remains for now unresolved and potentially an area of research for the next plan.

Worksheet Function and Formulas

With these preliminaries, tracing the formulas in the sample workbook should be straightforward. As is the custom, the discussion begins with column {{AQ}}, December 2009 through February 2010.

This section deals with the East and West, on- and off-peak quarterly average prices. Energy, cost, and dispatch calculations use these, as well as the decision criteria and elasticity calculations. This section does not address the decision criteria, however, because each decision criterion uses electricity prices differently. Therefore, Appendix L addresses each specific criterion separately. (See section "Decision Criteria," beginning on page L-80 of Appendix L.) This section also does not describe the worksheet formulas for load price elasticity, because Appendix L addresses those as well. (See the discussion "Loads" under the section "Multiple Periods," beginning on page L-59 of Appendix L.)

We begin with the calculation of flat[28] prices. A number of decision criteria, e.g., the decision criterion for price-responsive hydro, use flat electricity prices. The calculation of the electricity prices in {{AQ 224}} is

=AQ$207*4/7 + AQ$219*3/7

which is the average of on- and off-peak prices for electricity west of the Cascades, weighted by the number of on-peak and off-peak hours.[29]

[28] "Flat" market prices are average prices, where the average is with respect to on- and off-peak hours in whatever period is under discussion.
[29] There are 1,152 hours on peak in a standard hydro quarter and 864 hours off peak. See Appendix L for more background about standard months and quarters. Then, for example, 4/7=1152/(1152+864).

Tracing backward, the on-peak price in {{AQ 207}} has the formula

=AQ$204*(1.01)

This is the on-peak price for electricity east of the Cascades, with a 1 percent adder for losses and wheeling costs. The off peak price, AQ 219, has an identical formula that points to the off peak price for electricity east of the Cascades.

If we continue to trace the on-peak price, {{AQ 204}} has the formula

=AQ203+AQ200

This is the price adjustment in {{row 203}}, plus the unadjusted price in {{AQ200}}. The price adjustment in {{row 203}} does not contain any formulas. The RRP algorithm writes the values in this row. Appendix L, in the discussion of "RRP" from the section on "Multiple Periods" describes how this algorithm works to produce a price adjustment that balances energy requirements with energy sources.

The unadjusted on-peak, east-of-Cascades price in {{AQ 200}} uses the formula

=MIN(250, AQ$104*AQ$191*AQ$197)

This formula caps the east-of-Cascades prices at $250 a megawatt hour. The Council chose this ceiling on electricity prices because it reflects the current limit imposed by the Department of Energy on west-wide prices in 2002.

The expression AQ$104*AQ$191*AQ$197 in the previous equation captures the influence of local hydro generation, loads, and natural gas prices on electricity prices. Referring to equation (13), the adjusted electricity price is the product of the unadjusted electricity price, times two factors of the form

$$P_g(t)^{\alpha_1} \exp\{\alpha_3 H(t) + \alpha_2 L(t)\} \qquad (14)$$

One of the factors is the reciprocal of this expression and includes parameters that describe "normal" values for hydro generation, loads, and natural gas prices. The other factor has these values for the particular future. In the workbook model, the value in {{AQ104}} is the unadjusted electricity price. The term {{AQ 191}} has the form in equation 14 with the values for hydro generation, loads, and natural gas prices from the current future. The term {{AQ 197}} has the reciprocal of the form in equation 14, with the values for expected hydro generation, base case loads, and base case natural gas prices. In the following, the section first traces the construction of the value in {{AQ 197}}. It then traces the value in {{AQ 191}}, and finally it proceeds with the construction of the unadjusted electricity price in {{AQ 104}}.

The formula in {{AQ 197}} is

=1/AQ$194^0.44/EXP(0.000045*AQ$195-0.000014*AQ$196) which the reader will recognize as the constant 1/c in equation 13, page P-71. That is, {{AQ 194}} just points to the median forecast of natural gas prices in {{row 53}}. The cell {{AQ 195}} reconstructs the on-peak, west-of-Cascades load by multiplying the median load forecast by the on peak multiplier 1.14. (See discussion of this multiplier on page P-39, leading up to Figure P-25.) The value in cell {{AQ 196}} is the average on peak hydro generation for that period. The values in {{row 196}} are from reference [**20**].

The formula in cell {{AQ 191}} is

=AQ$178^0.44*EXP(0.000045*AQ$183-0.000014*AQ$188)

which is essentially identical, except that it references the values for hydro generation, loads, and natural gas prices that manifested this particular modeling future.

Hourly Behavior

The regional model assumes a lognormal standard deviation of hourly electricity prices that are 10 percent of the respective on- and off-peak quarterly averages. This means, for example, that if the average on-peak electricity price over the hydro quarter is $35/MWh,

- 99.7 percent of the hourly on-peak prices would fall below $47.25,
- 95.4 percent of the hourly on-peak prices would fall below $42.75,
- 68.3 percent of the hourly on-peak prices would fall below $38.68,
- 31.7 percent of the hourly on-peak prices would fall below $31.67,
- 4.6 percent of the hourly on-peak prices would fall below $28.66, and
- 0.3 percent of the hourly on-peak prices would fall below $25.93

The distribution of prices is not symmetric because of the nature of the lognormal distribution. That is, there is greater up-side variation than downside variation. It is also true that, while there is substantial variation in monthly and quarterly prices, daily prices correlate with monthly prices, and hourly prices correlated to daily prices. There is more information available, and therefore more price variation seen, on the longer time scales.

The last section of this chapter will address the correlations of hourly electricity price with those of other variables, such as natural gas price and loads.

Comparison with the Council's Electricity Price Forecast

The Council's electricity prices used in the final Fifth Power Plan and regional model L28 are from work that Council staff completed on October 21, 2004. (See Reference [**21**].) This section begins with a comparison of the Council's forecast with the independent term of the electricity price. Because this independent term represents the electricity price generated by the model before adjustments necessary to restore supply-demand balance, it is, in a sense, more directly comparable to the Council's price forecast. The final prices that resources see, however, can differ dramatically due to such adjustments. Therefore, the section also presents a statistical characterization across futures of the final, adjusted on- and off-peak prices for the Council's recommended resource plan.

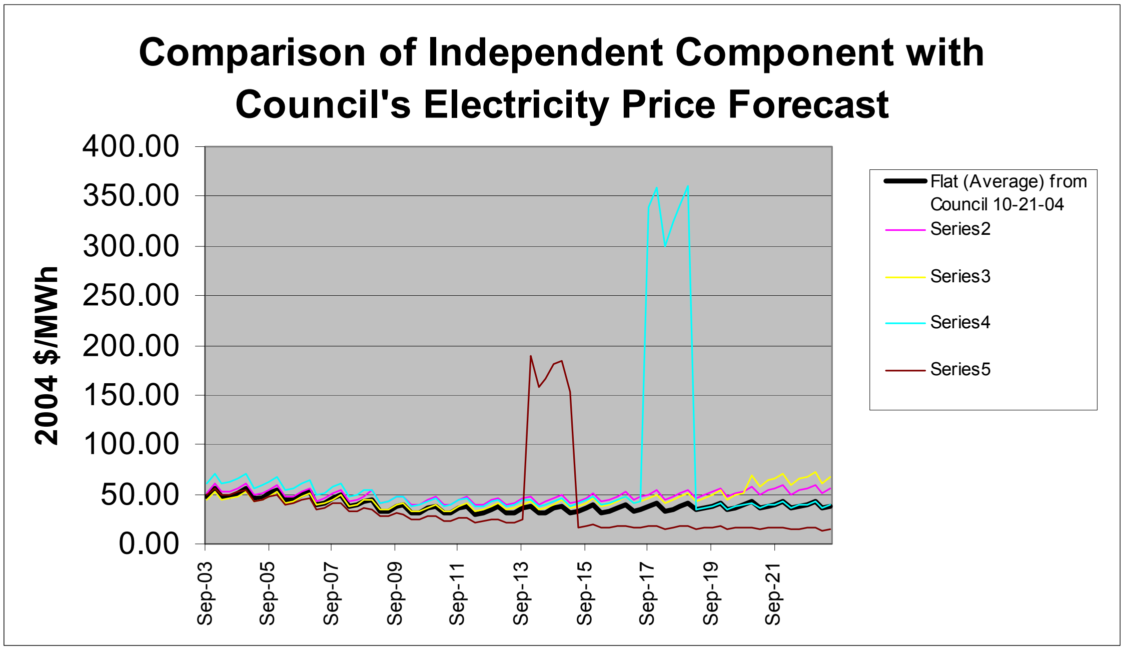

Figure P-52: Comparison of Independent Term Futures with Council's Electricity Price Forecast

The methods of principal factors, jumps, and specific variance described earlier produce the independent term of the electricity price. These use the Council's forecast as a median forecast. In Figure P-52, four random price futures appear along with the

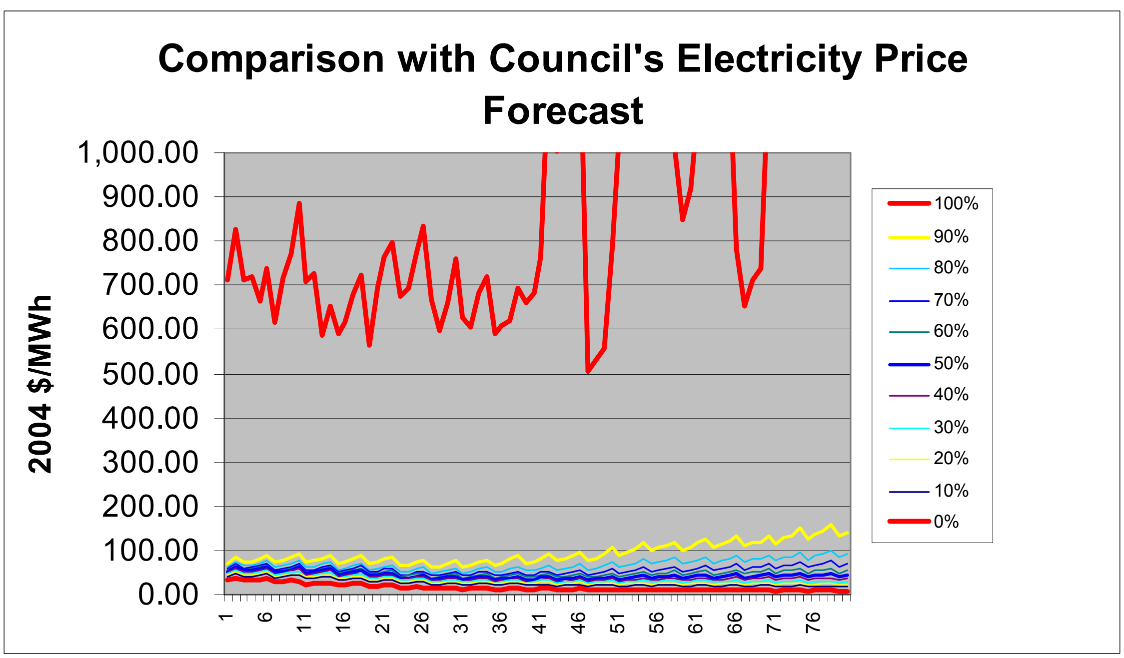

Figure P-53: Statistical Characterization of Independent Term Futures

Council's forecast (the heavier line). This figure presents the average of the Council's forecast over each quarter, on- and off-peak.

There are price series both above and below the Council's forecast, but two of the forecasts have jumps that last a couple of years. To get a more representative idea of the likelihood of these excursions, a statistical representation is helpful. Figure P-53 shows the price deciles for the 750 futures. It is clear that prices above \$150/MWh (2004\$) are rare, occurring less than 10 percent of the time in each quarter, but their magnitudes can

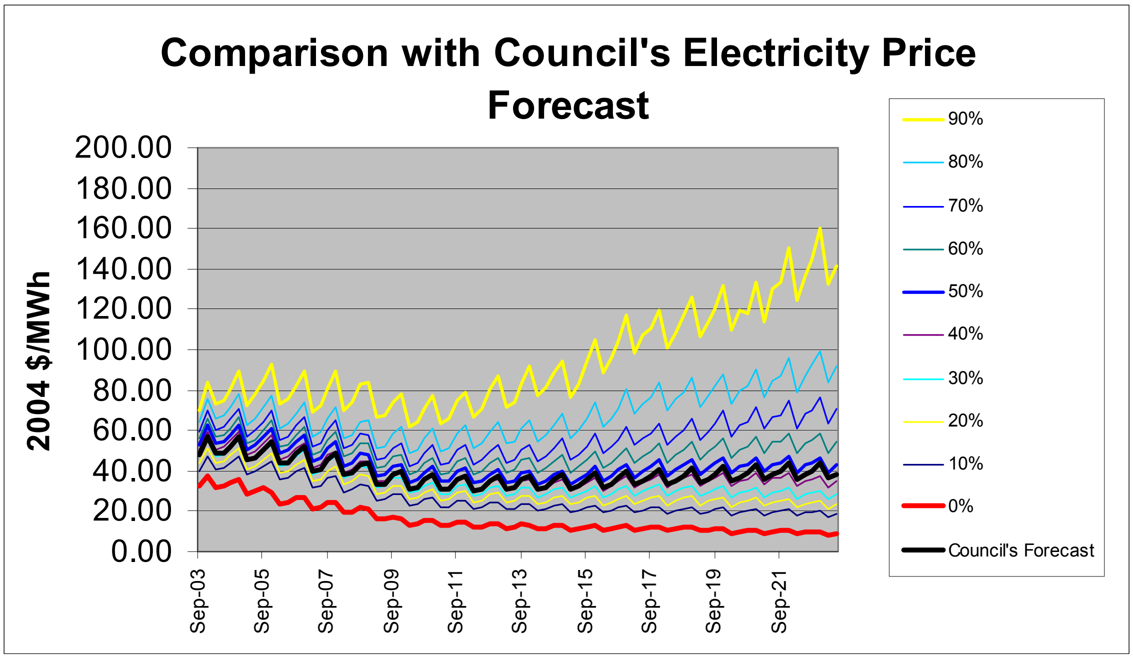

Figure P-54: Statistical Characterization of Independent Term Futures (Lower 90 Percent)

be quite significant. These low-probability events are largely due to the kinds of jumps illustrated in Figure P-52. Because the top decile dominates Figure P-53, the same information with that decile removed appears in Figure P-54.

One observation about the distribution of the regional model's electricity prices at this point is that the regional model's price median (50 percent decile) is slightly above the Council's forecast. The difference is small, less than \$6.29/MWh and averaging \$4.26/MWh. The reason for this difference is the influence of jumps. In early studies with electricity price, jumps had a recovery period that would cause their influence over time to average out. The recovery time was so long, however, that it precluded multiple jumps in a study. (One jump's recovery needed to finish before another jump could take place.) For this reason, the model uses a somewhat shorter jump recovery period, which produces a net lifting of median prices. This slight lifting effect, however, is not considered material to the analysis. One reason the effect is immaterial is that other influences on the independent term, described next, dwarf the lifting.

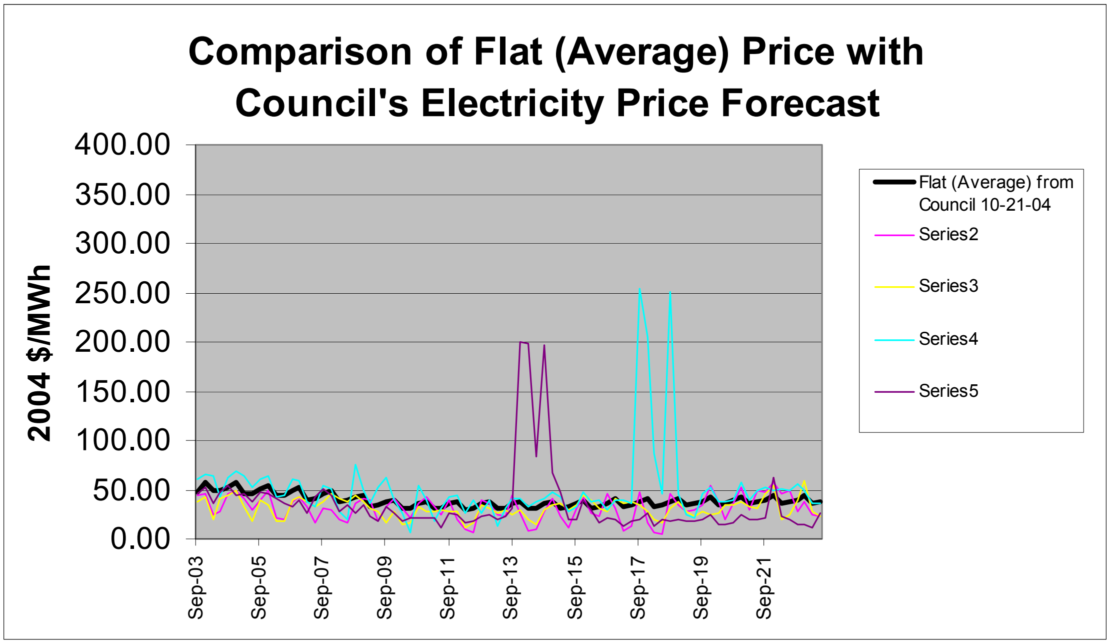

Figure P-55: Comparison of Futures after Adjustment with Council's Electricity Price Forecast

As described earlier, the influences of loads, natural gas price, and resource generation, including hydro generation, are significant in the regional model's electricity price. The effect is evident in Figure P-55 for the four futures appearing in Figure P-52. In Figure P-55, the prices are depressed in general from those in Figure P-52. This should not be too surprising. The recommended resource plan, to which these price futures pertain, has significant resources in most futures. The downward pressure on electricity due to surplus resources alone will produce this effect.

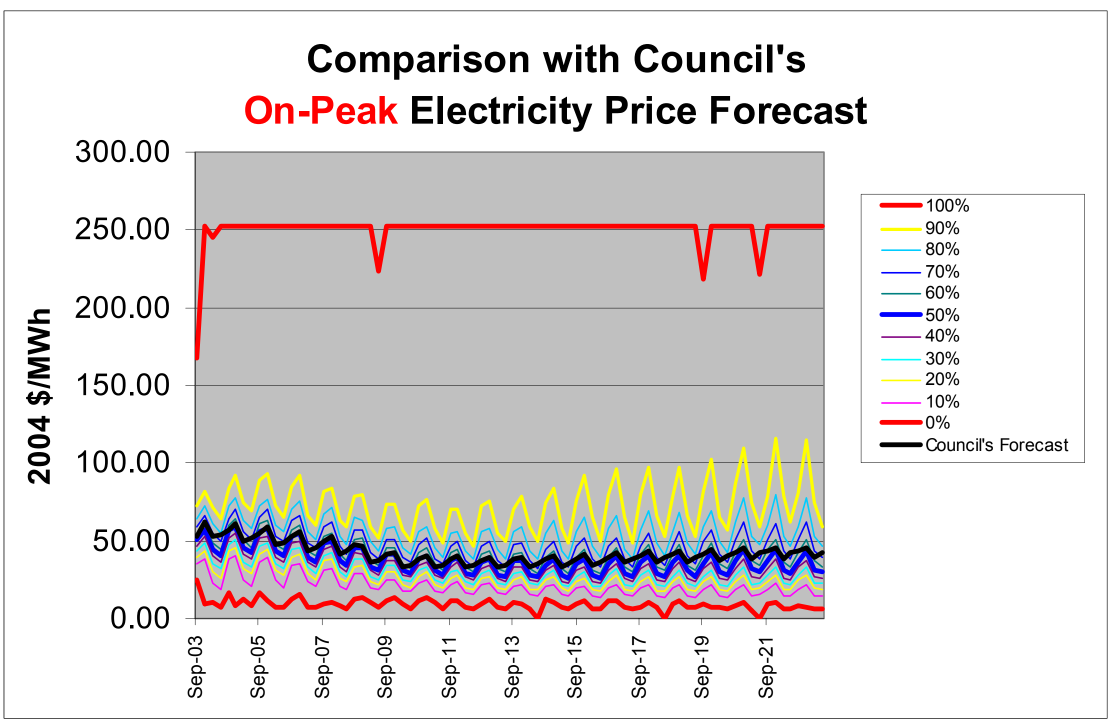

Figure P-56: Statistical Characterization of Futures' On-Peak Prices

A statistical comparison of the final on- and off-peak prices for the regional model to the Council's price forecast shows a similar pattern. While the median of independent term for electricity price is slightly above the Council's forecast, that for the regional model's on-peak price is slightly below that for the Council, as seen in Figure P-56. Another feature of the on-peak price distribution is that all prices are at or below $250/MWh. Actually, the ceiling price is slightly higher than $250/MWh, because the model assumes the cap applies to east of Cascades prices, and transmission costs cause the delivered price to west of Cascade loads to be higher. The reason for the ceiling is a cap imposed by the U.S. Department of Energy in June 2001.[30] The view of staff and advisors is that this cap, or something like it, is likely to remain in place for the foreseeable future.

In Figure P-57, the off-peak price deciles from the regional model appear next to the Council's off-peak power price. As expected, the deciles generally lie slightly below the corresponding on-peak price deciles.

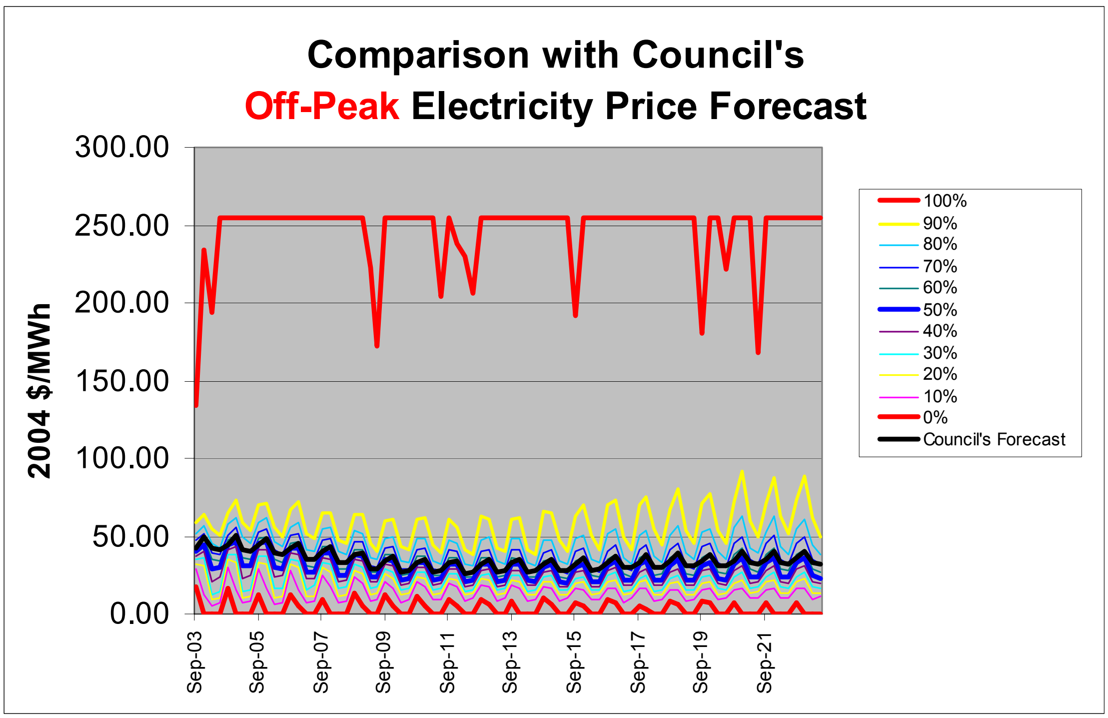

Figure P-57: Statistical Characterization of Futures' Off-Peak Prices

This concludes the discussion of electricity price and its associated uncertainty. A comparison of the regional model's prices with those of the Council's forecast shows some predictable differences. For the most part, however, there is agreement, and the behaviors of the regional model's price futures appear reasonable.

[30] See, for example, Federal Energy Regulatory Commission, "Commission Extends California Price Mitigation Plan for Spot Markets to All Hours, All States In Entire Western Region," news release, June 18, 2001, EL00-95-031, EL00-98-030 and - 033, RT01-85-001 and -033, EL01-68-000 and –001.

Forced outage rates

Unplanned outages affect the availability of power plants. Although, by definition, planners cannot forecast when these outages may occur for a specific plant, an ensemble of power plants have predictable behavior over a sufficiently long time period. This behavior has permitted the power generation industry to acquire estimates of forced outage rates (***FOR***s) for various kinds of generation technology.

If *H* is the number of hours in a sufficiently large period, and *h* is the number of hours we expect a plant to be unavailable due to forced outages, the ***FOR*** is defined to be *h/H*. The period must only be large enough for the ***FOR*** to have predictive significance. Unfortunately, this tells us very little about the frequency or duration of forced outages. That is, even if a planner were using the same period as that on which the statistic is based, he cannot tell how long or how frequently a plant should be out of service. Of course, the period a planner would use would typically be smaller than that of the statistical sample, further muddying the water. Typically, the planner simply derates each period's energy by the ***FOR***. Unfortunately, this eliminates the risk of extended outages that would nevertheless be consistent with the statistical value.

The traditional approach to modeling forced outages statistically is to use a binomial distribution. The binomial distribution represents events that are independent of each other and when these events have fixed likelihood. For existing power plants, creating a stochastic variable with this distribution is relatively easy. For new power plants, however, the situation is more challenging in the regional portfolio model. As the number of identical power plants increases, the availability of the ensemble of power plants becomes more predictable. Because each new plant actually represents an ensemble of plants in the regional model, and because the number of plants, or cohorts, changes not only from plan to plan but from future to future, creating exactly the right distribution of energy duration is not easy.

Because of these considerations, the regional portfolio model uses a simpler approach than incorporating a binomial distribution. Energy deration due to forced outages is a random variable with a symmetric triangular distribution, and with an average (and most likely) value equal to the ***FOR*** (see Figure P-58). The generation technology determines the expected availability of each plant in the regional portfolio model. A fall 2003 reassessment of regional power plant outage rates [**22**] form the basis for the technology values. A summary of the regional model's final values of FOR appear in Figure P-59.

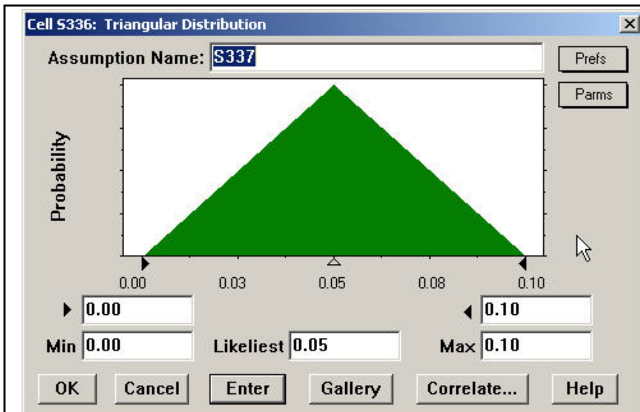

Figure P-58: Typical FOR Distribution (FOR=0.05)

In each future, the model makes a separate draw from the triangular distribution for each plant (or surrogate plant) for each hydro quarter. The energy of the plant over the period diminishes by a corresponding amount. For example, for the energy calculation for the plant "PNW West NG 5_006" in cell {{S339}}, one finds the references illustrated in Figure P-60. The reference to cell S336 in Figure P-60 is to this plant's FOR in this period. (The values in the assumption cell S336 happen to appear in Figure P-58.) As explained in Appendix L, the FOR must derate both the electric energy and the gas used.

SourcePType	Fuel Type	FOR
Hyd	Water	0.000
Wind	Wind	0.000
CCCT	MT gas	0.050
CCCT	PNW E. gas	0.050
CCCT	PNW W. gas	0.050
Biomass ST	Mill Residue	0.070
Coal	MT coal	0.070
Coal	N. NV coal	0.070
Coal	PNW E. coal	0.070
Coal	PNW W. coal	0.070
GT	PNW E. gas	0.070
GT	PNW W. gas	0.070
GTAero	PNW E. gas	0.070
GTAero	PNW W. gas	0.070
STCG	No. 2 FO	0.070
STCG	PNW E. gas	0.070
STCG	PNW W. gas	0.070
IC	No. 2 FO	0.080
IC	PNW E. gas	0.080
IC	PNW W. gas	0.080
Nuclear	WNP-2 nuclear fuel	0.090

Figure P-59: FOR Rates

For some plants, the model does not use this stochastic representation. For certain classes of resources, the model uses a simple capacity de-ration instead. These plants are those that are small, and would make trivial contribution to forced outages, and new units. For new units, the issue is the potential complexity, described above, associated with the changing number of units in the ensemble. Rather than introduce another source of complexity into the model that could influence the choice of new resources, the model takes the simplest approach.

S339 = {=SpreadOption($P339, S$46,S$207-$R337,S$68+0.059*S$74,(1-S336)*1152*$S335,(1-S336)*1152*$S335*$I335,1,0,0,0,$R$201,$R$55,$T$14)}

	Q	R	S	T	U	V	W	X
333		Resources						
334	PNW West NG 5_006							
335	Capacity_ID: PNW West NG 5 Cap	487	503	390	432			
336	Expected FOR	0.04746901	0.040443308	0.053736452	0.038572008	0.070589275	0.028776688	0.066424
337	ariable_Cost ($/MWh): PNW West NG 5 VOM	3.02						
338								
339	Energy(MWh)	169103.1	1223.0	0.0	0.0	94317.0	180308.7	27320
340	Cost ($M)	-0.8	0.0	0.0	0.0	-0.3	-0.9	-
341	Capacity Factor (%)	30.2%	0.2%	0.0%	0.0%	16.8%	31.1%	48.
342								

Figure P-60: Calculations in Cell {{S339}}

Where the model uses the stochastic representation, the same availability is used both on- and off-peak. This makes sense, as an outage would not discriminate between these subperiods.

Finally, we point out that FOR is the only aspect of a future that is ***not*** computed in the range of the worksheet reserved for such calculations,[31] although it could be, and arguably should be. Keeping it with the resource facilitates review and verification of resource performance.

[31] The discussion "Logic Structure of the Portfolio Model" on pages P-14 ff identifies the specific range.

Aluminum Price

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 prices of aluminum and 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 an 80 percent likelihood of all aluminum plants closing during the forecast period.

To represent aluminum price futures, the Council evaluated several approaches, and the approach that most closely matched historical price patterns is a geometric Brownian motion (GBM) process with mean reversion. Aluminum prices do not exhibit the seasonal shape that natural gas and electricity prices possess. Instead, they tend to wander away from a trend with quasi-cyclical excursions of varying regularity, as illustrated in Figure P-61. (See Reference [**23**].)

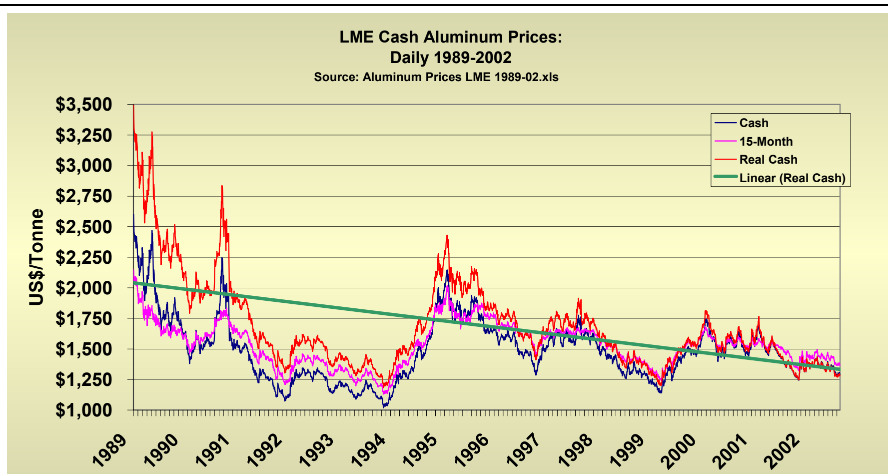

Figure P-61: Historical Aluminum Prices (2002 $/metric Tonne)

In Figure P-61, a linear regression line emphasizes the downward trend in aluminum prices that has been evident over the last 20 years or so.

The section "GBM with Mean Reversion," beginning on page P-23 describes the mathematical principles of the stochastic process. The regional model workbook implements the equations as follows (see also Figure P-62):

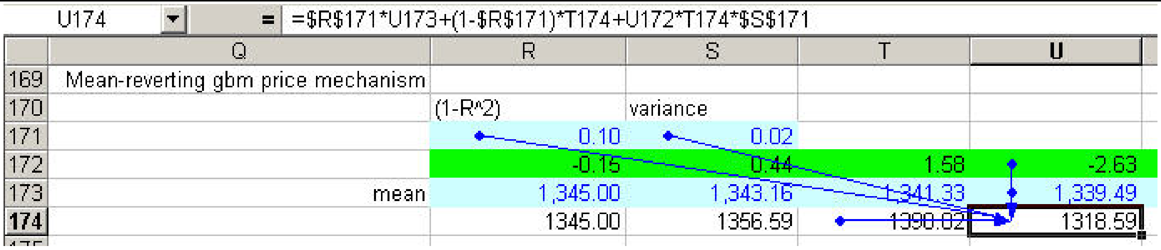

Figure P-62: Workbook Formula for GBM with MR

The formula

=$R$171*U173+(1-$R$171)*T174+U172*T174*$S$171

replicates the equation from the section "GBM with Mean Reversion,"

$$p_t + dp_t = p_t + a(b - p_t)dt + \sigma p_t dz = abdt + (1-a)p_t dt + dz \cdot p_t \sigma$$

where

p_t is stochastic variable in question

dp_t is the change in p_t from the previous step

dz is a drawn from a N(0,1) process

dt is the step size, which has value 1 for discrete processes

a is constant which controls the rate of reversion

b is the equilibrium level

σ is the standard deviation of p_t

(Note that the label "variance" in cell {{S 170}} of Figure P-62 is incorrect. The value in cell {{S 171}} is the standard deviation of the log-transformed aluminum prices. See Reference [**24**].) A Crystal Ball assumption cell provides an underlying Weiner process *dz* with the appropriate distribution (Figure P-63).

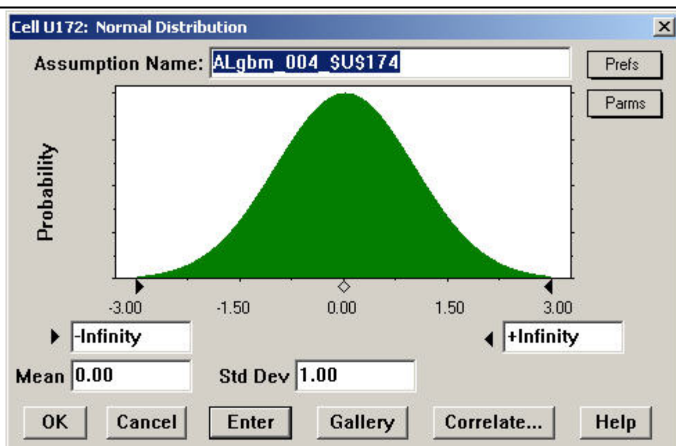

Figure P-63: Assumption Cell {{U 172}}

Figure P-64 illustrates the behavior of this process representation. The individual futures exhibit the same kind of irregular walk around the mean that does the historical data. The values are smoother, however, as expected from quarterly averages.

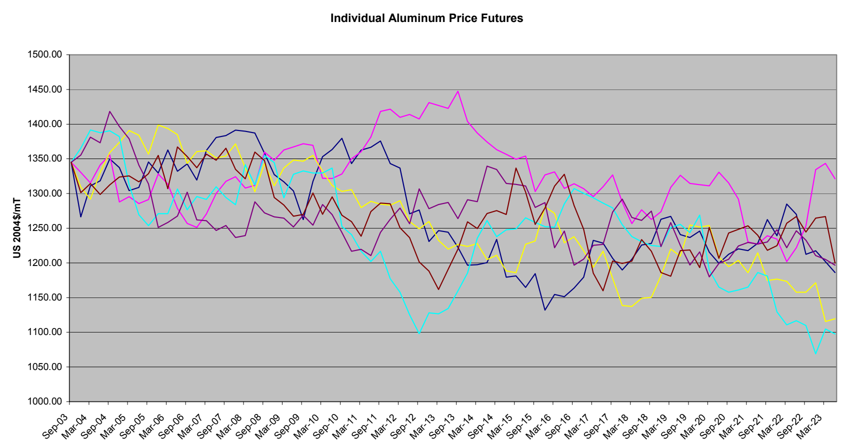

Figure P-64: Selected Aluminum Price Futures

Figure P-65 provides additional, statistical description of the aluminum price futures. It shows the quarterly deciles, plotted against the periods in the study. It is evident that the mean to which prices are reverting is trending down, consistent with the historical price behavior. The mean price descends from the May 2004 price of $1345/mT to $1200/mT (2004 $) by the end of the study (see Reference [**25**]).

CO_2 tax

A significant proportion of scientific opinion holds that the earth is warming due to atmospheric accumulation of greenhouse gases. The increasing atmospheric concentration of these gases 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 gases. 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 will occur without cost, future climate-control policy is 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 of using a proxy tax on power plant generation and economic value 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.[32] (See Reference [**26**].) The probability of any such tax during the forecast period is 67 percent. If enacted, the value for the carbon tax has a uniform distribution between zero and $15 per ton if it is enacted between 2008 and 2016; and between zero and $30 per ton if enacted thereafter (2004$). These draws are independent of other parameters, although other stochastic variables, like a production tax credit, depend on CO_2 tax. The two sections following this one describe the relationship.

The probability distribution of this stochastic variable was the subject of intense debate during the development of the plan. While authoritative studies[33] supported carbon tax as high as $100/ton$_{CO2,}$ the final values had as much to do with the principle of "thresholding" as with what the likely values might be. Specifically, increasing the CO_2 tax had little effect on the plans lying on the efficient frontier. (See the discussion of CO_2

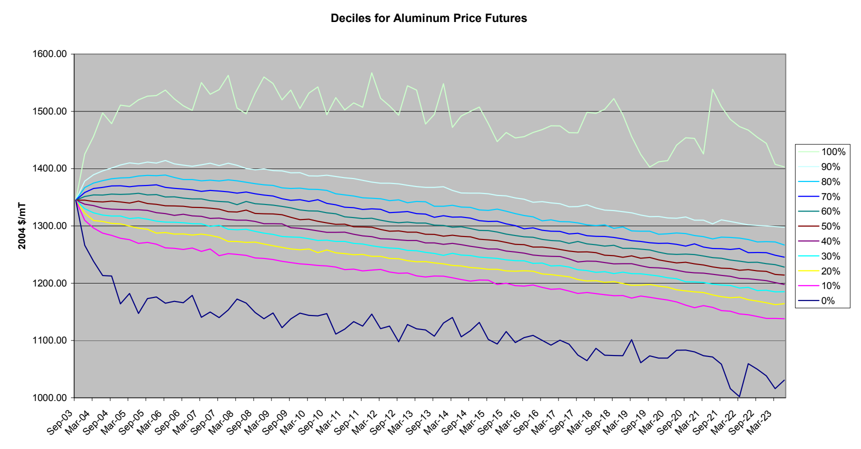

Figure P-65: Deciles for Quarterly Aluminum Prices

sensitivities in the section, "CO_2 Policy.") Using higher values would therefore have only token value and would render the model results questionable among those who do not believe higher taxes are likely. Few participants, on the other hand, could argue for smaller probability and magnitude of tax. One third of the futures had no tax at all.

As a basis for comparison, note that PacifiCorp is heavily reliant on coal-fired power, the cost of which would be especially sensitive to carbon tax. PacifiCorp has little motivation to argue for high CO_2 tax rate. The expected value tax rate in the regional

[32] At a May 20, 2003 meeting held in the Council's main Portland office, experts on carbon tax were reluctant to speculate on the likelihood or magnitude on any carbon tax. They did appear to be in agreement, however, that if the United States enacted a carbon tax, it would require the support of the executive branch of the U.S. Government. The change would likely arrive, therefore, with a change in administration.

[33] See, for example, ***MIT Joint Program on the Science and Policy of Global Change***, *Emissions Trading to Reduce Greenhouse Gas Emissions in the United States, The McCain-Lieberman Proposal*, Report 97, June 2003, available at http://mit.edu/globalchange/www/reports.html#r100

model until 2020, however, is less than the expected value forecast that appeared in PacifiCorp's 2004 IRP [**27**].[34]

Given what some might consider such a low expected CO_2 tax rate assumption, did the tax matter? It did, but for reasons that may require explanation. First, in a risk model, the extreme values are as important as the expected value, and the high end of the range exceeded what some would consider likely, as it should. Second, what drives much of the resource selection in the regional model is not a single source of risk, such as CO_2, but combinations of risks. Each independent source of risk adds to the expected net cost of the resource. For coal-fired power plants, for example, lack of planning flexibility, capital cost exposure, and load uncertainty were also issues affecting economic feasibility. The CO_2 tax assumption merely contributed to the coal-fired power plant risk.

Figure P-66 has six of the first CO_2 tax futures, although in two of those futures no tax arrives. In each future, there is at most only one arrival of taxes, and it occurs as a step. This is, in fact, the way the regional model represents CO_2 tax in all futures. However, the *maximum* size of the step depends on the year it happens, as mentioned earlier.

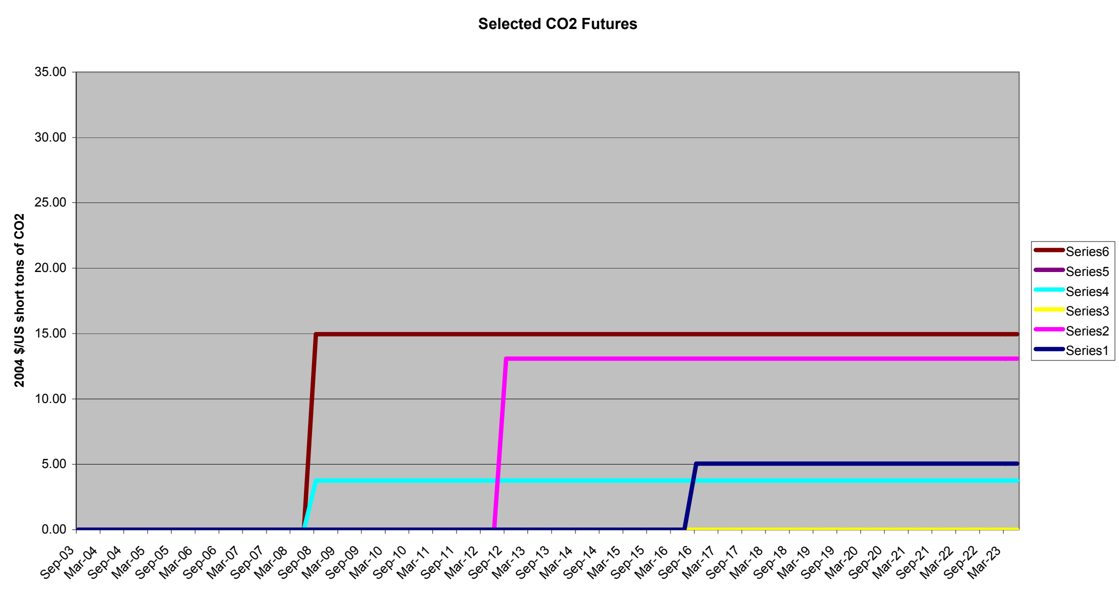

Figure P-66: Selected CO_2 Futures

Figure P-67 provides some descriptive statistics across periods. In addition to the deciles the reader has seen in prior illustrations, the graph includes the average CO_2 tax across all

[34] PacifiCorp used $8.00/ton$_{CO2}$ (2008 $) beginning in 2010 or about $7.38 in 2004 dollars using PacifiCorp's inflation assumption of 2.02 percent. Their study discounted this value in the first two years only to $3.69 (2004 $) in 2010 and to $5.54 (2004 $) in 2011. (See Table C.7, and supporting discussion in Appendix C, page 37 of the PacifiCorp 2004 IRP, Technical Appendix.) The regional model's expected tax rate grows and surpasses PacifiCorp's by less than $0.46 only in the last three years of the study.

futures. (A dotted line identifies the average.) One of the striking features of this graph is the non-appearance of the deciles below 40 percent. Those deciles all lie on the zero-tax line. On reflection, however, this is consistent with the earlier observation that approximately a third of the futures contain no tax.

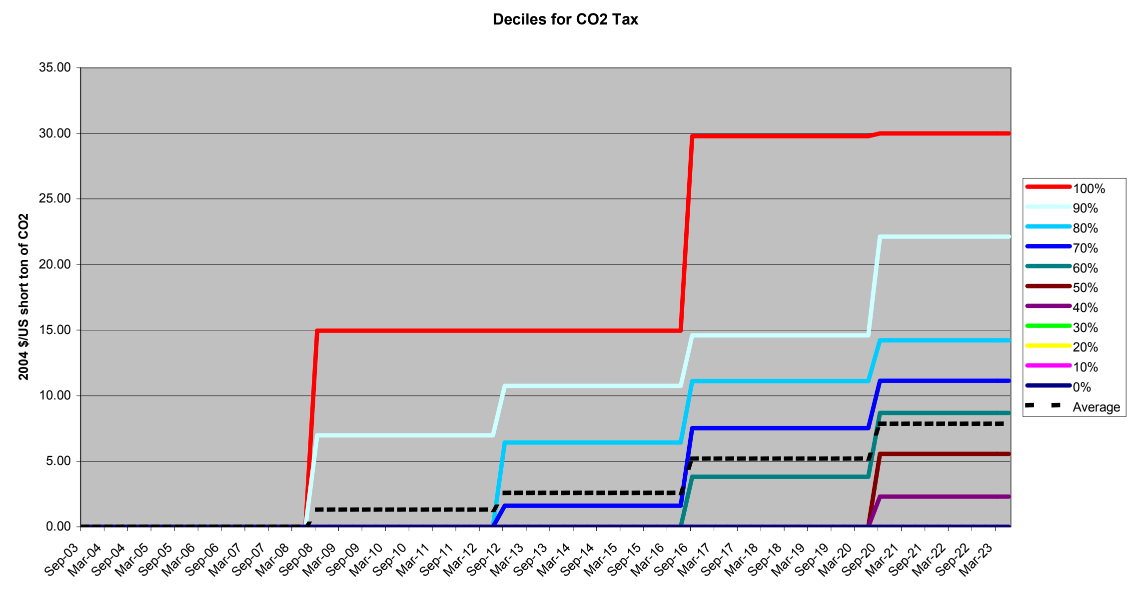

Figure P-67: CO_2 Statistics

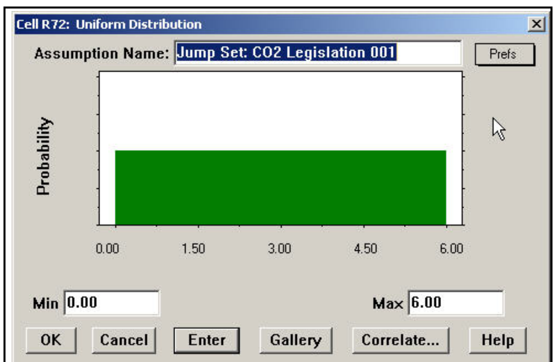

Figure P-68: Timing

To capture this behavior in the workbook, only two Crystal Ball assumption cells are necessary. The first one, {{R72}}, illustrated in Figure P-68, controls the timing of step. It is a uniform distribution from 0.0 to 6.0. The explanation for the range of this random variable becomes evident in a moment. The second assumption cell, {{S72}}, illustrated in Figure P-69, determines the size of the step.

The model first determines in which column any step takes place, as shown in Figure P-70. The formula in cell {{T72}}, for example, is

=IF(T$46>4+INT($R72)*16, $S72,0)

This formula compares the period ({{T46}}) to one of the values 4, 20, 36, 52, 68, or 84, which {{R72}} determines. That is, the value 1.009516835 belongs to the interval (1,2]. There are six such intervals, all with equal likelihood. The function converts the value to 1.0, which formula in {{T73}} converts to 20. (If the cell {{R72}} has value 6.0, then {{T73}} would be 100, which has probability zero.) These period values correspond to the period September through December of each election year. If the column's period number exceeds this value, it assumes the value in cell {{S72}}, which will determine the size of the step.

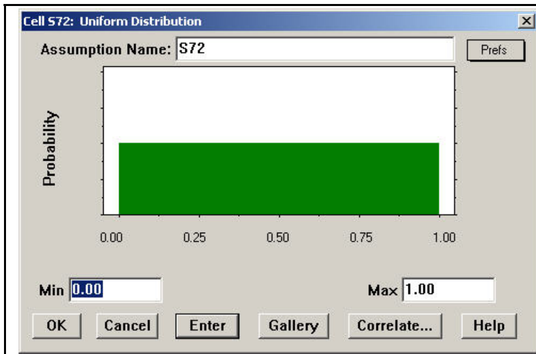

Figure P-69: Size

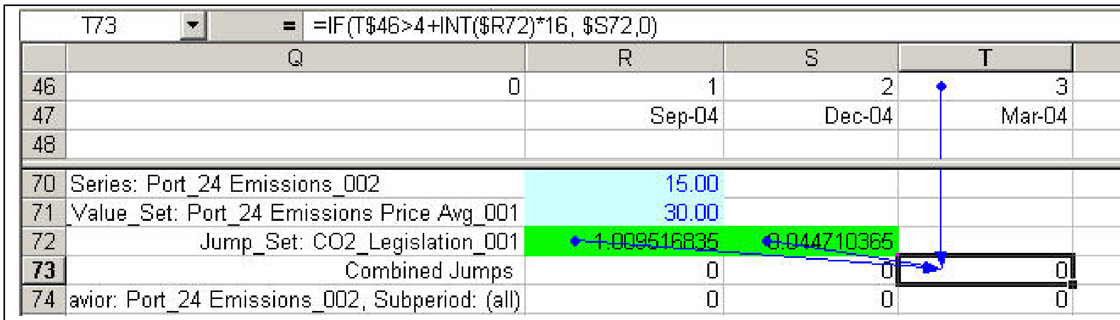

T73 = =IF(T$46>4+INT($R72)*16, $S72,0)

	Q	R	S	T
46	0	1	2	3
47		Sep-04	Dec-04	Mar-04
48				
70	Series: Port_24 Emissions_002	15.00		
71	Value_Set: Port_24 Emissions Price Avg_001	30.00		
72	Jump_Set: CO2_Legislation_001	1.009516835	0.044710365	
73	Combined Jumps	0	0	0
74	avior: Port_24 Emissions_002, Subperiod: (all)	0	0	0

Figure P-70: Normalized Step

At this point in the calculation, the values in {{row 73}} have the value in cell {{S72}} if they belong to periods after the first occurrence of any step. Otherwise, they have the value 0.0.

The task remaining for formulas in {{row 74}} is to properly scale these values to the real tax rate. The formulas are of the form

=IF($R$72<1,0,IF($R$72<3,T73*$R$70,T73*$R$71))

and Figure P-71 illustrates the references. The first "if" test prohibits any tax from appearing during the George W. Bush administration. This was a modification made later in the development of the model. It effectively decreases the probability of a tax in the study period. The second "if" test scales the range of the tax to $15/ton before 2016 and to $30/ton subsequently.

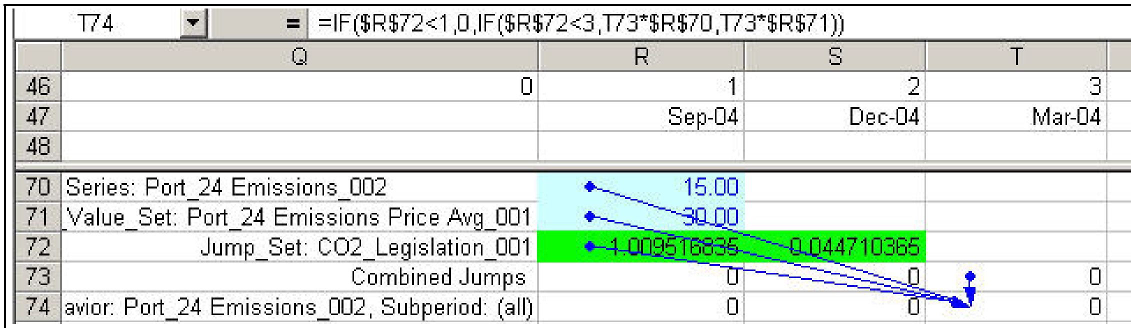

T74 = =IF($R$72<1,0,IF($R$72<3,T73*$R$70,T73*$R$71))

	Q	R	S	T
46	0	1	2	3
47		Sep-04	Dec-04	Mar-04
48				
70	Series: Port_24 Emissions_002	15.00		
71	Value_Set: Port_24 Emissions Price Avg_001	30.00		
72	Jump_Set: CO2_Legislation_001	1.009516835	0.044710365	
73	Combined Jumps	0	0	0
74	avior: Port_24 Emissions_002, Subperiod: (all)	0	0	0

Figure P-71: Nominalized Step

Production Tax Credits

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, Congress extended it to the end of 2005, retroactive to the beginning of 2004. In addition, in October, they extended the scope of qualifying facilities 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 a project's operation. The longer-term fate of these incentives is uncertain. The original legislation contains a provision for phasing out the credit as the cost of qualifying resources becomes competitive with electricity market prices. 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. Termination always takes place before the wind energy-cost forecast declines to 30 mills/kWh in 2034 (2004$). That is, there is never a modeling future where a PTC would extend 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[35]. 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.

A three-step process determines the PTC value that the regional model will use in a given future and period. In the first step, a formula like

=IF(T46>$R76,0,9.9)

in cell {{T79}} determines whether the wind plant should be commercially viable. Figure P-72 illustrates the references. The label in {{Q79}}, "PTC (after commercial

T79 = =IF(T46>$R76,0,9.9)

	O	P	Q	R	S	T	U
46			0	1	2	3	4
47				Sep-04	Dec-04	Mar-04	Jun-04
75							
76			PTC, JK workbook	104.3030438			
77			Integration	5.82	10.76		
78			Greentag Value	3.5	4.5		
79	**Conversion (#CO2/kWh)**		PTC (after commercial viability test)	9.90	9.90	9.90	9.90
80		**1.28**	Carbon Tax ($/MWh)	0.00	0.00	0.00	0.00
81			PTC (after CO2 tax effect)	9.9	9.90	9.90	9.90
82			VOM	1.13	1.13	1.13	1.13
83			Greentag Value	3.50	3.51	3.53	3.54

Figure P-72: PTC Calculation, Step 1

viability test)," is misleading. Federal politics would determine viability, and commercial competitiveness is one of several issues. As mentioned above, the PTC could go away almost immediately, if it became unpopular for any reason. The PTC may also outlive its original purpose if political or economic forces support retention. The distribution of a random variable describing this lifetime must therefore have some small, positive value in the near term and in years after renewables would become competitive.

[35] 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 9,000 BTU/kWh.

This model compares Council forecasts of wind generation fixed costs to its electric market prices to estimate when renewables would become competitive. In an outboard calculation (Reference [**28**]), staff estimated wind would achieve economic competitiveness in 2016. This assumes an electric price of $40/MWh in that year and wind generation costs that decline at about 1.7 percent per year (Reference [**29**]). Moreover, staff assumed the chance of the PTC surviving when wind generation cost fell to $30/MWh in 2034 would be nil, so the model uses a triangular distribution for the lifetime of the PTC. The year 2016 corresponds to the 52nd period (hydro quarter), so the

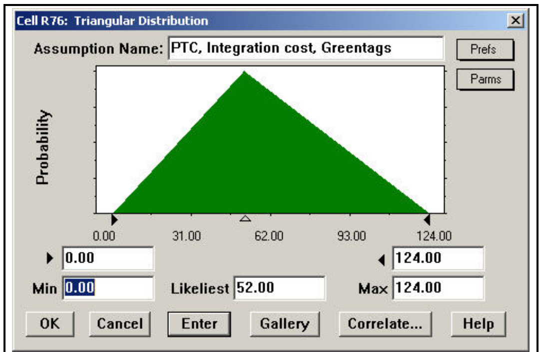

Figure P-73: Political Viability Distribution

distribution has 52 as its mode; the year 2035 corresponds to the 124th period, so that value determines the maximum value. Because the study only extends 80 periods, there is a substantial probability that the PTC does not disappear due to political non-viability during the study.

The formula in cell {{T79}} stipulates that if the period exceeds the value of the random variable, the PTC is zero; otherwise it has a real levelized value $9.90/MWh in 2004 dollars (Reference [**30**]). This value corresponds to the current credit of roughly 1.7 cents/kWh in year 2000 dollars, using Council assumptions for wind capacity factor and inflation. Staff elected not to make the PTC value a random variable and saw no compelling reason to assume this would either increase or decline over time.

The second step of the process to determine the PTC value in the regional model is an examination of any CO_2 tax in the period. The cell {{T80}} is typical and contains

=T74*$P$80/2

(Cell references appear in Figure P-74.) This formula converts the tax in $/US short ton (2004 $) to $/MWh using the value in {{P80}}.[35] The conversion factor is in pounds of CO_2 per kWh, so the conversion is

$/MWh = $/ton • tons/pound • pounds/kWh • kWh/MWh or
$/MWh = $/ton • pounds/kWh • 1000/2000

The last term in the product gives rise to the factor of two in the denominator of the formula in cell {{T80}}.

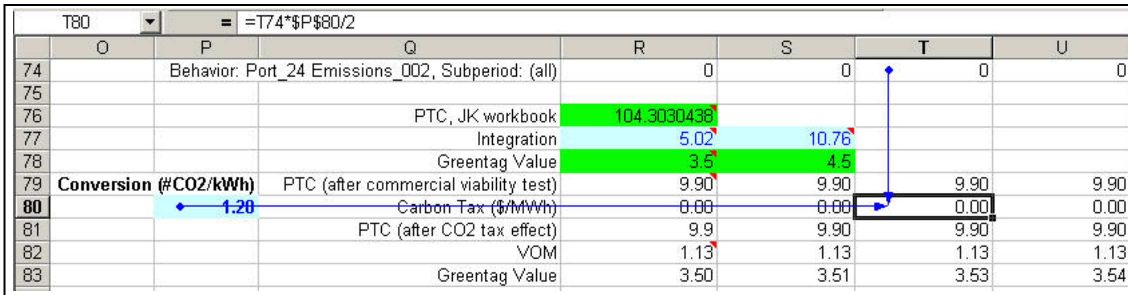

T80 = =T74*$P$80/2

	O	P	Q	R	S	T	U
74		Behavior: Port_24 Emissions_002, Subperiod: (all)		0	0	0	0
75							
76			PTC, JK workbook	104.3030438			
77			Integration	5.02	10.76		
78			Greentag Value	3.5	4.5		
79	Conversion (#CO2/kWh)		PTC (after commercial viability test)	9.90	9.90	9.90	9.90
80		1.20	Carbon Tax ($/MWh)	0.00	0.00	0.00	0.00
81			PTC (after CO2 tax effect)	9.9	9.90	9.90	9.90
82			VOM	1.13	1.13	1.13	1.13
83			Greentag Value	3.50	3.51	3.53	3.54

Figure P-74: Carbon Tax Effect (Step 2)

The third and final step of the process to determine the PTC value is the "PTC offset" due to any CO_2 tax. In the draft plan, the PTC went away in any future where any positive CO_2 tax occurred. The issue that arose between the draft and final plan was, "Would the PTC go away entirely even if the CO_2 tax were very small?" The problem was that the combined support for renewables could undergo a discontinuity, a net drop, if the CO_2 tax were very small. This struck the Council as unrealistic.

To address this matter, new logic provided for the remaining PTC to be a function of the magnitude of the CO_2 tax. Figure P-75 illustrates the PTC remaining. In terms of support for wind generation, the PTC corresponds to a \$15.47/ton CO_2 tax, given Council assumptions. With the new logic, if the CO_2 tax that arises is less than half of this, the PTC remains in place; if the tax is fifty percent higher than this, it disappears entirely. Between those values, it declines dollar for dollar with the tax rate. Figure P-76 shows the combined advantage relative to gas-fired generation provided by the CO_2 tax and the PTC. Note that no discontinuity exists for the combined support.

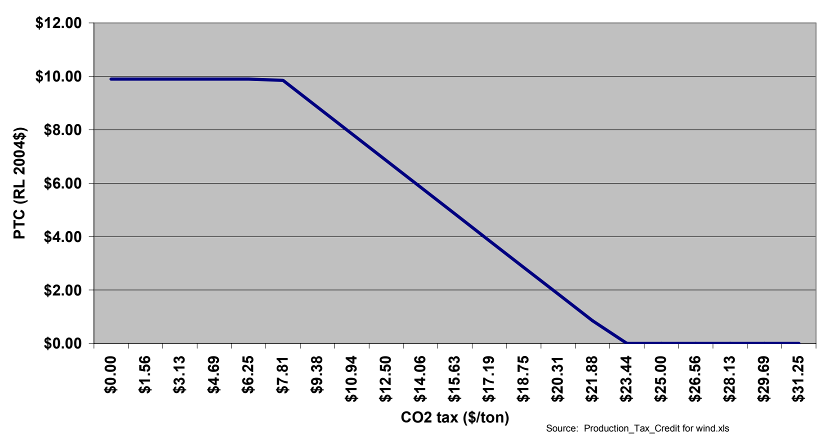

Figure P-75: PTC as a Function of CO_2 Tax

Figure P-77 shows how the workbook implements the PTC with formulas, such as that in cell {{T81}}. Again, if the tax has become politically non-viable, the PTC from cell {{T79}} in this example is zero.

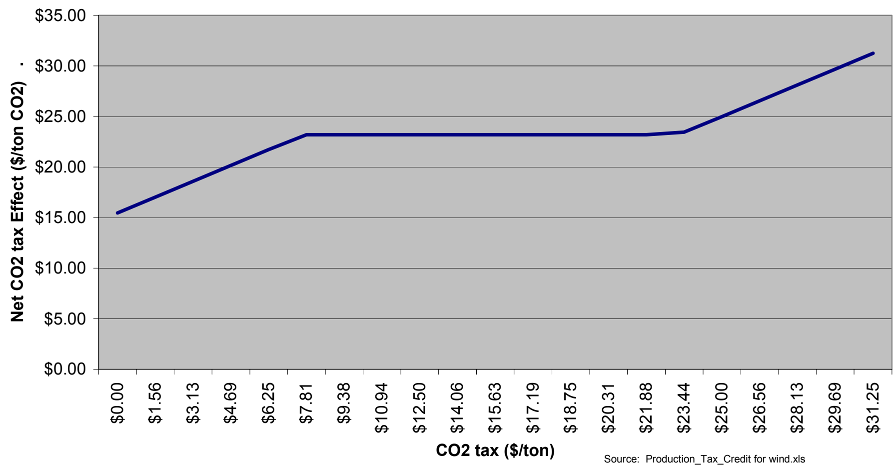

Figure P-76: Combined PTC and CO_2 Tax Effects

T81 = =IF(T80<0.5*T79,T79,IF(T80>1.5*T79,0,1.5*T79-T80))

	O	P	Q	R	S	T	U
74		Behavior: Port_24 Emissions_002, Subperiod: (all)		0	0	0	0
75							
76			PTC, JK workbook	104.3030438			
77			Integration	5.02	10.76		
78			Greentag Value	3.5	4.5		
79	Conversion (#CO2/kWh)		PTC (after commercial viability test)	9.90	9.90	9.90	9.90
80		1.28	Carbon Tax ($/MWh)	0.00	0.00	0.00	0.00
81			PTC (after CO2 tax effect)	9.9	9.90	9.90	9.90
82			VOM	1.13	1.13	1.13	1.13
83			Greentag Value	3.50	3.51	3.53	3.54

Figure P-77: Transition Logic (Step 3)

Figure P-78 characterizes the deciles for the PTC before adjustment for CO_2 tax. As

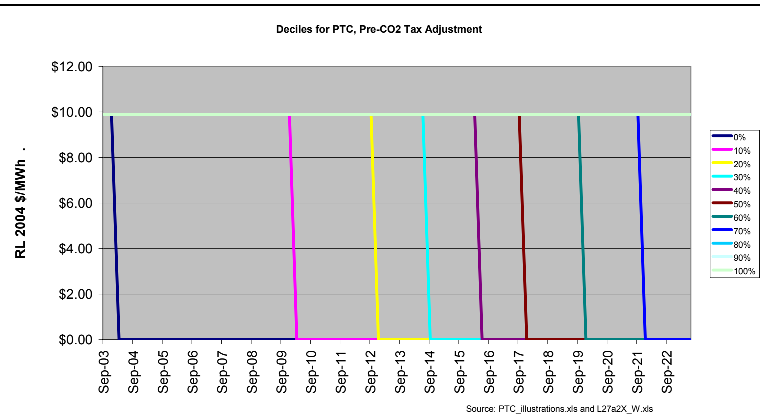

Figure P-78: Deciles for PTC before the CO_2 Tax adjustment

expected, the median value is around 2016, although the median is not the mode for the distribution in Figure P-73.

Figure P-79 has deciles for the final PTC, after the CO_2 tax adjustment. The effect of the tax is evident in each of the decile curves, with greater effect visible in out-lying years. The average of the final, quarterly values is a dotted line in this Figure. It also behaves as expected. Appendix L documents the final use for PTC value.

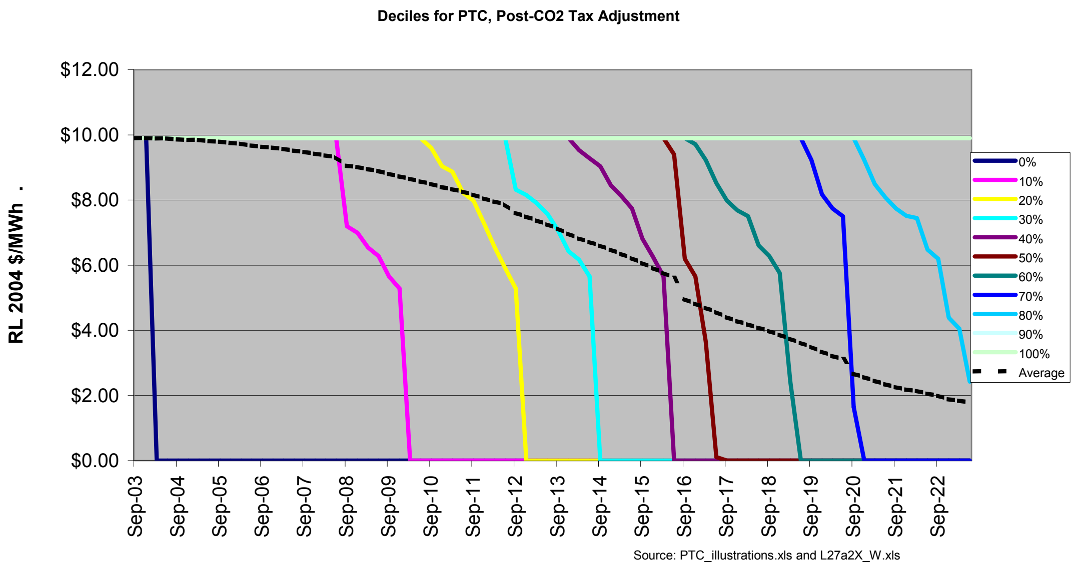

Figure P-79: Deciles for PTC After the CO_2 Tax Adjustment

Green Tag Value

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,
- paid for renewable energy components of utility supply portfolios, and
- paid for renewable acquisitions to meet requirements of renewable portfolio standards and system benefit charges.

Tag value varies by resource and was between $3 to $4 per megawatt-hour for wind power when the Council approved the final plan.

In the model, green tag value can start the study period anywhere 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 $8 per megawatt-hour (2004$). (See Reference [**31**].) 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.

In the workbook, the green tag value is a simple linear function of time. First, the model draws random variables for the starting value and the ending values. Figure P-80 illustrates the Crystal Ball assumption cells, {{R78}} and {{S78}}, respectively, responsible for providing those values.

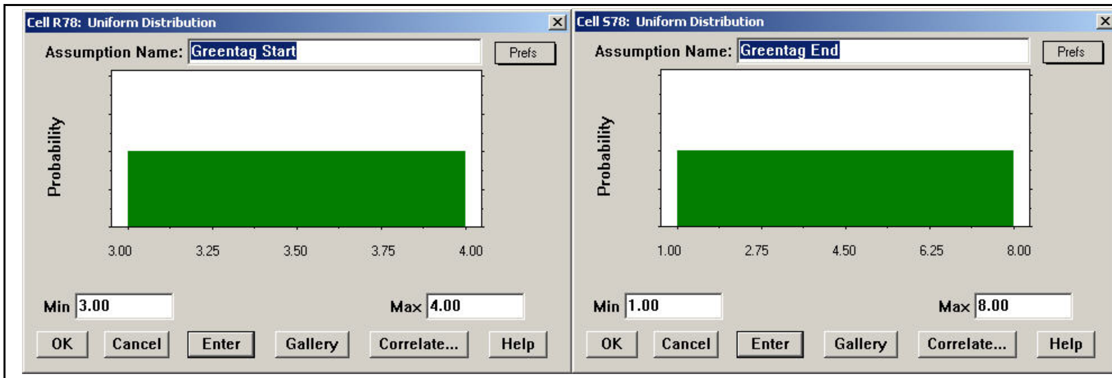

Figure P-80: Green Tag Starting and Ending Values

The model then creates a straight-line function over periods, as illustrated by the formula in Figure P-81.

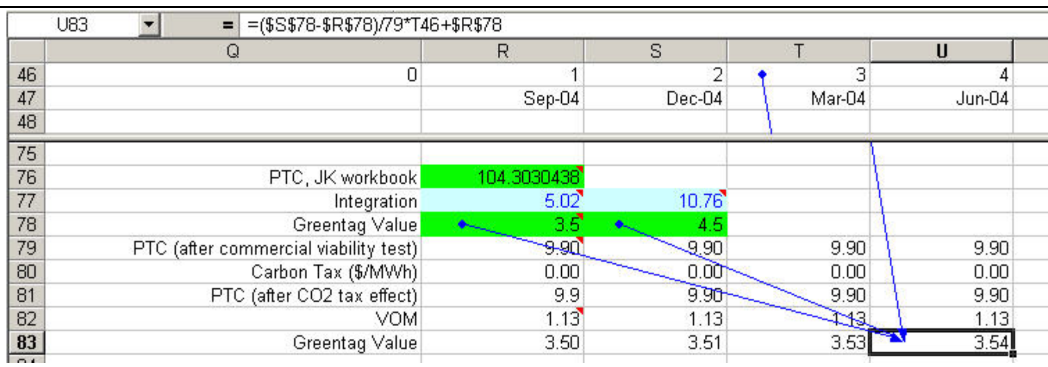
U83 = =($S$78-$R$78)/79*T46+$R$78

	Q	R	S	T	U
46	0	1	2	3	4
47		Sep-04	Dec-04	Mar-04	Jun-04
48					
75					
76	PTC, JK workbook	104.3030438			
77	Integration	5.02	10.76		
78	Greentag Value	3.5	4.5		
79	PTC (after commercial viability test)	9.90	9.90	9.90	9.90
80	Carbon Tax ($/MWh)	0.00	0.00	0.00	0.00
81	PTC (after CO2 tax effect)	9.9	9.90	9.90	9.90
82	VOM	1.13	1.13	1.13	1.13
83	Greentag Value	3.50	3.51	3.53	3.54

Figure P-81: Formula for Green Tag Value Over Periods

The decile summary for this stochastic variable is particularly uncomplicated and appears in Figure P-82:

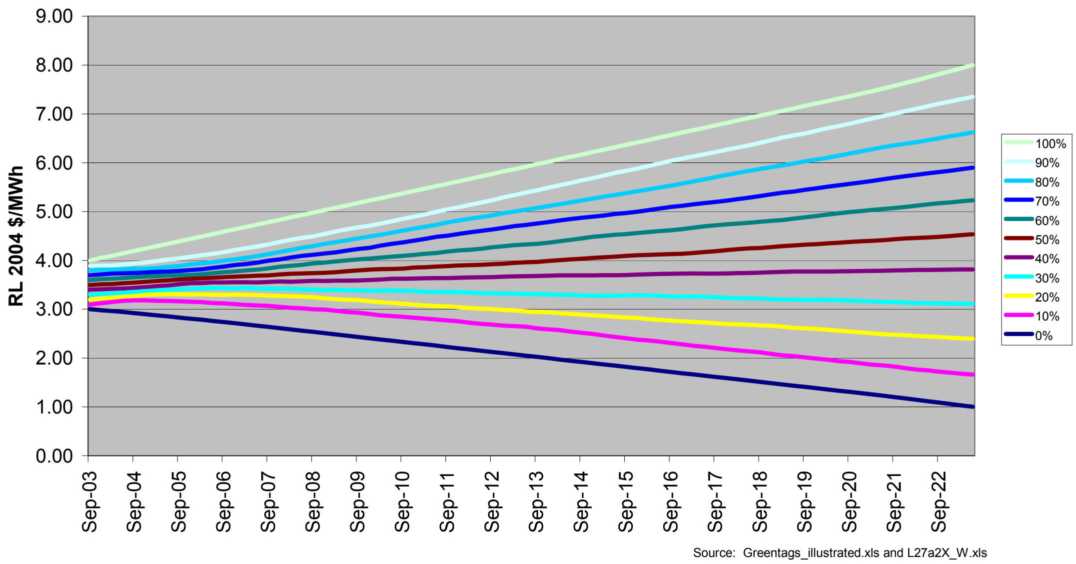

Figure P-82: Deciles and Average for Green Tag Value

Appendix L documents the final use for green tag value and how it is incorporated, along with PTC and variable operations and maintenance, into the cost of wind generation.

Correlations

Correlations among variables are typically different at different time scales. For example, load may have positive correlation with electricity prices on an hourly time scale, but on an annual average scale have negative correlation. This negative correlation stems from demand elasticity. Consequently, this section deals with correlation among key variables at different time scales.

The regional model explicitly addresses three time scales. The first is hourly correlation, within a quarterly period, referred to here as intra-period correlation. The second is correlation of quarterly averages. The third is correlation that exists on the scale of multiple periods. The first situation has its own section, while the second and third situations are combined. If it is essential to discriminate between the second and third types of correlation, the section distinguishes them in context.

There are also explicitly modeled correlations and those correlations that arise from assumptions, choices, and constraints in the model. The latter includes the relationship between electricity price and the amount of resource that is available due to the selection of a particular specific plan. (See the discussion of "RRP Algorithm" in Appendix L, page L-51.) It also extends to the relationship between electricity price and resource parameters, like the CO_2 tax. Because these relationships depend on variables that may or may not be representative for particular situations, however, this section does not attempt to characterize such correlations.

Short-term Correlations

The correlation of values assumed within each period appears in Figure P-83. More accurately, these are correlations of values within each subperiod. The distinction is important. Note, for example, that there is no correlation assumed between hydrogeneration energy and load or between hydrogeneration and market price. In fact, as much hydrogeneration as possible is produced on-peak, when market prices are high, which would result in high correlation. The solution to this apparent paradox is that the model already captures such correlation by distinct treatment of these variables in subperiods. The correlation table in Figure P-83, properly speaking, is any correlation net of subperiod modeling.

	Market Price	Non-DSI Load Flat Energy	PNW West - NG Price	Hydrogeneration Energy
Market Price	1.00	0.95	0.60	0.00
Non-DSI Load Flat Energy	0.95	1.00	0.00	0.00
PNW West - NG Price	0.60	0.00	1.00	0.00
Hydrogeneration Energy	0.00	0.00	0.00	1.00

Figure P-83: Correlation of Hourly Values

Because of how the regional model captures energy and cost, any temporal correlation of a variable with itself (autocorrelation) at the hourly scale is not relevant. The value of thermal dispatch over a subperiod, for example, is the sum of hourly values.

Correlation of natural gas price with electricity prices is significant to estimating the cost and value of thermal dispatch, as well as a forecasting capacity factor. An hourly correlation of 60 percent is taken as representative. Because of the many sources of interaction between load and electricity market price, this correlation is 0.95. All other correlations are zero. These values appear in the regional model at range {{R14:T16}}, shown in Figure P-84. Because hydrogeneration has no correlation with the other variables, its presence is not necessary. Because the correlation matrix is symmetric, this table includes only the values above the diagonal.

	Market Vol Price	Non-DSI Load Flat Vol	PNW West - NG variable cost vo
Market Vol Price	1	0.95	0.60
Non-DSI Load Flat Vol		1	0.00
PNW West - NG variable cost vol			1

Figure P-84: The Regional Model's Correlation Table

Long-term and Period Correlations

There are essentially three, explicit long-term correlations: the effect of natural gas price, loads, and hydrogeneration on electricity price; the effect of electricity price on loads; and the autocorrelation (chronological correlation) of variables with themselves. The regional model handles correlations of period averages for distinct variables through sensitivities, that is, a linear adjustment of one variable's average by another variable's average. It captures autocorrelations either through principal factors (Page P-26) or, in the case of aluminum price, through GBM coefficients (Page P-23).

Modeling correlation between averages of distinct variables as sensitivities is consistent with the correlation simulation described in the section "Simulating Values for Correlated Random Variables" on page P-24. Recall that for electricity, there remains a significant random term, the "independent" term, which provides uncorrelated behavior.

This appendix describes the effect of natural gas price, loads, and hydrogeneration on electricity price in section "The Influence of Loads, Natural Gas Price, and Hydro Generation," beginning on page P-69. It outlines the effect of electricity price on loads in the treatment of electricity price uncertainty, under the subsection "The Application to Load Elasticity," starting on page P-73.

Risk Measures

This chapter describes risk measures and the treatment of risk. It begins with a discussion of risk measures generally and considerations that led the Council to select the risk measure used in the regional model, $TailVaR_{90}$. It examines alternative risk measures and explains how each one relates to the $TailVaR_{90}$ risk measure.

This examination leads us to the following observations. 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. This correspondence is not accidental. It probably does not hold for individual utilities. The correspondence stems from the impact that adding a large amount of regional resources can have on regional prices. Individual utilities, on the other hand, are typically price takers whose supply actions do not affect market prices.

Background

It may be useful to define what the Council means by risk.

Risk is a measure of the expected severity of bad outcomes.

A specific example of a measure of risk is the average of outcomes in the "bad" tail of a distribution of costs, as illustrated in Figure P-85. In this case, bad outcomes are outcomes that are more expensive. This definition distinguishes the Council's risk measure from several in common use. For example, some use the standard deviation of the distribution of outcomes as a risk measure. The standard deviation, however, does not measure bad outcomes per se. The Council considers the standard deviation a measure of

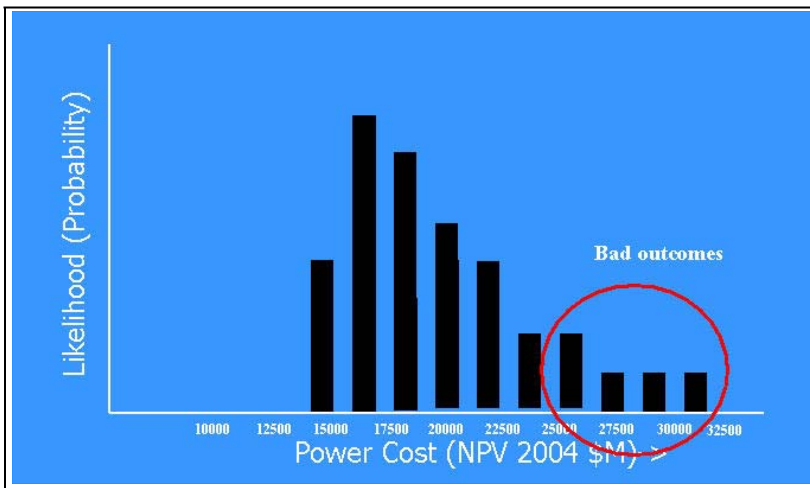

Figure P-85: Bad Outcomes

predictability, not risk.

There are several reasons for the selection of this definition of risk. First, the Council believes a measure should not penalize a plan because the plan produces less predictable, but strictly better outcome. Consider, for example, Plan A and Plan B, which have cost outcomes distributed as illustrated in Figure P-86. Plan B has more predictable outcome but every outcome is worse (more expensive) than any outcome for Plan A. The Council would not consider Plan A riskier than Plan B. Even if the distributions overlapped, but for each future (game) Plan A did better than Plan B, the Council would not consider Plan A riskier than Plan B.

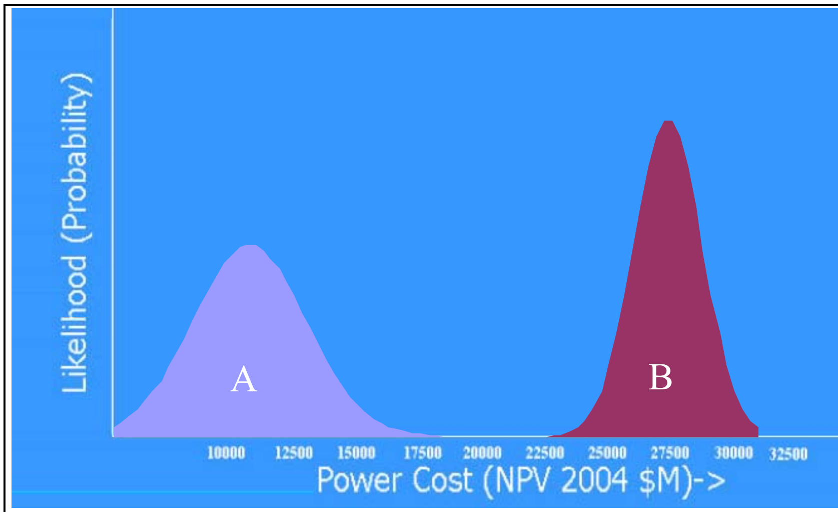

Figure P-86: Plans A and B

When confronted with situations like that which Figure P-86 illustrates, it is tempting to dismiss the problem because the average costs for Plan B are obviously worse than those for Plan A. No decision maker, it is argued, would fall into the trap of choosing the "less risky" Plan B over Plan A. That may be true in this situation, but consider the following example.

One plan produces the distribution of costs shown in Figure P-85; another plan creates the distribution in Figure P-87. (The section discussing cost distributions for the regional study describes the characterization of these distributions as those associated respectively with the resource surplus "price taker, and with deficit plan.") The distributions are mirror images of one another, reflected around the mean. Because they are mirror images of each other, they obviously have the same average cost and standard deviation. A decision maker using average cost and standard deviation would therefore not be able to discriminate between them. Comparing the distributions directly, however, reveals that the first distribution has a much greater likelihood of bad outcomes than the second. (See Figure P-88). The Council would consider the first plan riskier than the second.

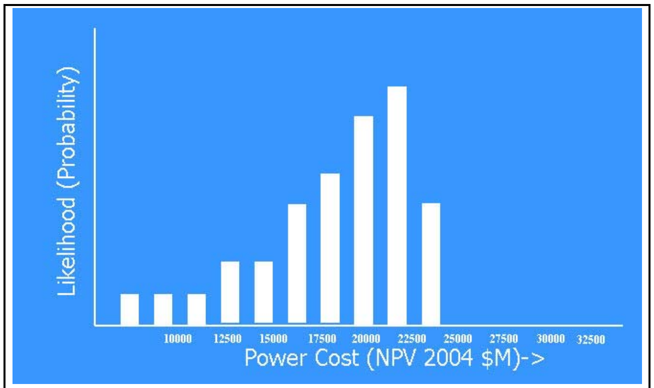

Figure P-87: Cost Distribution for Price Taker with Surplus Resources

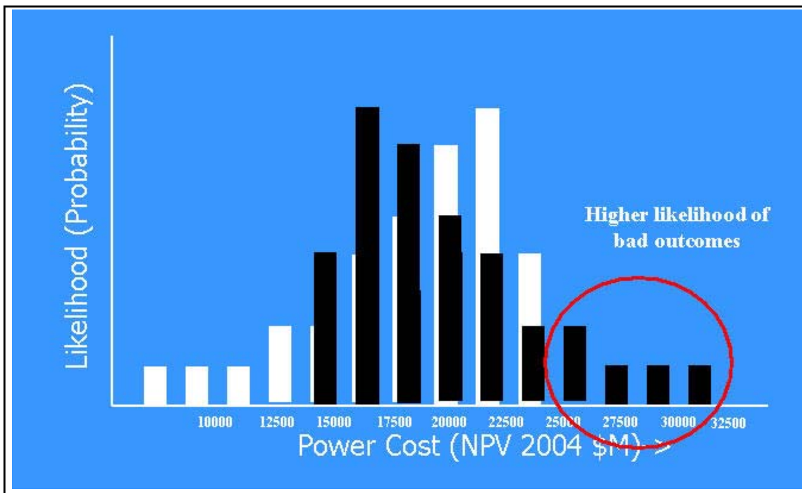

Figure P-88: Cost Distributions for a Price Taker, a Resource-Surplus Plan (white) and Resource-Deficit Plan (black)

Another reason to choose the definition of risk that the Council has is because it can be less expensive to reduce only expected severity of bad outcomes. Homeowner's fire insurance, for example, limits the economic damage that would otherwise take place in an accident. The insurance premiums, however, are typically much less expensive than the alternative of fire-proofing the home and its contents.

Finally, improving predictability by reducing the standard deviation may come at the cost of eliminating good outcomes as well as bad. In the example of fire insurance, for example, neither the fire insurance nor the alternative of fireproofing the home improve the outcome in fortunate circumstances. There is either a premium to pay or the cost of fire-proofing. The cost of fireproofing, however, impacts good outcomes much more.

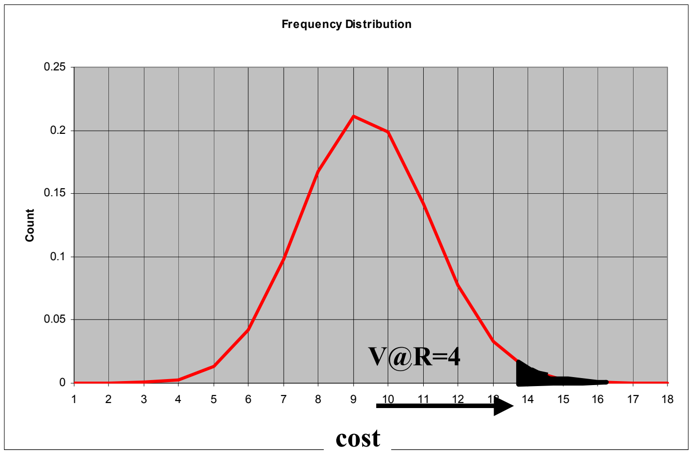

Figure P-89: Value at Risk

Some measures of risk recognize the logic of reducing bad outcomes but fall short in other regards. Value-at-Risk or VaR (sometimes V@R) is an example. Value-at-risk is a risk measure popular with investment and trading companies. VaR estimates the loss on a portfolio possible over a given period. Specifically, VaR_{95} is the loss exceeded with less than five percent likelihood. The loss is relative to some benchmark, such as the mean of the distribution. In Figure P-89, the probability distribution represents the possible costs[36] associated with a project over the next month, denominated in millions of dollars. The black tail of the probability distribution represents 5 percent of the area, and the 95th quantile is \$13.5M. If the expected cost is \$9.5M, the VaR is \$4M.

The problem with VaR is that it does not capture the value of portfolio diversification. To illustrate this, consider a simple situation where the good outcome has zero cost and the bad outcome has a cost of \$1. Consider two instruments (X_1 and X_2) with independent but identically distributed costs, sampled across ten futures (games) as shown in Figure P-90. Each instrument has a one-in-ten chance of producing a bad outcome. Each instrument has a VaR_{85} of zero, because more than 85 percent of the outcomes are zero (or less). The portfolio combining these two independent instruments, however, has a VaR_{85} of 1.0, which is a riskier VaR level. That is, the portfolio is riskier, as measured by VaR_{85}, than the individual instruments! This is contrary to the concept of diversification.

[36] Note that we could have used the example of losses on a portfolio of investments, operating expenses incurred by a company, or a host of other cases. The principle of measuring bad outcomes is the same.

Future	X_1	X_2	X_1+X_2
1	0.00	0.00	0.00
2	0.00	0.00	0.00
3	0.00	0.00	0.00
4	0.00	0.00	0.00
5	0.00	0.00	0.00
6	0.00	0.00	0.00
7	0.00	0.00	0.00
8	0.00	0.00	0.00
9	0.00	1.00	1.00
10	1.00	0.00	1.00
VaR@85%	0.00	0.00	1.00

$$0 = VaR(X_1) + VaR(X_2) < VaR(X_1 + X_2) = 1$$

Figure P-90: Outcomes for Two Instruments in a Portfolio

The problem with VaR is it gives no indication of *how bad* the outcomes are within the bad tail. In fact, any risk measure that reports only statistical quantiles suffers from this shortcoming.

Coherent Measures of Risk

Experts in investment and risk management recognized the problems just described, and in the 1990s produced a class of risk measures that addresses them.[37] A ***coherent*** measure $\boldsymbol{\rho}$ of risk is a function from outcome distributions to the real numbers. It has the four following mathematical properties. These properties make the measure useful for properly ranking choices. They also address the issues raised earlier. The property of Monotonicity, for example, guarantees that if all of the outcomes for a given plan are better, then that plan will not have greater risk. The property of Subadditivity guarantees that portfolio diversity reduces risk. In the following, $\boldsymbol{\lambda}$ and $\boldsymbol{\alpha}$ are real number-valued constants.

[37] In 1999, Philippe Artzner, Universite Louis Pasteur, Strasbourg; Freddy Delbaen, Eidgen*f*ossische Technische Hochschule, Zurich; Jean-Marc Eber, Societe Generale, Paris; and David Heath, Carnegie Mellon University, Pittsburgh, Pennsylvania, published “Coherent Measures of Risk” (***Math. Finance*** 9 (1999), no. 3, 203-228) or http://www.math.ethz.ch/~delbaen/ftp/preprints/CoherentMF.pdf

- Subadditivity – For all random outcomes (losses) X and Y,

$$\rho(X+Y) \leq \rho(X)+\rho(Y)$$

- Monotonicity – If $X \leq Y$ for each future, then

$$\rho(X) \leq \rho(Y)$$

- Positive Homogeneity – For all $\lambda \geq 0$ and random outcome X

$$\rho(\lambda X) = \lambda\rho(X)$$

- Translation Invariance – For all random outcomes X and constants α

$$\rho(X+\alpha) = \rho(X) + \alpha$$

The Council's measure of risk, $TailVaR_{90}$, is coherent [**32**]. It is defined to be the average of the ten percent worst outcomes, as illustrated in Figure P-91.

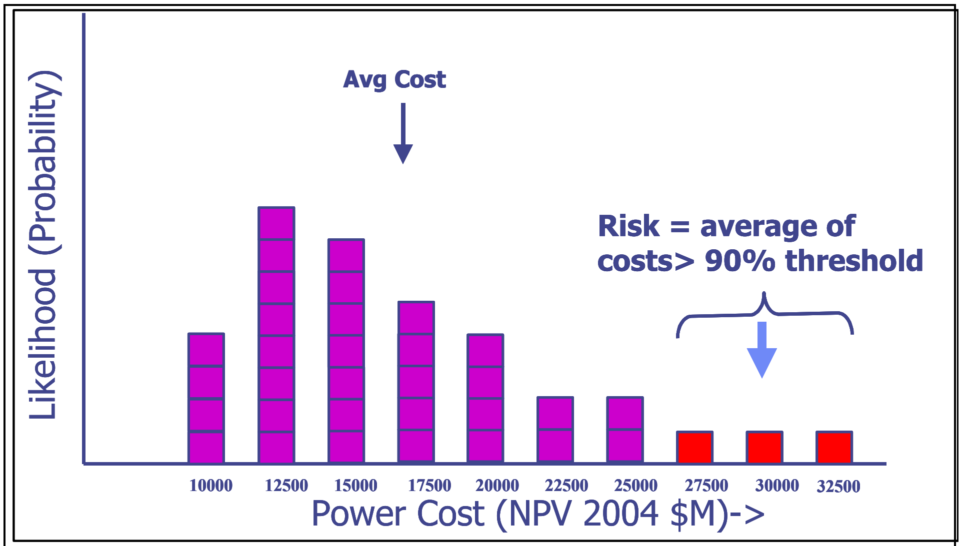

Figure P-91: Plan Cost Distribution, Average Cost and Risk (TailVaR90)

$TailVaR_{90}$ is a measure of risk associated with *economic efficiency*. The Northwest Power and Conservation Council is required to develop a 20-year power plan under the Pacific Northwest Electric Power Planning and Conservation Act to assure the region of an adequate, efficient, and reliable power system. Previous and current Council studies use net present value (NPV) as a measure of economic efficiency. NPV is demonstrably better for this purpose than alternatives, such as B/C ratios and internal rate of return (IRR). Because the primary measure is one that relies on NPV, it stands to reason that bad outcomes are those with unfavorable NPV. Consequently, $TailVaR_{90}$ is fashioned to measure the expected severity of unfavorable NPV.

$TailVaR_{90}$ distinguishes between the two distributions illustrated in Figure P-88. It is reasonable to expect, therefore, that the results obtained using this measure would not compare well with those obtained using a non-coherent measure of risk, like standard deviation. Surprisingly, non-coherent and coherent measures gave comparable results in regional studies. The next section explains why this is so.

Distributions of Cost for Regional 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. This potential depends only on how high the market price for electricity goes. As the following explains, however, this is not the case for the region as a whole.

An example of the cost distribution situation for price takers with surplus resources appears in Figure P-87. The source of risk for utilities with surplus resources is low market prices for electricity.

The following shows that low market prices are associated with the right-hand tail in the figure. With low market prices, the utility and its customers are better off if the utility buys its electricity from the market. This leaves the utility with the cost of "stranded resources," that is, plants that customers are still paying for but are not using. The size of this risk may be large, but it is limited. Market price for electricity may go down significantly, but it obviously cannot go below zero. This means costs beyond meeting requirements out of market purchases will not be greater than the fixed costs of unused resources. Therefore, total costs have an upper bound, as illustrated in the simple example shown in Figure P-92.

Figure P-92 shows the total costs of a simple system over a period, say a year, if electricity prices were to remain fixed at the value on the horizontal axis. This system has a load, and there is a cost of meeting that load in the electricity market. The dark purple, dotted line illustrates that cost. The system has a single generator that costs $50M/year in fixed costs and a dispatch price[38] of $30/MWh. Significantly, the size of the generator is twice the size of the load. The generator costs are the solid, dark blue line. When electricity price exceeds the generator's dispatch price, the generator creates value that offsets its fixed costs. The value of the generator in the electricity market increases dollar for dollar, with each dollar that the electricity price exceeds the dispatch price. The total costs, shown by the solid yellow line, are maximum at the dispatch price of the turbine. For prices higher than that, the turbine value offsets the cost of serving the load; for lower prices, lower purchase costs reduce total cost.

[38] The dispatch price is the electricity price that would cause the generator to cover just the cost of fuel and any other cost of operation that depends only on the amount of energy generated.

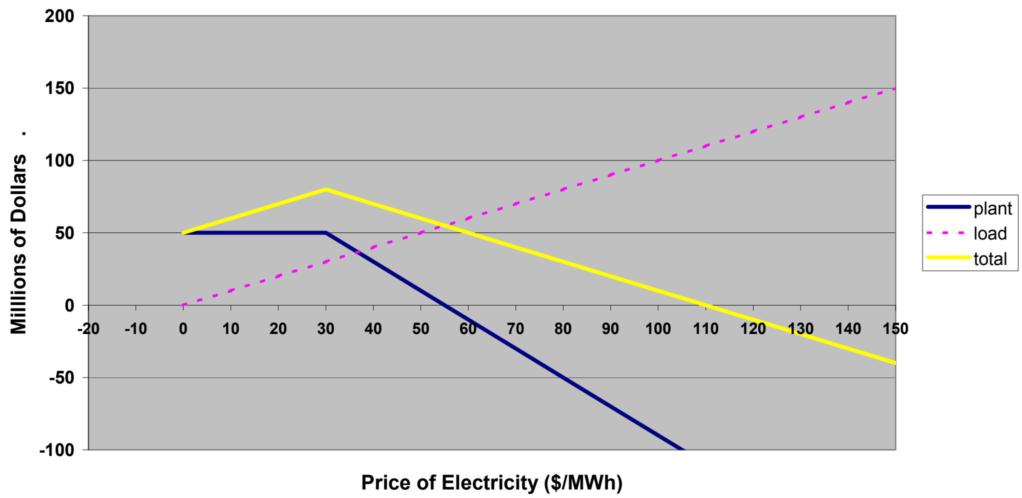

Figure P-92: Example of Total Costs for a Simple Surplus System

If electricity prices were fixed, the total cost could be read off the vertical axis of Figure P-92. For electricity prices that have a distribution instead of being fixed, however, there is a corresponding total cost distribution. The cost distribution has a tail extending to the left (*lower* costs) in Figure P-87, corresponding to *higher* electricity prices, because of the relationship shown in Figure P-92. Net costs can even become negative if prices are high enough, as Figure P-92 suggests.

The cost distribution situation for price takers with *deficit* resources is similar, except costs are now bounded below and *unbounded above*. For the simple example illustrated in Figure P-93, the load is larger than the plant. Now, however, higher costs correspond to *higher* electricity prices. If electricity prices have lognormal distribution, the distribution will have an unbounded tail extending to higher prices. This situation leads to a total cost distribution resembling that in Figure P-94. That cost distribution now has a tail pointing in the direction opposite that of Figure P-87.

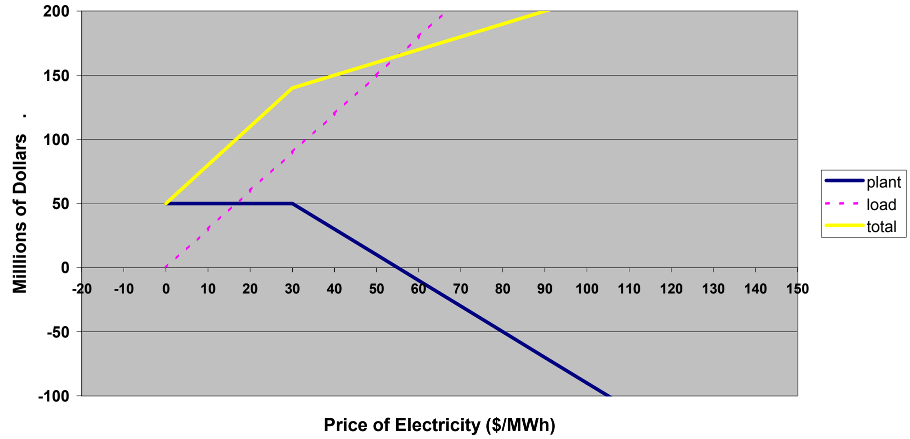

Figure P-93: Example of Total Costs for a Simple Deficit System

The region's cost distribution, it turns out, never resembles that for the surplus system. The preceding examples assume that utilities are price-takers, that is, the utility's surplus does not dampen electricity market prices. The aggregate regional resource situation, however, can affect market prices. Resources surplus to the regions requirements, after

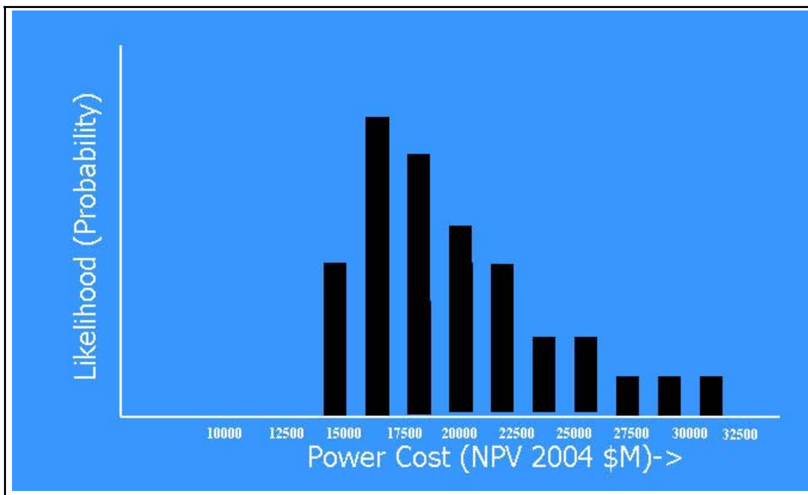

Figure P-94: Cost Distribution for Price Taker with Deficit Resources

exports, depress price. Effectively, the price range in Figure P-92 is capped on the high side, trapping the costs in positive territory. The final distribution for costs will tend to be more symmetric in this case than it would be for a deficit region. The width of the distribution may become quite small, but the mean will go up due to fixed costs. The "good," right-hand tail that is present in Figure P-87, however, does not materialize.

Because distributions like that in Figure P-87 never arise in the regional study, the mean cost is higher than the median cost. This has relevance to the question of the metric chosen for central tendency. Some would argue that the median is a better measure of central tendency than the mean for risk analysis. The next section is a brief digression from the topic of risk measures to address that issue.

Median and Mean Costs

Is the median a better measure of central tendency than the mean for risk analysis? The median future is a future where an equal number of better and worse futures lie above and below. The ordering here is by the cost of the future. Thus, only relative order of the futures determines the medium. In contrast, the mean is more sensitive to the values of futures. It tells us nothing about the likelihood of futures that lie above and below. Now, the region will face only one future. Which future? For that matter, on which face will a rolling die land? It is a matter of the likelihood of landing on each face, not the value of the faces. The mean cost of future, by the same token, 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[39]. This all tends to argue for the use of the median.

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. The Council chose the mean to a certain extent because it is simpler to communicate than the median.

Fortunately, it does not make much difference which of the two measures of central tendency we choose. Distributions for outcomes of plans exhibit a strong relationship between the two measures. Figure P-95 shows that the mean and median values track very closely.

The mean value is consistently above the median, reflecting the observation that 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. 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 P-95 at the lowest average cost.

In conclusion, while the median might be a better measure of the central tendency than the mean for decision making under uncertainty, using the mean will give the same results in terms of the construction of the feasibility space and selection of plans. For studies of regional costs, distributions are skewed in the same direction as resource-deficit plans, and the mean and median have a strong correspondence.

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

This section suggests that, because distributions for regional cost are always skewed in one direction, non-coherent measures like standard deviation might give comparable results to those obtained by $TailVaR_{90}$. Returning to the topic of risk measures, the next section asks, how representative is $TailVaR_{90}$?

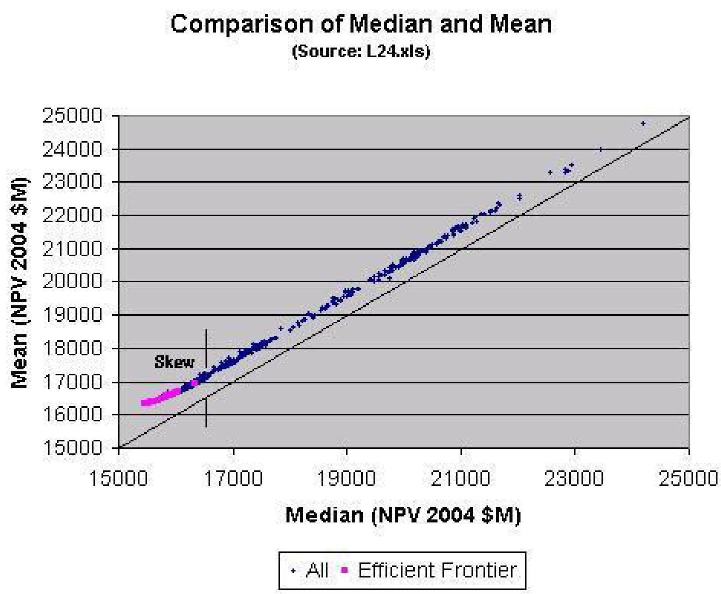

Figure P-95: Mean versus Median

Perspectives on Risk

Many alternatives exist for measuring the risk. Each study performed with the regional model recorded a host of alternative risk measures, as well as both the mean and median cost. Figure P-96 illustrates the standard report, which Appendix L describes in detail. Risk measures for each plan appear on the right-hand side of this report and include:

- $TailVaR_{90}$
- Standard deviation
- $CVaR_{20000}$
- VaR_{90}
- 90^{th} Decile
- Mean (over futures) of maximum (over 20 years) of annual cost increases
- Mean (over futures) of standard deviation (over 20 years) of annual costs

(The figure simplifies the report, leaving out some columns and rows, to provide a more comprehensive view of the report.) Subsequent, out-board studies examined alternative sources of risk, such as relative exposure to bad market conditions and variation in average power cost.

This section reviews this information, extracted from the final plan. This section asks:

- How representative of alternatives is $TailVaR_{90}$? Would the Council have made a different choice of plans if it had used some other measure of economic risk?
- Given that the Council chooses a plan from among those on the efficient frontier, do other measures help the selection?
    - o How do conventional measures of reliability, like loss-of-load probability (LOLP), vary along the frontier?
    - o Do other perspectives on risk, such as cost volatility, give us a way to further refine the selection?

	A	B	C	D	E	AQ	AR	AS	AT	AU	AV	AW	BG	BH	BI
1	**														
2	* Analysis of *														
3	* OptQuest.log *														
4	* with *														
5	* Analysis of Optimization Run_L27A2.xls *														
6	**														
7	Sim	Cnsrvn_Lo	Cnsrvn_Di	RM	CCCT_C	GCC_CY1	IGCC_CY1	Mean	Std_Dev	Median	TailVaR90	CVaR20	Cnsv_Cst:	Mean	
8	1706	0	5	0		0	0	23647.44	6602.989	22295.37	37435.84	26851.5	22.61565	F	
9	1873	0	5	5000		0	0	23653.11	6379.034	22272.49	36955.9	26707.6	22.48798	F	
[illegible]	1697	0	5	5000		0	0	23730.78	6250.586	22396.39	36729.42	26541.8	22.38016	F	
122	1233	10	5	5000		425	425	24430.97	5596.721	23208.61	35880.81	26168.4	22.92024	F	
123	1234	10	5	5000		425	425	24435.55	5593.984	23207.23	35880.22	26171.7	22.91977	F	
124	1232	10	5	5000		425	425	24440.25	5591.099	23205.85	35879.08	26175.3	22.91493	F	
125	948	5	25	5000		425	425	24508.42	5569.806	23311.89	35870.65	26202.5	24.06349	F	
126	3	0	0	0		0	0	23661.51	6599.311	22301.92	37450.84	26836.8	22.24873		x
127	1726	50	5	0		0	0	23847.44	6496.037	22455.33	37438.75	26864.0	[illegible]		x
1499	775	5	5	5000		425	425	24420.24	5604.747	23205.11	35887.48	26191.7	22.41363		x
1500	912	5	5	5000		425	425	24686.57	5485.664	23486.43	35882.99	26243.5	22.25928		x
1501	922	5	25	5000		425	425	24459.83	5598.099	23261.39	35881.38	26237.1	24.06979		x
1502	1230	10	5	5000		425	425	24467.54	5575.609	23260.49	35880.23	26198.7	22.90646		x
1503	60	50	50	0	12	1700	1700	31340.27	6122.014	29994.17	44043.84	31386.0	29.59224		
1504	24	50	50	0	12	1700	1700	30878.93	[illegible]	29537.35	43619.37	30924	29.59777		
2009	264	35	30	5000		425	425	25013.59	5386.828	23829.93	35987.55	26329.6	26.69841		
2010	807	15	5	5000		425	425	25090.18	5395.79	24027.38	35985.2	26423.2	23.24394		
2011	630	0	0	5000		425	425	24824.09	5489.656	23716.65	35963.71	26344.6	21.46526		
2012															

Figure P-96: Alternative Risk Measures (Right Hand Side) from Appendix L

The section first examines economic risk measures. These derive from distributions of net present value study costs and include coherent and non-coherent measures of risk. It then reviews measures of cost volatility. Cost volatility here refers to year-to-year variation in both going-forward costs and total costs, including embedded costs. It also refers to the consistency of factors that would affect rates, such as imports of expensive energy. Finally, the section addresses two conventional measures of engineering reliability, LOLP and resource-load balance.

Alternatives to TailVaR_{90}

As explained earlier in this chapter, measures of NPV distribution are the most appropriate risk measures, given the task of the Council's plan. Measures of NPV distribution are a kind of "economic efficiency" risk. Among alternatives to TailVaR_{90} for measuring such risk are 90th quantile, standard deviation, VaR_{90}. These examples happen to be non-coherent measures of risk. Even among coherent measures of economic risk, however, there are unlimited choices for such measures.

CVaR_{20000}, for example, is a coherent measure of economic risk, and earlier Council studies used it as the primary risk measure. CVaR_{20000} is the average of costs exceeding $20,000 million. The concept is that if decision makers can deem an economic threshold as undesirable, the average of costs above that threshold makes a reasonable measure of risk.

CVaR_{20000}, however, has several shortcomings. Most important, the Council does not have an a priori vision of what that threshold should be. The CVaR_{20000} measure even complicates the process of studying cost distributions to arrive at such a threshold. A distribution may shift dramatically with the introduction of new assumptions. If plan distributions for the base case and change case fall on one side or the other of the threshold, CVaR_{20000} cannot discriminate between them any better than the mean. Finally, because the threshold is a subjective assessment by the decision maker, selecting a threshold introduces another assumption to defend and debate.

TailVaR_{90} addresses these issues and affords additional benefits. Because the value of TailVaR_{90} is never less than the 90th quantile, for example, the Council can make statements about the likelihood of "bad" outcomes. That is, futures with TailVaR_{90} costs

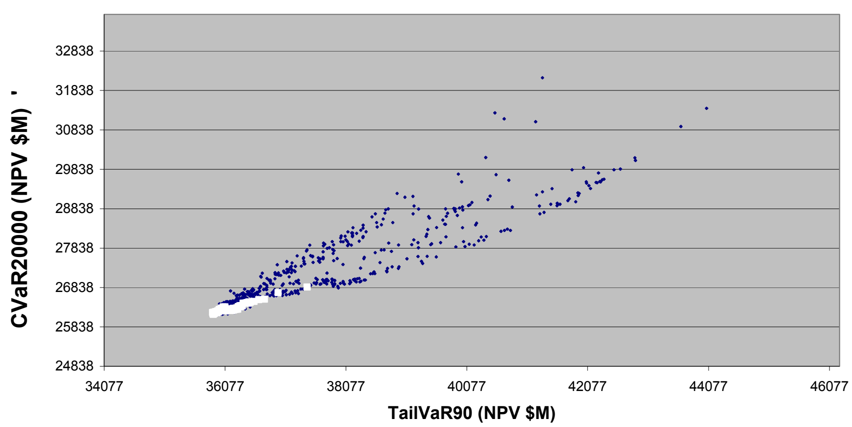

Figure P-97: Relationship, CVaR20000 to TailVaR90

or greater are expected with less than 10 percent probability.

One measure of how well $CVaR_{20000}$ compares with $TailVaR_{90}$ is their correlation. If they produce the same rank of outcomes, they provide effectively the same information. If the two measures are plotted against one another, as in Figure P-97 [**33**], the points would fall on a strictly monotonic curve. (See, for example, Figure P-95.) The dispersion of points around the monotonic curve is an indication of their correlation. Figure P-97 suggests that the correspondence between $CVaR_{20000}$ and $TailVaR_{90}$ is rather weak, in general. The white points in the bottom left-hand corner of the distribution, however, correspond to the efficient frontier, using $TailVaR_{90}$. On that efficient frontier, the correspondence is quite good.

Figure P-98 reconstructs the feasibility space using $CVaR_{20000}$. Again, the white points are the efficient frontier constructed by using $TailVaR_{90}$. Evidently, it does not make any difference whether we construct the efficient frontier using $TailVaR_{90}$ or $CVaR_{20000}$.

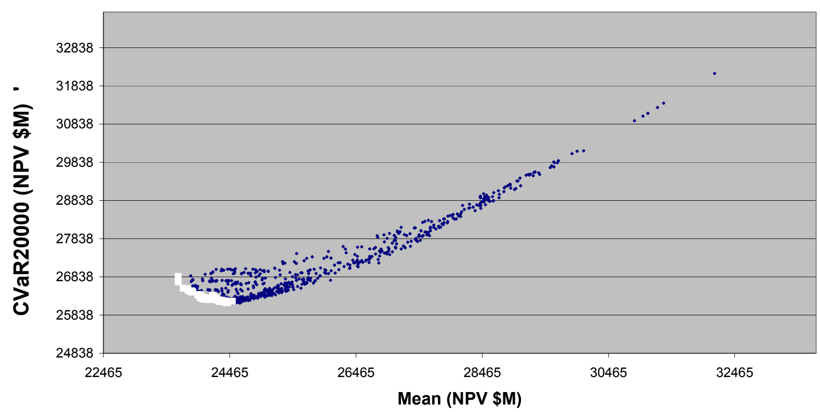

Figure P-98: Feasibility Space, using CVaR20000

90th Quantile

Non-coherent measures do not correspond well, in general, to $TailVaR_{90}$. Figure P-99 plots the 90th quantile against $TailVaR_{90}$.

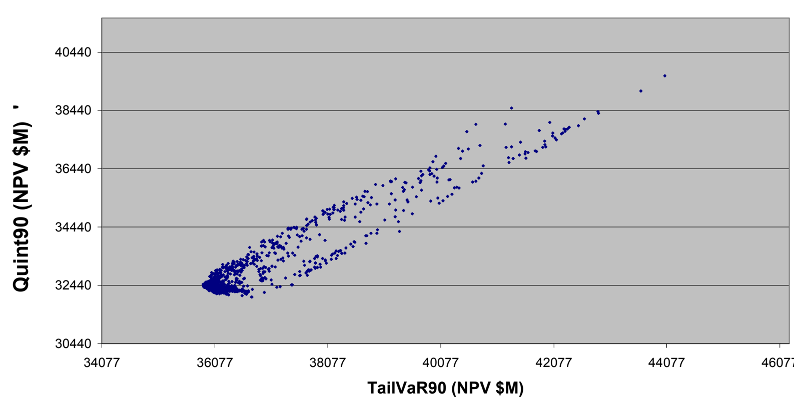

Figure P-99: Relationship, 90th Quantile to TailVaR90

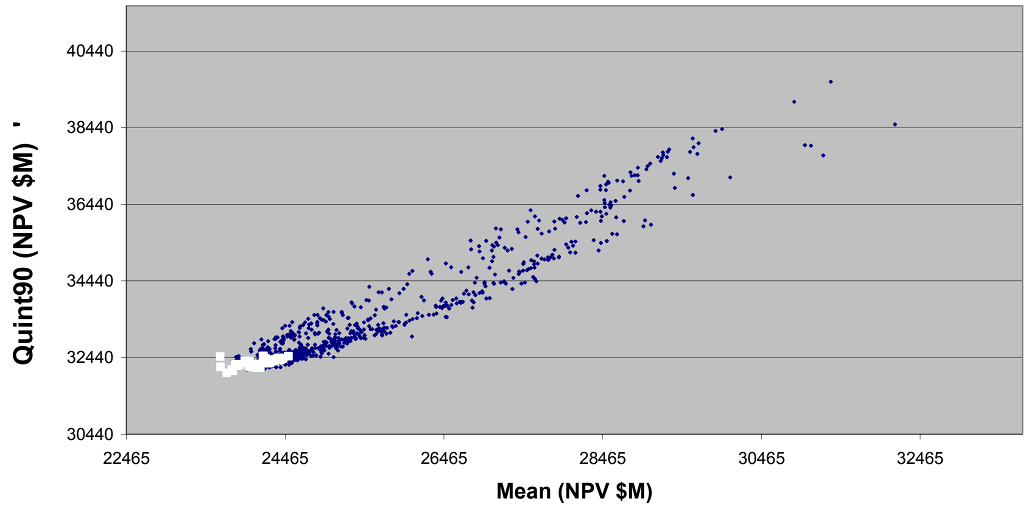

Figure P-100: Feasibility Space, using 90th Quantile

The relationship is clearly much weaker than for CVaR_{20000}. Figure P-100 makes it clear that the efficient frontier using the 90th quantile does not correspond to that using TailVaR_{90}. The efficient frontier using TailVaR_{90} is clearly well within the set of dominated points. It is reassuring, however, that the efficient frontier using TailVaR_{90} is contained in the set of nearly efficient points using the 90th quantile. It appears that plans that are efficient with respect to TailVaR_{90} are efficient, or nearly efficient, with respect to the 90th quantile.

Standard Deviation

Standard deviation bears virtually no relationship to TailVaR_{90}, as illustrated in Figure P-101. Fortunately, because the cost distribution for the region is always skewed in the

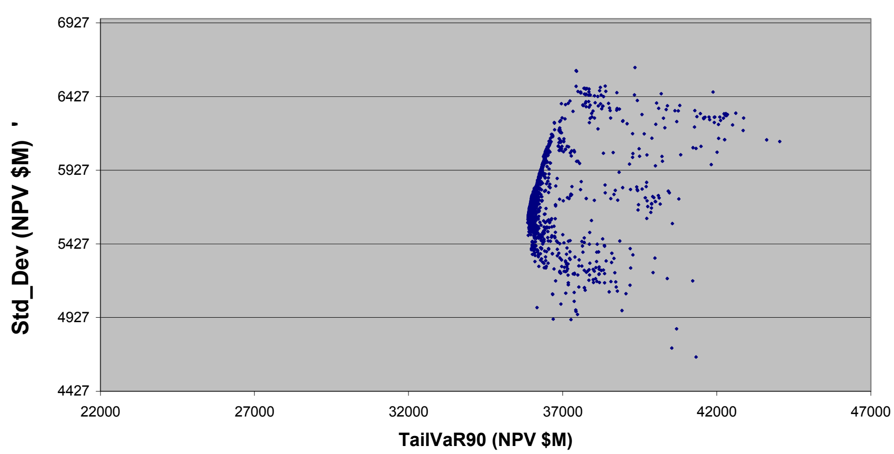

Figure P-101: Relationship, Standard Deviation to TailVaR90

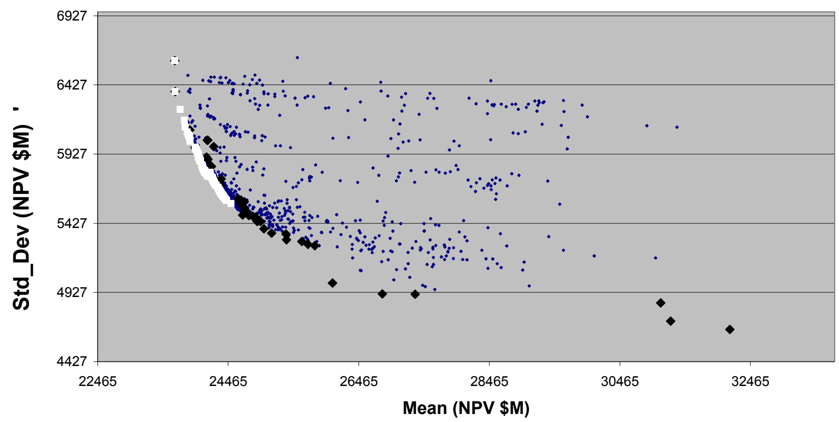

Figure P-102: Feasibility Space, using Standard Deviation

same direction, plans that are efficient using $TailVaR_{90}$ have least standard deviation for each level of cost. Consequently, those plans that are efficient using $TailVaR_{90}$ are also efficient using standard deviation, as Figure P-102 illustrates. In fact, the sequence of plans along the efficient frontier closely follows that for the efficient frontier using $TailVaR_{90}$. For example, the least risk plan, Plan D, is also least risk – among the white points – using standard deviation. (See Figure P-103.)

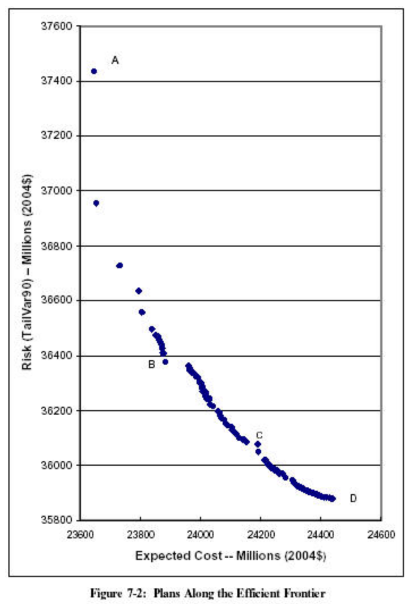

Figure P-103: Chapter Seven's Figure 7-2

There are a good number of plans, identified by the black diamonds in Figure P-102, that are efficient with respect to standard deviation, but not efficient with respect to $TailVaR_{90}$. This raises the obvious question, "Would the Council have selected another plan if they used standard deviation?"

It is unlikely that the Council would have chosen any of the Black Diamond plans. The reason, simply stated, is that these plans perform worse under a preponderance of futures than the plans corresponding to the white points.

For example, Plan E in Figure P-102 has substantially better standard deviation than Plan D ($380 million smaller). If we compare the total system cost of Plan E in each future against the cost of the corresponding future for Plan D, we can construct the illustration of the sorted differences appearing in Figure P-104. While Plan E is more predictable as measured by standard deviation, it produces a better outcome in less than two percent of the futures. The number of futures with significant difference is half of that. In over 80 percent of the futures, the outcome for plan E is over $1 billion worse than that for Plan D. The ability of $TailVaR_{90}$ to discern plans that perform better

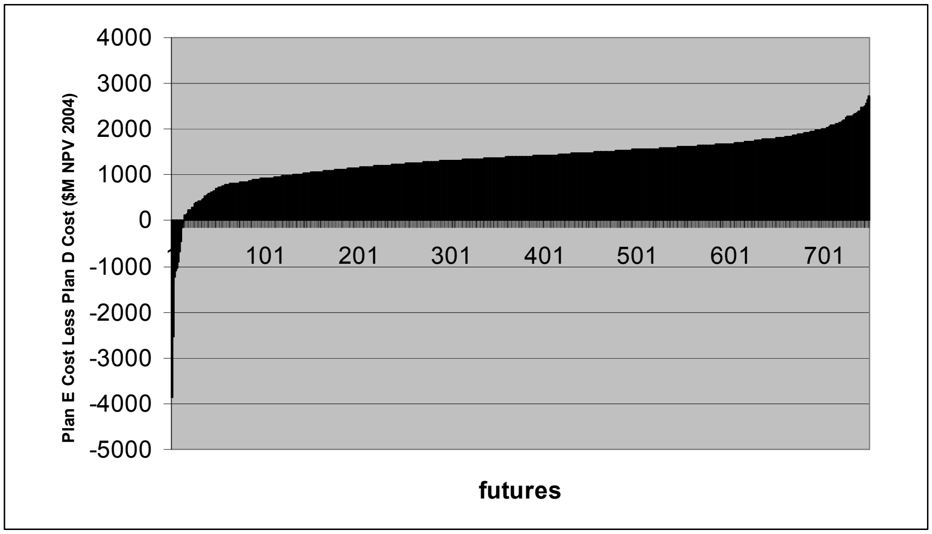

Figure P-104: Cost Differences Between Plans D and E, by Future

in the vast majority of futures is directly related to the property of monotonicity shared by all coherent risk metrics.

<u>VaR_{90}</u>

The definition of Value-at-risk (VaR) appears earlier in this chapter. It is a risk metric that, like standard deviation, primarily measures the width of the distribution. It should not be too surprising, therefore, that its correspondence to $TailVaR_{90}$ resembles that of standard deviation. (See Figure P-105.) The correspondence of the efficient frontier to that defined using TailVaR90 (white points) is not as clean as it is for standard deviation. Nevertheless, plans that are efficient with respect to $TailVaR_{90}$ are efficient or nearly efficient with respect to VaR_{90}. The efficient frontier in Figure P-106 below the white dots has the same explanation as the corresponding area for standard deviation. Once again, the conclusion is that it is unlikely the Council would have chosen plans from the efficient frontier of Figure P-106 below the plans illustrated with white points.

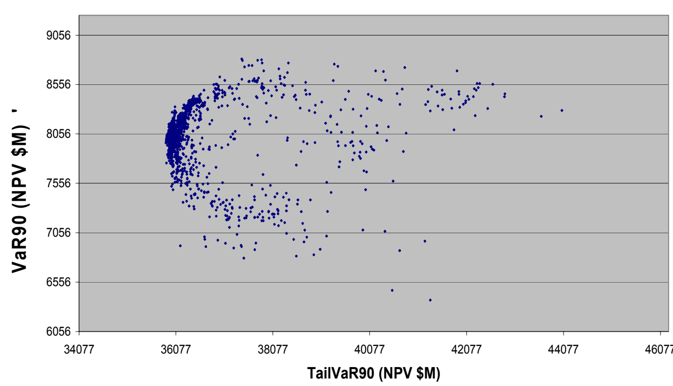

Figure P-105: Relationship, VaR90 to TailVaR90

Cost Volatility

Economic efficiency can hide a multitude of sins. Costs over the study period can produce low net present value while still exhibiting large volatility. Cost volatility is undesirable because it can produce sudden and unexpected retail rate increases.

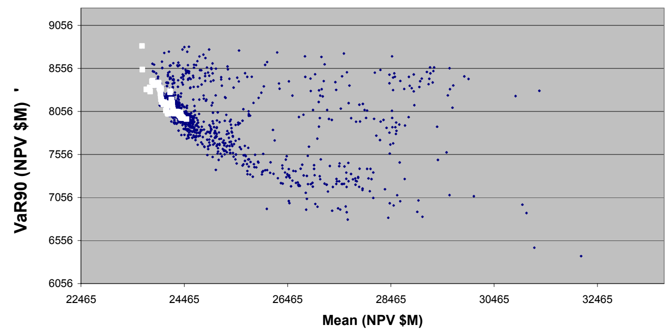

Figure P-106: Feasibility Space, using VaR90

There are several questions one can ask about cost volatility. First, how do the plans along the efficient frontier perform with respect to cost volatility relative to those plans that are not on the efficient frontier? Second, what kind of variation in cost volatility exists among plans on the efficient frontier? Third, what are some of the key drivers of cost volatility?

There are many ways to define cost variability. The next section considers several types of cost volatility and explains the purpose of each.

Average Incremental Annual Cost Variation

Figure P-107 illustrates that the relationship between the mean cost variation and TailVaR_{90} is quite weak. Mean cost variation is the average, across the 750 futures, of the standard deviations for changes in annual costs across the study:

$$\frac{1}{750}\sum_{i=1}^{750}\sigma_i$$

where

$$\sigma_i \equiv \sigma\left(\frac{C^i{}_t}{C^i{}_{t-1}} - 1 \ , t = 2, \cdots, 20\right)$$

and

$C^i{}_t$ is the cost in year t of future i

This tends to be a weak indicator of volatility for a couple of reasons. This standard deviation uses the first half of the study, when there are virtually no differences among plans. Averaging over futures tends to water down this metric as well.

There are a few things that we can discern, however, from Figure P-107. Plans on the TailVaR_{90} efficient frontier[40] (white points) all tend to lie in a narrow range of mean cost variation. That is, by this measure it does not really matter which plan from the efficient frontier we choose. It is notable that there are many plans with less mean cost variation. These are associated with more expensive plans and surplus resources. Resource shortage and electricity market price volatility increase cost variability; surplus resources will lower cost volatility because they tend to dampen wholesale electric market prices. Also, while the plans on the TailVaR_{90} efficient frontier are not among those with the lowest cost variation, they are also not among those with the greatest.

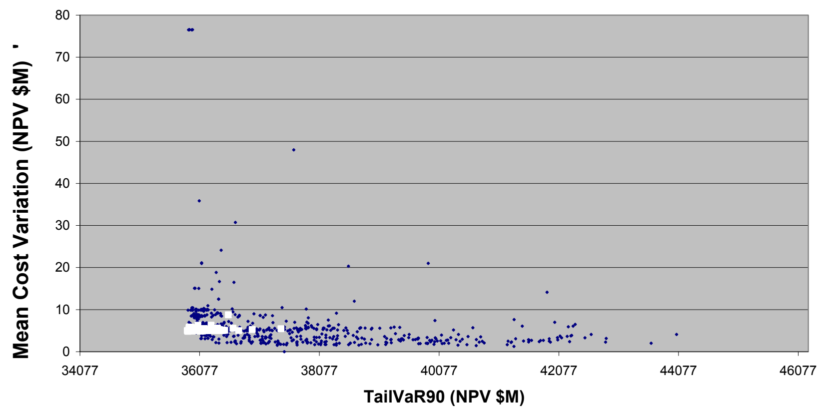

Figure P-107: Relationship, Average Incremental Annual Cost Variation to TailVaR90

<u>Maximum Incremental Annual Cost Increase</u>

A slightly more sensitive measure of cost volatility is the maximum increase in annual costs, averaged across futures:

$$\frac{1}{750}\sum_{i=1}^{750}\max_t\left(\frac{C^i{}_t}{C^i{}_{t-1}}-1\ ,t=2,\cdots,20\right)$$

Figure P-108 compares these values to TailVaR_{90}. We still see roughly the same pattern that was evident for average incremental annual cost variation. If we expand the area around the TailVaR_{90} efficient frontier, we can see that there is a very weak relationship between the two measures. Figure P-109 includes a regression line that emphasizes this weak relationship.

[40] In this chapter, the TailVaR_{90} efficient frontier refers to those plans that are on the efficient frontier if they were in a plot of plan mean cost against plan TailVaR_{90}.

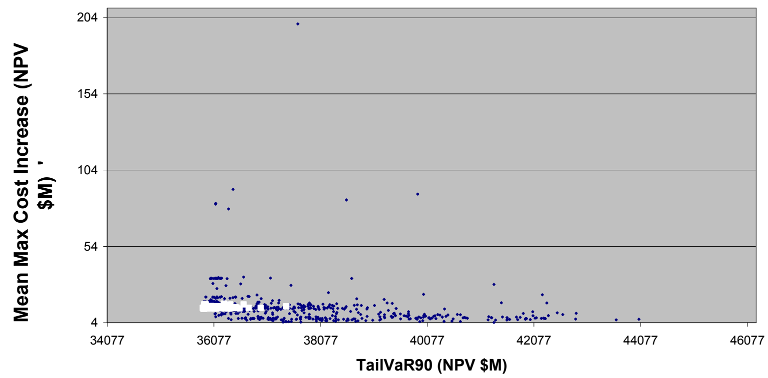

Figure P-108: Relationship, Maximum Incremental Annual Cost Increase to TailVaR90

Council staff investigated several alternative measures for cost volatility. The next section describes several of the more successful results.

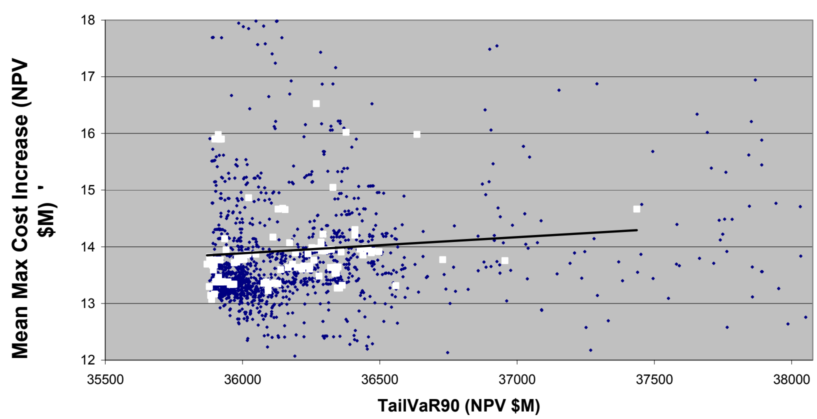

Figure P-109: Magnification of Previous Figure

Average Power Cost Variation (Rate Impact)

Counting the number of futures with cost increases that exceed in a given level yields a more stable measure of cost volatility. This section describes how the four scenarios identified in Figure P-103 perform under this measure.

Figure P-110 [**34**] shows the percent of futures where cost increases exceed the levels on the horizontal axis. While the preceding discussions of annual cost volatility used only variable costs and forward-going fix costs, Figure P-110 includes system embedded costs of about \$7 billion per year.[41] Including this embedded cost reduces the cost volatility, compared to the statistics in the previous section, but it provides values that more closely correspond to total power costs and retail rates.

In Figure P-110, the horizontal axis are cost increases calculated by dividing each year's costs by the costs in the first year of the study:

$$\frac{C_i}{C_1}$$

This provides some insight into how the costs vary with respect to current circumstances. The graph suggests that there is significant improvement in moving from the least-cost Plan A to Plan B. In particular, the likelihood of cost increases exceeding 30 percent is half of that for the least-cost plan. Plans B, C, and D (the least risk plan) all have comparable cost volatility.

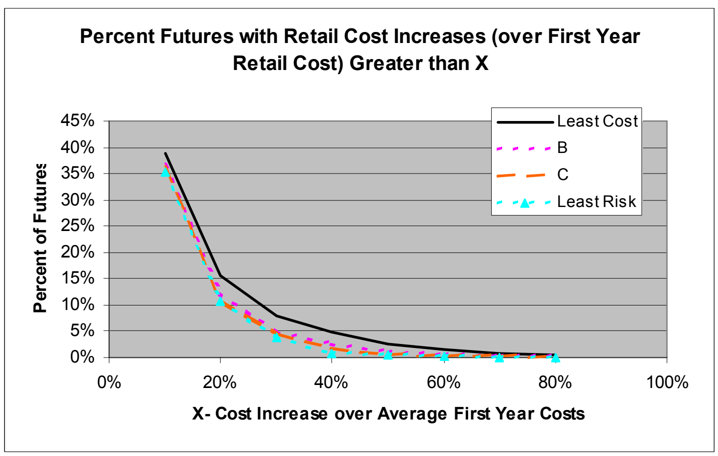

Figure P-110: Exceedance Probability: Increase Over First Year's Cost

An alternative way of measuring cost variation is to look at the difference in costs from year-to-year and compare that change to costs in the first year of the study:

$$\frac{C_i - C_{i-1}}{C_1}$$

This provides an idea of rate shock, while "normalizing" the denominator. Without normalizing the denominator, cost increases expressed as percentage change would

[41] Staff attempted to adjust the embedded costs from year-to-year for depreciation. In real terms, these costs decreased by 3 percent per year.

appear to be different when the change in annual cost expressed in dollars is the same. The results for this analysis appear in Figure P-111. They suggest the same conclusions as the previous figure, although the halving in likelihood is now for percentage cost increases over 40 percent, instead of 30 percent.

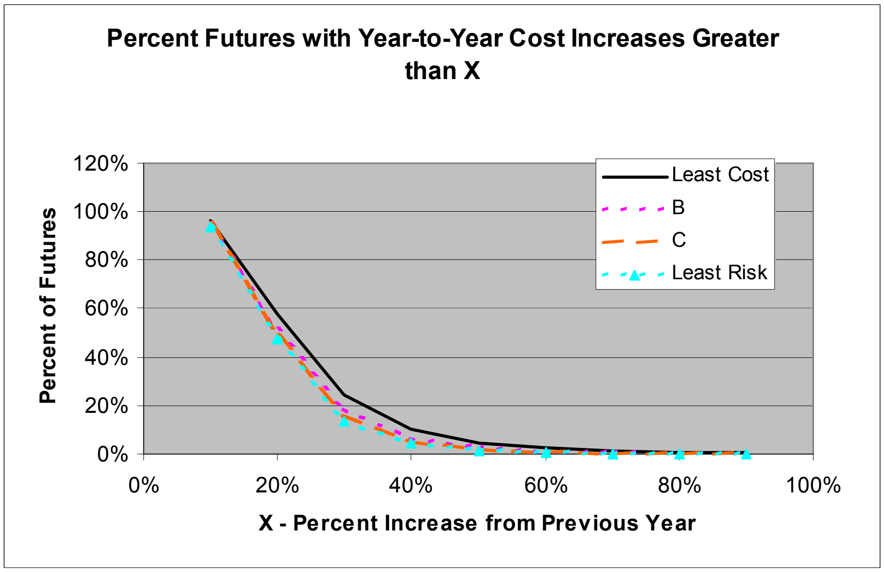

Figure P-111: Exceedance Probability: Changes Over First Year's Cost

Finally, Figure P-112 uses a simple cost change from year-to-year:

$$\frac{C_i - C_{i-1}}{C_{i-1}}$$

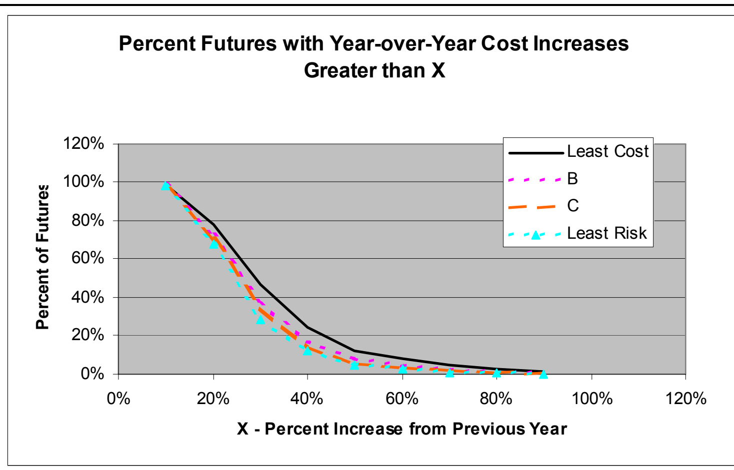

Figure P-112: Exceedance Probability: Changes Over Previous Year's Cost

The conclusions from this figure would be the same as the previous one.

By these measures, we see substantial reduction in cost volatility in going from the least-cost plan to any of the other three plans. Cost volatility among the three lower risk plans clearly decreases as TailVaR$_{90}$ risk decreases, but those three provide roughly similar results.

It is reasonable to ask what is driving the cost volatility. Figure P-107 and Figure P-108 suggests that the fix costs associated with new power plants are not the source of cost variation. In fact, plans with more resources seem to have less cost variation. This points to a source of risk that is prominent among the Council's concerns: electricity market price risk. While market price uncertainty can contribute to risk, it is not in itself a source of bad outcomes. The region needs to be in a deficit situation and importing energy for high market prices to produce larger cost and rate volatility.

Imports and Exports

Figure P-113 [**35**] shows the difference in substantial imports between the least-risk plan and the least-cost plan. For the purpose of this illustration, substantial imports are those exceeding 1,500 MW-quarters. Imports are identical until about the year 2013, when power plants begin to appear in the least risk plan. As expected, there is more import in the least-cost plan, exposing the region to high electricity market prices.

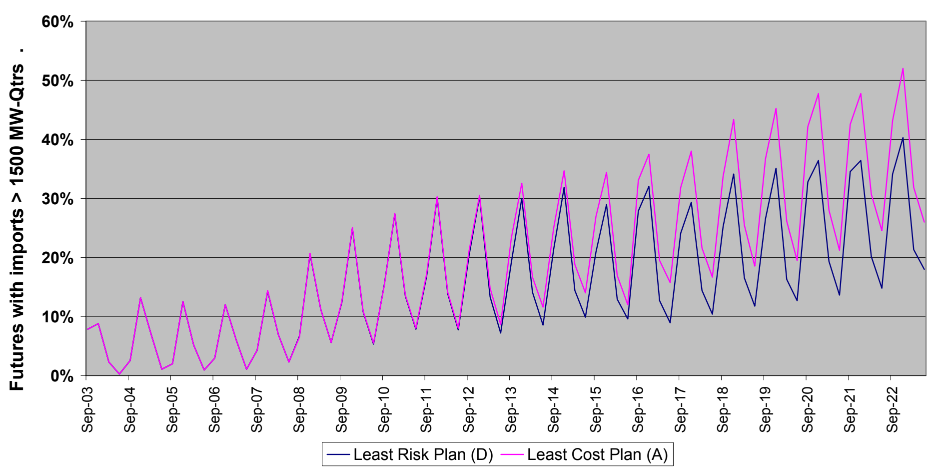

Figure P-113: Frequency of Futures with Large Imports

Exposure to Wholesale Market Prices

Figure P-114 [**35**] examines specifically those events where there are both substantial imports and high market prices for electricity (over $100 per MWh). This figure suggests several conclusions. The least-cost plan introduces substantially greater likelihood of incurring costs associated with high market prices than the least-risk plan. This is due to both the higher likelihood of high market prices with the least-cost plan and the higher likelihood of substantial imports. Plan B reduces the likelihood of exposure to expensive imports by half, but not to the extent of the least-risk Plan D. It is also notable that the least-risk plan maintains the likelihood of regional exposure to wholesale market prices at roughly the same level throughout the study.

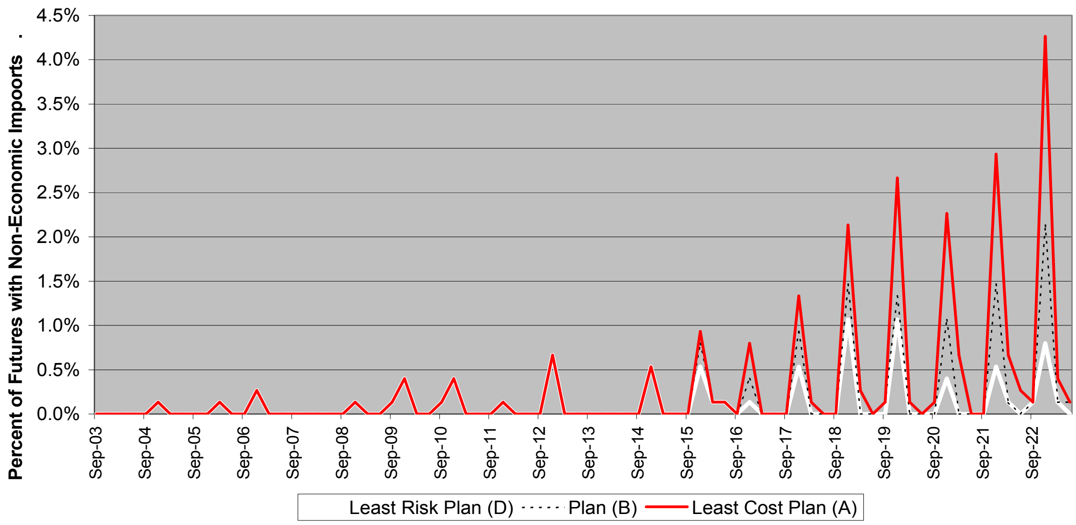

Figure P-114: Likelihood of Exposure to Non-Economic Imports

This concludes the analysis of cost volatility among plans in the feasibility space, and among plans on the efficient frontier, in particular. This analysis suggests that plans on the efficient frontier do not have the least cost volatility, but they do possess moderate cost volatility. Economically inefficient, resource-surplus plans have lower cost volatility. Among the plans on the efficient frontier, cost volatility decreases with plan risk, as measured by $TailVaR_{90}$. Most of the volatility, however, diminishes passing from the least-cost plan (Plan A) to Plan B (Figure P-103). Plan B has substantially more wind development than Plan A, and it lacks the IGCC plant and late CCCT development of Plans C and D.

Engineering Reliability

Many of the concepts introduced with the regional model are new to decision makers in the regional power planning community. Economic risk metrics, in particular, may be unfamiliar. As we will see, economic risk metrics appear to be more conservative than engineering risk metrics, in the sense that plans with satisfactory economic risk are also adequate and reliable. Nevertheless, there is no guarantee that a plan that has good economic characteristics must have high reliability from an engineering perspective. It stands to reason that decision makers will want to confirm that plans along the efficient frontier meet traditional measures of engineering reliability.

Energy Load-Resource Balance

It is challenging to relate the results from the regional portfolio model to other system planning models. Other models cannot capture certain events and behaviors, such as the regional model's dynamic reaction to unforeseeable futures. To better communicate the results of the regional model, Council staff nevertheless examined questions typically put to system planning models like, "What is the loss of load probability associated with this plan?" or "What kind of a energy resource-load balance does that plan produce?"

The last question is the genesis of this section. At first glance, answering the question should be easy. There is, after all, a plan of construction and an expected load forecast. The difference between these, expressed in energy, should characterize the resource-load balance, shouldn't it? Actually, no, because the plan is a schedule of earliest construction. Which and how many plants eventually come on-line and the energy requirement both depend on the future.

Moreover, because the plan is essentially a schedule of options to build resources, the number and size of resources grows relative to the expected load. That is, the energy resource-load balance – what we occasionally refer to as the "energy reserve" – is growing relative to the load. The reason for the growth in reserve is that, further out there is greater uncertainty. With growing uncertainty about fuels, loads, taxes, and so forth, it becomes cost effective to have more options to respond to that uncertainty. That is, the plan may have both a coal-fired and a gas-fired power plant as options in outlying years because the model will develop one or the other, depending on circumstances, but presumably not both.

One way to make the regional model results to a certain extent comparable is to examine the energy reserve on a future-by-future basis. For the recommend plan, a sample on annual energy reserve from the 750 futures appears in Figure P-115.

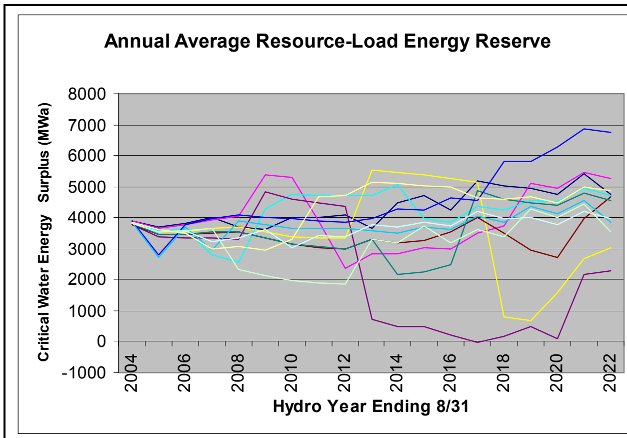

Figure P-115: Resource-Load Balance for the Plan

What the figure shows are 12 futures with wildly varying reserve margins, which can rapidly change from relatively high, positive values to negative values. These sudden excursions are typically associated with business cycles, the return or departure of smelters, changing contract levels, and power plants coming into service. Major sources of load and production variability, weather and hydrogeneration stream flows, do not influence this picture. This figure reflects a planning energy reserve margin. Planning studies typically disregard those sources of variation. Instead, this figure shows energy reserve margin using weather-

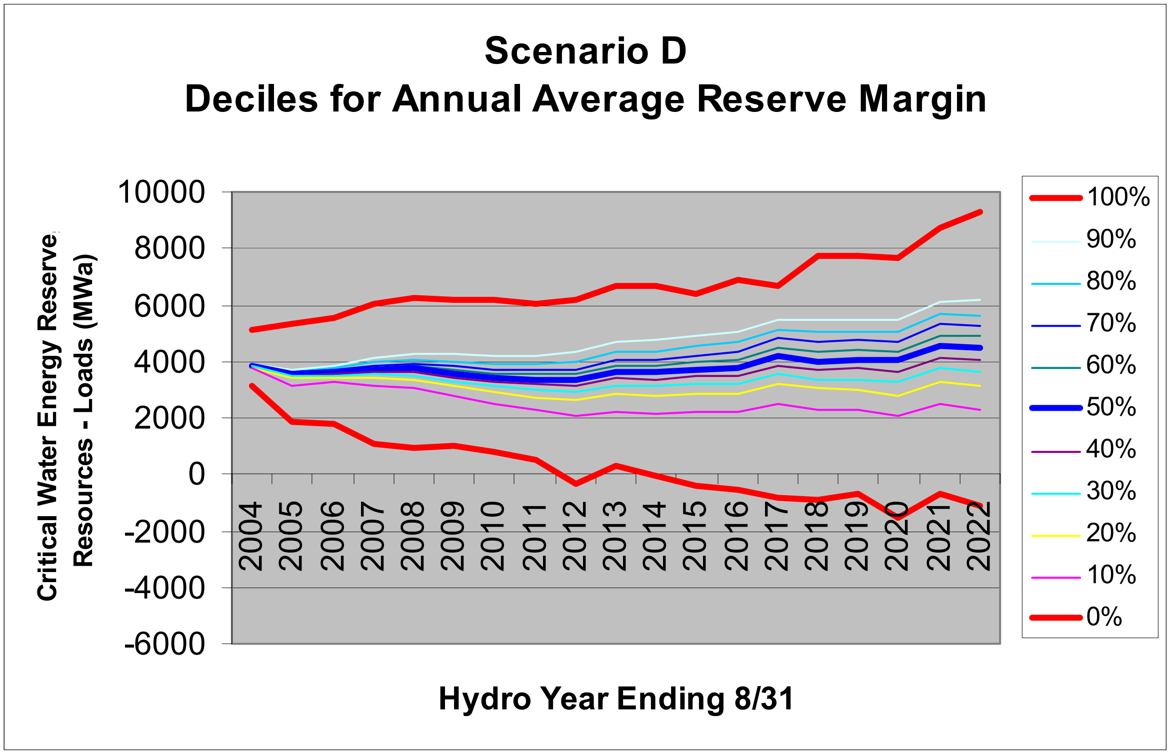

Figure P-116: Statistical Summary of Reserve Margin, Least-Risk Plan (D)

adjusted loads and critical water assumptions. (Critical water is the lowest hydrogeneration energy due to historical stream flow variation.)

Figure P-115, however, does not provide a sense of what kind of patterns may exist over all the 750 futures. To see those patterns, statistical summaries are necessary. We emphasize again here that these statistical summaries may be misleading and require interpretation. (See the subsection "Comparison with the Council's Load Forecast" of the section of the "Uncertainties" chapter dealing with "Load.")

Figure P-116 shows the recommended plan's quarterly deciles for critical water energy reserve [**36**]. Chapter 7 of the Council's Fifth Power Plan discusses four resource plans selected from along the efficient frontier. These plans are illustrated in Figure 7-2, which is reproduced below (Figure P-103) for easy reference. Figure 7-2 refers to the recommended plan as Scenario D.

What is evident from Figure P-116 is the median energy reserve stays about where it is today, perhaps a few hundred average megawatts higher. This is consistent with the observation that the region is currently surplus of resources on an expected value basis. Also, the upper and lower bounds, the "jaws" so to speak, become wider farther along in time. This illustrates one of the facts highlighted earlier in this section: greater uncertainty merits greater contingency planning.

Finally, the lower jaw moves into negative energy reserve only in outlying years. From studies that the Council has performed on regional reliability, a deficit of 1,000 MWa would still produce a reasonably reliable system as measured by loss of load probability. This graph suggests that economic reliability is more conservative, requiring more total resource, than engineering reliability measures. This stands to reason, because engineering reliability ignores the costs of the plants providing such reliability. In fact, inefficient or costly resources will be supporting the system during shortages.

Several technical assumptions are material to interpreting these figures. First, IPP energy, totally about 3,250 MWa, is included in the reserve margin calculation. Several regional planning organizations, such as the Pacific Northwest Utility Coordinating Council (PNUCC), do not include regional energy not under contract. The portfolio model's primary focus, however, is cost and market price. Disregarding IPPs would distort market price calculations. Moreover, the Council believes uncommitted IPP could be economically diverted in the winter from displacement sales outside the region.

Second, the energy associated with generation resources is discounted by maintenance but not by forced or unplanned outages. This is consistent with industry practice. Finally, the reserve calculation includes firm regional contracts and sales, according to the BPA White Book, and assumes 11,650 MWa for critical water hydrogeneration.

Intuition suggests that lower cost, higher risk plans on the efficient frontier would have lower energy reserves. Working along the efficient frontier through Scenarios C and D

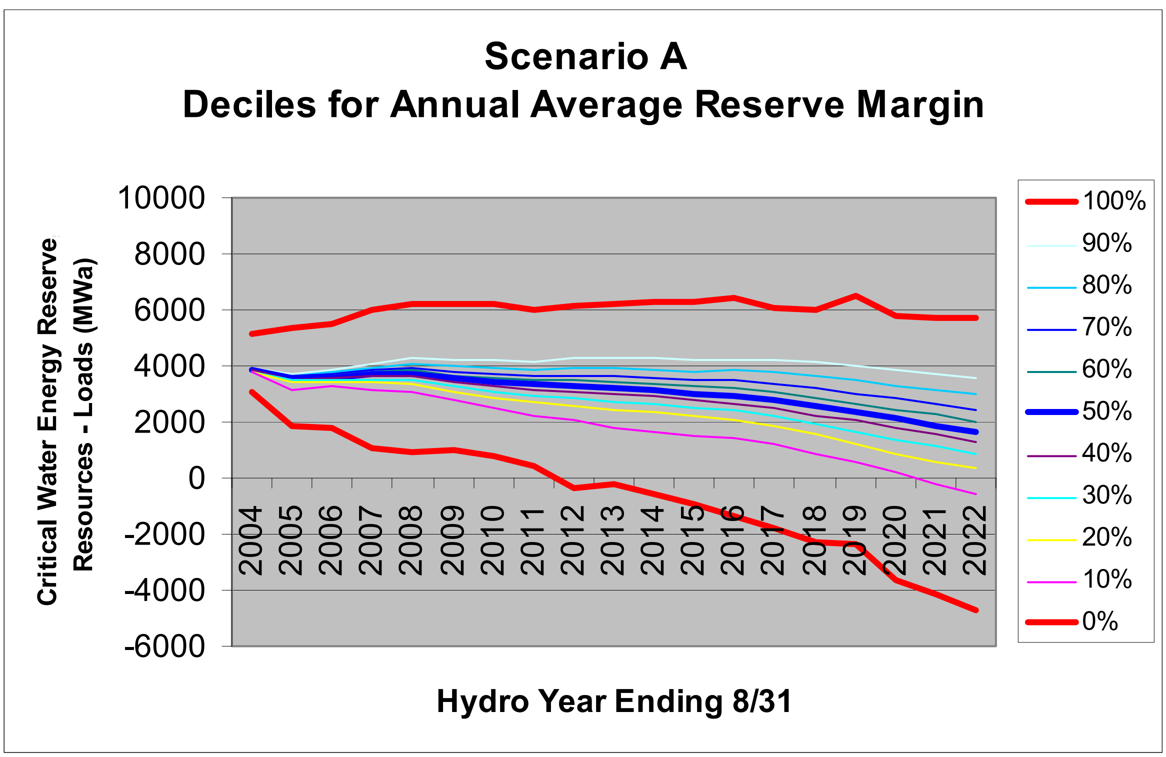

Figure P-117: Statistical Summary of Reserve Margin, Least-Cost Plan (A)

(Figure P-118) to Scenario A (Figure P-117), this pattern is evident. All of the plans start out in a similar situation, which existing resource and load dictate. Only after about the year 2010 do the energy reserves differ significantly. The region is in a surplus situation until then, and no resources or other actions – except for small differences in

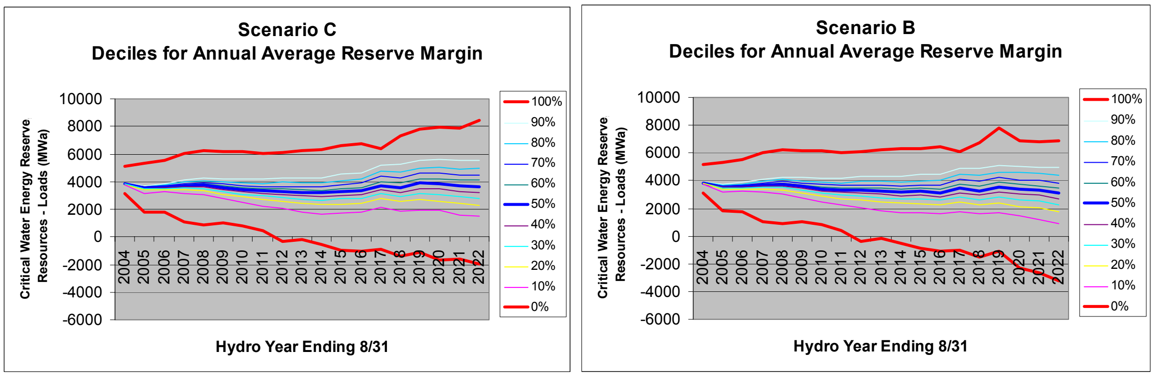

Figure P-118: Energy Reserve Statistics for Scenarios C and B

conservation – differentiate the scenarios.

The least-cost plan, Scenario A, has a reserve margin that falls roughly 2200MWa between 2010 and the end of the study. In the least-cost plan, only inexpensive conservation enters the plan. It is not evident that Scenario A's energy reserve margin stabilizes during the study.

Loss of Load Probability (LOLP)

For the reasons described in the previous section, it's difficult to make a direct comparison of loss of load probability using the regional model to the LOLP computed using a traditional model. The problem stems from the diversity of futures that the regional portfolio model uses to estimate cost and risk. Nevertheless, Council staff used the GENESYS model to analyze Plans A, B, C, and D using a single, representative future [**37**]. For this future, fuel prices and loads are identical to the benchmark values used in the regional model. Conservation and smelter loads are the average values across futures in the regional model. Power plant construction proceeds without interruption, and all power plants are in service on the earliest feasible date.

The three lower risk plans all produced zero loss of load probability across the years in the study. Only the least-cost plan produced nonzero values.

Figure P-119 [**38**] shows the loss of load for the least-cost plan in average megawatt-seasons on the vertical axis and the exceedance probability of the horizontal axis. There are five exceedance curves in this figure, corresponding to the study years 2008, 2010, 2013, 2018, and 2023. The horizontal, heavy black line is the Council's threshold for a significant event. The Council considers events smaller than 10 megawatt-seasons too small to be of concern. System operators can probably take some extraordinary measure to deal with such events, short of curtailing loads. The vertical, dashed red line is the Council's threshold for event likelihood. In principle, it is impossible to build a completely reliable system. Therefore, it becomes necessary to define a likelihood below which loss of load events are acceptable.

The Council considers a plan reliable if events do not simultaneously exceed the two thresholds. Referring to Figure P-119, the reliability of the system in a year corresponding to the curves is adequate if the curve does not enter the upper right hand quadrant defined by the two thresholds.

Clearly, the least-cost plan (Plan A in Figure P-103) is adequate by this definition in every year. The maximum loss of load probability associated with the least-cost plan occurs in those two years just before an early combined-cycle and wind power plant come online and again near the end of the study. Loss of load probability, by the Council's definition, reaches 4 percent in those two years.

This section on engineering reliability opened with a comment that economic risk assessment appears to be more sensitive than engineering reliability assessment. In this section, studies show that plans – even least-cost plans – that are on the efficient frontier pass engineering reliability planning criteria. This stands to reason because engineering reliability criteria, such as those presented in this section, ignore cost. Engineering criteria use prices to assure that the system operates in a realistic manner, using merit dispatch, but if load is not lost, there is no penalty. Economic risk metrics, on the other hand, will warn planners in advance that the last remaining, most expensive resources in the supply stack are maintaining reliability. Sufficiently high prices and penalties, moreover, will signal any event relevant to engineering reliability. In this sense, then, we conclude that economic risk assessment tend to be more sensitive than engineering reliability assessment.

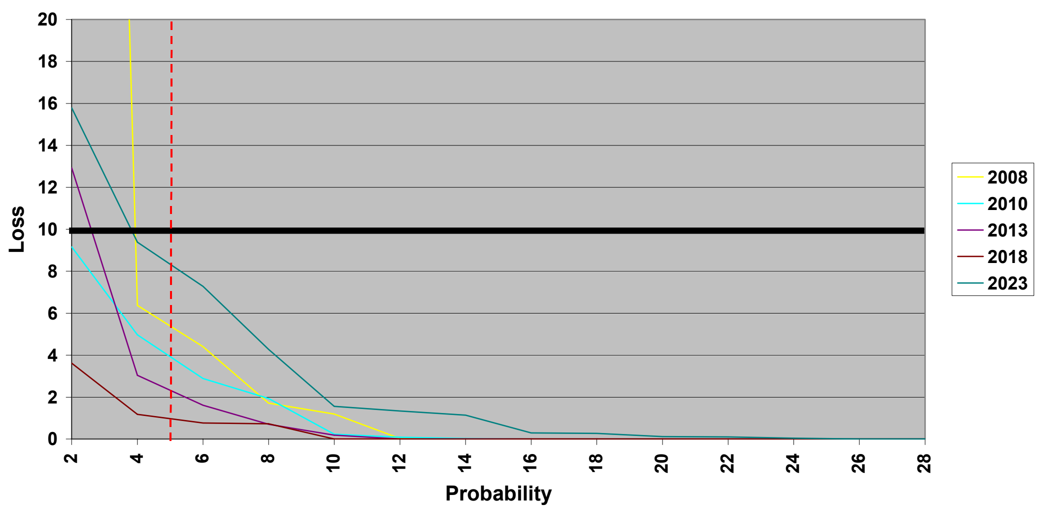

Figure P-119: LOLP for Least-Cost Plan

A Final Risk Consideration

Reviewing the results presented in this chapter, it would be reasonable to choose Plans B, C, or D. Plan A clearly has more risk and cost volatility. While Plan D has lowest risk, Plans B, C and D have comparable performance. All three plans call for substantial amounts of wind, which is absent in Plan A. Plan C adds more CCCT capacity later in the study; Plan D begins the construction of an IGCC coal plant in 2012.

One source of risk not discussed above, however, is the risk of premature commitment. Most planners understand that it would be imprudent to commit to any decision earlier than necessary. More time brings more information and perhaps additional options. This is the reason why plans typically comprise an *action plan*, focusing on the immediate commitments, and the rest of the plan, which addresses activities that can be reassessed and decided later.

The selection of Plan D costs nothing now and reduces premature commitment risk. Specifically, it implicitly calls for re-evaluation of alternatives earlier than would Plan B or Plan C. The coal plant and CCCT units have longer lead-times than do the wind units, and the wind units in Plan D arrive earlier and in larger number. By selecting Plan D, the Council has signaled a re-evaluation of the plan no later than 2009, three years before the earliest construction date 2012. Three years are necessary for the siting and licensing process of a IGCC plant. If the IGCC were to be located in a transmission constrained region like Idaho or Montana, which is a strong possibility, transmission studies need to begin immediately. Transmission has an even longer lead-time. If the Council were to choose Plans B or C, instead, no re-evaluation probably would be necessary until 2012, and transmission may not be as much of an issue. If the region waited until 2012 to evaluate its situation and then discovered it needed an IGCC plant, however, the delay could be costly.

Summary and Conclusions

This section addressed $TailVaR_{90}$ as a risk measure for the region. The section introduced coherence measures of risk and explained their advantages. ($TailVaR_{90}$ as a coherent risk measure.) It explored alternative measures of economic risk and evaluated how representative $TailVaR_{90}$ is with respect to each. It concludes that $TailVaR_{90}$ is representative, in the sense that the other risk measures examined would have produced the same, or substantially the same, choice of plans for the efficient frontier.

This section also examined the plans on the efficient frontier using cost volatility and engineering reliability planning criteria. Plans on the efficient frontier do not have the smallest cost volatility, but they do possess moderate cost volatility. Economically inefficient, resource-surplus plans have lower cost volatility. Among the plans on the efficient frontier, cost volatility decreases with plan risk, as measured by $TailVaR_{90}$. Most of the volatility, however, diminishes passing from the least-cost plan (Plan A) to

Plan B (Figure P-103). All of the plans on the efficient frontier appear to be reliable with respect to loss of load probability (LOLP) and resource-load balance.

Finally, while Plans B, C, and D have similar performance with respect to cost volatility and engineering reliability planning criteria, Plan D permits the Council to minimize premature commitment risk at no cost. For this reason, the Council selected Plan D as its preferred resource plan for the Council's Fifth Power Plan.

Sensitivity Studies

This chapter presents the results of detailed sensitivity analyses. The Council performed over 160 studies to understand how the conclusions of the model depended on assumptions, such as model structure, natural gas price, carbon penalty, the rate of conservation implementation, the feasibility of wind generation, and alternative decision criteria. Valuation studies are a specific kind of sensitivity analysis. The Council performed studies to value conservation, demand response, wind, and the gross value of independent power producers' power plants. Each study requires producing a feasibility space: about 1,400 twenty-year plans, each evaluated using 750 futures. In all, this work represents approximately 160 million twenty-year studies of hourly Northwest power-system operation.

The sensitivity studies appearing in this section do not all use the same base case or model. Because preparing feasibility spaces is time-consuming, typically requiring a day of computer simulation and a comparable amount of time for analysis, this section presents only the last completed studies. This should not be a limitation, however, to understanding the influence or effect in question. The Council performed these sensitivities with several models and varying sets of assumptions. After studying the results from multiple studies, typically a strong and intuitive pattern emerges. This chapter will present those patterns.

The reader should pay *little* attention to the absolute cost and risk values associated with the feasibility spaces, therefore. The base case values will depend on the model logic and assumptions, which may change dramatically from summary to summary. Instead, the reader should pay attention to the *change in location and shape* of the efficient frontier, between the sensitivity case and *its* corresponding base case. Each section below will present these side by side, with the base case illustrated with blue points and the change case illustrated in red points.

In the following, the format of each section will be

- Brief description of the issue
- Description of the workbook modeling
- Results from the efficient frontier
- General observations and conclusions

High Natural Gas Price

In this sensitivity [**39**], the average natural gas price was $1.50/MMBTU higher than in the base case. The purpose is to understand the implications if the median of the natural gas price distribution were higher than used in the base case.

There is no adjustment to electricity price, including through sensitivity parameters, described in the section beginning on page P-69. The benchmark prices in {{row 53}} are $1.50/MMBTU higher.

Figure P-120 shows the displacement of the efficient frontier due to the increased price of natural gas. The base case is in blue (light blue frontier) and the sensitivity case is in red (yellow frontier). Much of this displacement, of course, is due to existing gas-fired resources. The new resources in the plans along the frontier have little influence on total system cost.

What is of interest is the makeup of plans on the efficient frontier. Perhaps not surprising is that wind generation and conservation develop more across the entire efficient frontier, while CCCTs are less popular. More surprising is that coal is not on the efficient frontier in the high-gas price case, and there remains a substantial amount of CCCT siting and licensing. (There is no coal on the efficient frontier in this base case, either, because it is from an early generation of studies.) Despite low coal prices and the fact that there is no uncertainty assumed for coal price, this change in the distribution of probabilities for gas price futures seems to have little effect on the attractiveness of coal-fired generation. That CCCTs remain attractive can be understood from two factors. First, wind generation development is capped, and additional capacity of some sort is required. Second, new gas-fired generation is more efficient than existing gas-fired generation. Consequently, the newer units can economically displace older gas-fired units.

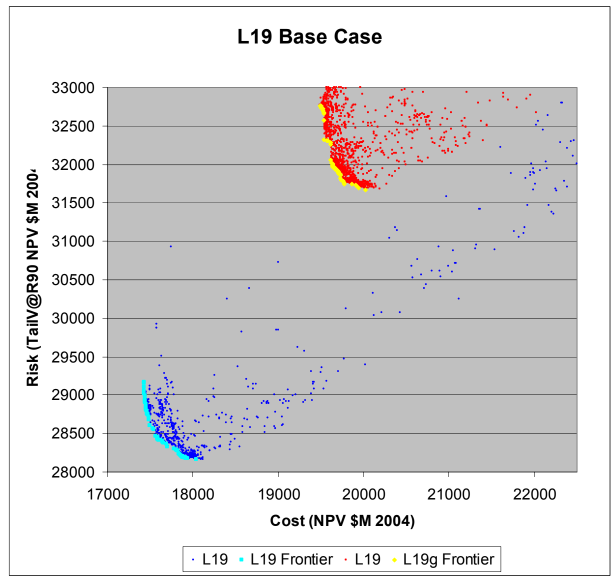

Figure P-120: High Natural Gas Price

Reduced Electricity Price Volatility

Electricity price volatility does not affect the value of all plants or plans equally. Non-dispatchable plants like wind and conservation, for example, are unaffected by such volatility. Only the average price of electricity determines their hourly value. On the other hand, volatility is a major determinant of the value of high-heat rate combustion turbines, such as SCCTs, and of demand response. Volatility also will increase the value of reserve margin strategies. Volatility can affect decision criteria differently. Thus, it is important to understand the influence of volatility assumptions.

In one study [**40**], Council staff cut in half the four parameters that control the jump size and principal factors for the independent term of electricity price:

Principle Factor constant offset (R94): (-0.5, 0.0, 0.5) ← (-1.0, 0.0, 1.0) triangular distribution
Principle Factor growth (R96): (-0.58, -0.33, 0.42) ← (-0.83, -0.33, 1.17) triangular distribution
Jump 1 Size (S99): (0.1.25) ← (0,2.50) uniform

Jump 2 Size (S100): (0.1.25) ← (0,2.50) uniform

As one might suspect, the average cost and risk declined with reduced electricity price volatility (Figure P-121). With this change, the premium for conservation disappeared, except at the most risk-averse end of the efficient frontier. The CCCTs are not developed. Contrary to the situation described at the opening of this section, there is *more* development of SCCTs. The lower probability of futures with high electricity prices would tend to make the fully allocated cost of such power less expensive. Wind develops somewhat less extensively across the efficient frontier, perhaps for the same reason as for other capital-intensive resources.

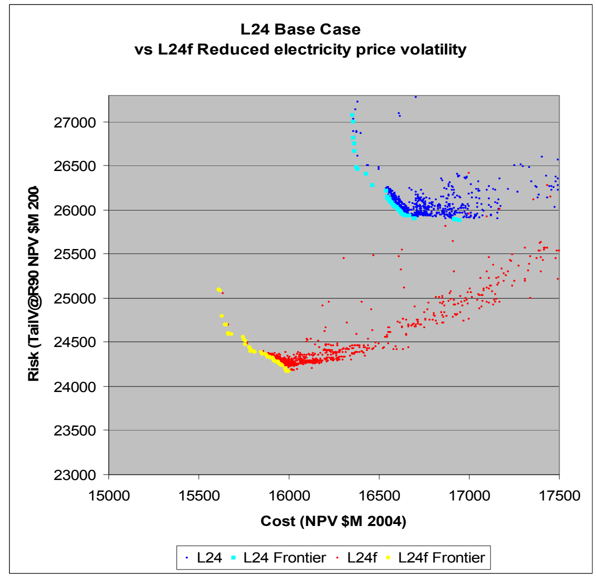

Figure P-121: Reduced Electricity Price Volatility

CO_2 Policy

Because of the prominence of debate over climate change, its possible causes, and its possible effects, the Council performed numerous analyses with alternative assumptions regarding the magnitude and likelihood of a CO_2 tax. (See also the discussion of CO_2 uncertainty.) Some decision makers may not share the view of CO_2 tax uncertainty adopted by the Council. These studies can perhaps help inform those decision makers about the credibility of regional model results.

No CO_2 Tax or Incentives for Wind

One view of the future might be that, scientists will determine that climate change is unrelated to manmade activities. Moreover, clean fossil fuels will become cheap and abundant. There is no chance of any CO_2 tax in this world and renewable energy has no value. Consequently, there is no chance for continuing the PTC, and green tag value falls to zero.

Note that base case modeling for CO_2 tax, PTC, and green tags already allows for *futures* such as this. This sensitivity study, however, posits that there is ***no possibility*** of positive values for these uncertainties.

Four separate studies examined the consequences of this set of assumptions. The latest [**41**] found new wind generation constructed in less quantity and much later, if at all, along the efficient frontier. Instead, a modest amount (400MW) of coal-fired capacity can begin construction around 2013 in about half the plans, those nearer the least-risk plan. The plans also have greater incentive for and more extensive deployment of lost-opportunity conservation. There is slightly less CCCT development at the least-cost end of the efficient frontier and slightly more development at the other end.

Many assume that the possibility of a CO_2 tax is coal generation's biggest risk. This study shows that eliminating CO_2 tax alone does not make coal a leading candidate for new capacity, even assuming low and stable fuel cost. The section "Conventional Coal" elaborates on the regional model's study results for that generation technology.

Higher CO_2 Tax

For a study that incorporates higher levels of CO_2 tax [**42**], the Council chose one of the tax scenarios that appear in an MIT analysis of the proposed 2003 McCain-Lieberman Act.[42] The study implements the McCain-Lieberman schedule for CO_2 tax (MIT Study, Table 4, page 17, Scenario 5), which is \$25/ton (2010), \$32/ton CO_2 (2015), \$40/ton CO_2 (2020), all in 1997\$. These levels are converted to 2004\$ by annual inflation of 2.5 percent and are converted to piecewise linear function of time. The resulting schedule appears in Figure P-122. This high level of tax is deterministic and is present in all futures with the same fixed schedule. (The regional model workbook implements the tax by pasting the values in this figure into row {{74}}, the final value for the CO_2 tax future.)

The feasibility space, illustrated in Figure P-123, shifts significantly up and to the right. The additional expected system cost associated with this sensitivity is about \$9 billion (NPV 2004\$). Discretionary conservation takes a big step forward in this sensitivity, increasing both the recommended premium for development and the amount delivered. CCCTs and wind develop extensively, even in the least-cost plans. The incentive for new CCCT capacity is the displacement of older, less efficient units. Wind development by the end of the study for

Inflation annual rate			Conversion to 2004\$
2.5%			1.188686
Calendar Year	Period	1997\$	2004\$
2004	2	1.00	1.19
2004	3	2.00	2.38
2004	4	3.00	3.57
2004	5	4.00	4.75
2005	6	5.00	5.94
2005	7	6.00	7.13
2005	8	7.00	8.32
2005	9	8.00	9.51
2006	10	9.00	10.70
2006	11	10.00	11.89
2006	12	11.00	13.08
2006	13	12.00	14.26
2007	14	13.00	15.45
2007	15	14.00	16.64
2007	16	15.00	17.83
2007	17	16.00	19.02
2008	18	17.00	20.21
2008	19	18.00	21.40
2008	20	19.00	22.59
2008	21	20.00	23.77
2009	22	21.00	24.96
2009	23	22.00	26.15
2009	24	23.00	27.34
2009	25	24.00	28.53
2010	26	25.00	29.72
2010	27	25.35	30.13
2010	28	25.70	30.55
2010	29	26.05	30.97
2011	30	26.40	31.38
2011	31	26.75	31.80
2011	32	27.10	32.21
2011	33	27.45	32.63
2012	34	27.80	33.05
2012	35	28.15	33.46
2012	36	28.50	33.88
2012	37	28.85	34.29
2013	38	29.20	34.71
2013	39	29.55	35.13
2013	40	29.90	35.54
2013	41	30.25	35.96
2014	42	30.60	36.37
2014	43	30.95	36.79
2014	44	31.30	37.21
2014	45	31.65	37.62
2015	46	32.00	38.04
2015	47	32.40	38.51
2015	48	32.80	38.99
2015	49	33.20	39.46
2016	50	33.60	39.94
2016	51	34.00	40.42
2016	52	34.40	40.89
2016	53	34.80	41.37
2017	54	35.20	41.84
2017	55	35.60	42.32
2017	56	36.00	42.79
2017	57	36.40	43.27
2018	58	36.80	43.74
2018	59	37.20	44.22
2018	60	37.60	44.69
2018	61	38.00	45.17
2019	62	38.40	45.65
2019	63	38.80	46.12
2019	64	39.20	46.60
2019	65	39.60	47.07
2020	66	40.00	47.55
2020	67	40.40	48.02
2020	68	40.80	48.50
2020	69	41.20	48.97

Figure P-122: Adapted CO_2 Schedule

[42] Paltsev, S., J.M. Reilly, H.D. Jacoby, A.D. Ellerman & K.H. Tay, ***Emissions Trading to Reduce Greenhouse Gas Emissions in the United States: The McCain-Lieberman Proposal***, MIT Joint Program on the Science and Policy of Global Change, Massachusetts Institute of Technology, June 2003 http://web.mit.edu/globalchange/www/ (Link to file server)

most plans is unchanged, however, because the Council assumes wind is constrained by availability and system integration limits. Not too surprising, coal-fired generation is nowhere near the efficient frontier.

This study is the latest of three performed using different base case assumptions. All studies resulted in roughly the same outcomes.

CO_2 Tax with Varying Levels of Probability

As mentioned in the Uncertainty chapter, some carbon tax is present in about two-thirds of futures. With no CO_2 tax and few incentives for wind, coal-fired generation begins to make an appearance on the efficient frontier. (See the discussion on page P-133.) In early studies [**43**], CO_2 tax was not tiered as it is in the draft and final plans, and the probability of a CO_2 tax was higher. In an attempt to threshold the conditions that favor alternative plans, various modeling studies [**44**] examined the effect of reduced probability of CO_2 tax and increased natural gas price. These studies shaped the representation for CO_2 tax used in the final plan.

Examining the study least favorable to wind generation and most favorable to coal-fired generation, wind still demonstrated a relative advantage. Even with only a 25 percent probability of a CO_2 tax by the end of the study and an increase of $1.50/MMBTU in natural gas prices, no coal plants appeared on the efficient frontier.

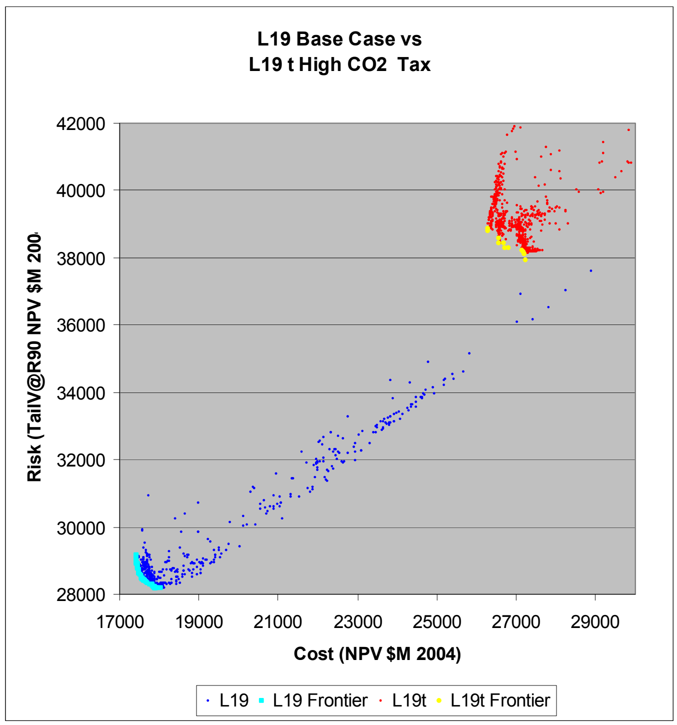

Figure P-123: Feasibility Space for High CO_2 Tax

These studies convinced the Council that in the kind of risk analysis the regional model performs, the "tail events" can be and often are more important than expected value events. The model does not choose coal in plans at the least-cost end of the risk-cost trade-off curve, because relying on the market and not building resources minimizes expected cost. The model does not choose coal in plans at the least-risk end of the curve, often because futures where CO_2 tax appears and planning flexibility is important hurt the performance of such plans.

To model these studies, the uniform distribution in the assumption cell {{R72}}, which controls in which period a CO_2 tax of any size occurs, has a larger range. Extending the upper value of the uniform distribution to 20 from six effectively reduces the chance any tax will start before the end of the study. (See pages P-85 ff for a description of how the model uses this parameter.)

Independent Power Producers

In studies performed before release of the draft plan, the regional model considered Independent Power Producer (IPP) generation as part of the region. During the vetting process, however, the Council realized that this was not consistent with how previous Council plans have defined the region. Specifically, the Council has taken the "region" to be the *ratepayers* in the area specified by the Act. It is also not consistent with how other regional utility planning organization, such as the Pacific Northwest Utility Coordinating Council (PNUCC) account for IPP plants.

Equating the region to its ratepayers is key to how the regional model performs its economic evaluation. The fully allocated costs of power plans belonging to regional utilities eventually pass to regional ratepayers. With public utilities and co-ops, the flow of plant expenses and profits back to ratepayers is relatively direct and evident. For privately held utilities, the flow may be less obvious to some observers. Assuming perfect regulation, the shareholders of private utilities receive only the return of and the return on capital investment in plants, called the "ratebase." These returns occur over time through the utility rates. This means that revenues *unrelated* to ratebase, the profits and losses from power plant operation, flow back to *ratepayers*, not to shareholders. Thus, the economic situation would be the same for a municipal utility.[43]

With IPP generators, however, the situation is different. Profits and losses from power generation of merchant plants flow to shareholders, who are *not* generally ratepayers of the region. The ratepayers effectively pay prevailing market prices for IPP power. (Perhaps forward contracts markets or ancillary services markets are more appropriate than wholesale firm energy markets in a given situation, but the principle is the same.)

With that clarification, the Council changed how the regional model captures the role of IPP plants. In the draft and final plan, IPP plants contribute only to the energy balance in the region. The model ignores IPP costs and profits.

This does not mean, however, that the IPP units have no influence in the results. Because the model constrains regional imports and exports, the model changes electricity prices as necessary to balance supply and demand. (See "The Market and Import/Export Constraints" and, in particular, the subsection "RRP Algorithm" in Appendix L for an explanation of this process.) To the extent that there are additional sources in the region to balance requirements, therefore, the likelihood of lower electric market prices increases. The lower expected market prices, in turn, flow through to the region. Regional utilities will buy and sell into the market to balance their respective load, and

[43] There are, of course, financing, governance, and other differences, but the model tries to deal with those through the calculation of real levelized costs (see Appendix L). The discussion here is only about whether the construction and operating revenues pass to the ratepayer.

the additional IPP generation will extend the depth of supply in that market. The Council believes this approach more accurately models the role IPPs serve in the region.

This situation comports with information provided by the IPP industry and with publicly available data on IPP plant dispatch. Much of the information about the role of IPPs in the region appears in the Overview and Chapter 2 of the Council's Fifth Power Plan. In summary, about 3,000 MW of IPP capacity remained uncommitted as of December 2004, when the Council adopted the final plan. Spot sales into the market remained vigorous, nevertheless, with plants averaging about a 50 percent capacity factor over the previous year.

How does the workbook model represent IPP operations? Appendix L has several pages of description, under the chapter "Resource Implementation and Data." The reader will find both the model's data and formulas there.

This section describes two studies the Council performed to better understand the role of the IPPs in the region. The first looks at the value of IPPs to the region. The second examines the effect of out-of-region purchasers contracted for all of the IPP capacity.

IPP Value

A study [**45**] attempted to estimate how much the IPP generation would be worth to the region, assuming all of the costs and benefits flowed through to regional ratepayers. This might arise, for example, if regional utilities contracted for all the output of the IPP plants.

The study, however, suffers from a serious difficulty. The approach is simply to include the operating costs and benefits of IPP plants, much as in early study work. This is acceptable for existing regional plants, because the construction costs are embedded and do not change from plan to plan or from future to future. For the region, however, the cost of acquiring the IPP plants is not "sunk"; it is not embedded. In fact, depending on the price that utilities would pay to acquire the output of the IPP plants, any benefit will shift between regional ratepayers and IPP shareholders. If the price were high enough, the region might not see any benefit, or it might see a net disadvantage.

The problem does not go away if the study simply chooses an arbitrary allocation of benefit. The modeling issue is more delicate than that. At some price, which we have no easy way to determine beforehand, plans including IPP purchases or contracts will not appear on the efficient frontier of the feasibility space. Substitutes for risk mitigation will become competitive. Moreover, because the IPP plants are not homogeneous – Centralia coal plant is among them, for example – plants will not appear on the efficient frontier in aggregate. The problem would then become one of determining a threshold acquisition price for each resource. That threshold price, in turn, depends on which other IPP resources appear on the efficient frontier. Of course, nothing assures us that the price the region might be willing to pay for a plant's output would be acceptable to the current owner.

There are other difficulties. Acquiring IPP output can adversely affect utility financing, for example. Chapter 2 of the plan mentions some of the more prominent reasons why utilities might not choose to contract for IPP output.

With these caveats, this sensitivity study implicitly assumes IPP owners give the plants to the region for free. There is no acquisition cost. This sets a (rather unrealistic) cap on the potential value of acquiring the plants for the region. Figure P-124 illustrates the reduction in cost (about $4 billion) and risk.

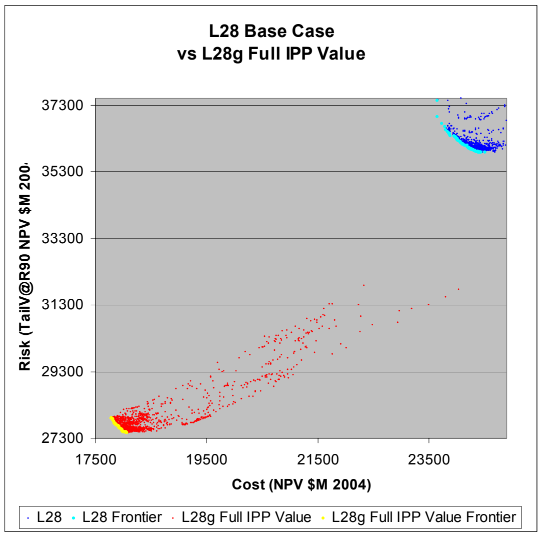

Figure P-124: Acquiring the Full Value of IPP Plants

Purchasing the output of the IPP plants pushes off most of the schedules for new resources, including conservation and wind. Because of their reliance on fossil fuels and natural gas in particular, however, IPP units cut the schedule of wind by half, but some wind remains on the frontier. In the least-risk case, 2,500MW of wind appears before the end of the study.

In the workbook, capturing the cost and benefit of the IPPs amounts to modifying the NPV cost adjustments described in Appendix L. Appendix L uses the example of the on-peak values for the surrogate plant "PNW West NG 3 006" which appear in row {429}. From the Appendix L discussion of valuation costing and of the thermal dispatch UDF, the value is the negative cost appearing in this row. The formula in cell {CV429} discounts these values to the first period:

=0.434512325830654*8760/8064*NPV(0.00985340654896882,$R429:$CS429)*
(1+0.00985340654896882)

The factor of roughly 0.4345 discounts the value of the plant, because about 43.45 percent of the plant belongs to a utility in the region and the rest of the plant (56.55 percent) is IPP. For the sensitivity study, this leading coefficient becomes 1.0, as do those for any plant that is partially or completely IPP.

Contracts for Sale of IPP Energy Outside of the Region

Participants in the public process of reviewing the regional resource plan asked, "What would happen if the output of the IPP plants were contracted outside the region?" The concern is that the region might suddenly find itself substantially short of resources. If that situation were possible, the region might need to acquire additional resources as protection against that contingency.

The initial approach to modeling this situation was simply to add contract exports and to make corresponding counter-scheduling adjustments to import/export constraints (see "Contracts" in Appendix L). This creates a heavier load for region, which the IPPs should incur. There must be an additional adjustment, however, to the economics. Increasing the load alone increases the *economic* obligation of the *region*. That is, the region sees an increase in the load it must serve. This is incorrect, however, as the IPPs are incurring the economic obligation, not the region. Thus to correctly model the situation, the energy out of the region must be increased, but the load used for economic value (or cost) of existing contracts to the region should remain *as in the base case*. With this representation, any effect for the region is due to electricity market price increases due to IPPs no longer contributing to the market.

Figure P-125 shows the results of this study [**46**]. Effectively, the sensitivity case and base case feasibility spaces are lying on top of one another. Within the repeatability of this tool, there is no discernable difference. The plans along the efficient frontier are essentially the same, as well.

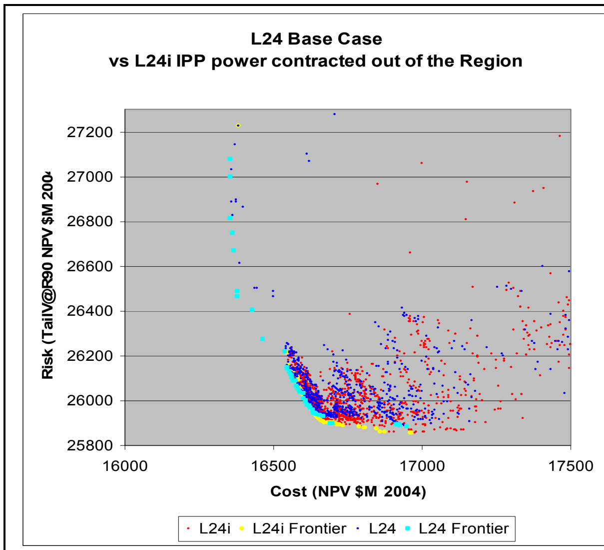

Figure P-125: IPP Energy Contracted Out-of-Region

Why would there be no change? At least market prices should rise, as mentioned earlier, increasing the cost to the region. In fact, what happens is that the model counter schedules contracts. The final dispatch of IPP units does not depend on the contract terms or initial contract obligations, but only on the IPP plant economics relative to the other plants in the region, i.e., a plant's place in the system merit order. The market price for electricity, in turn, depends primarily on the dispatch of the plants in the region. (See Appendix L for a discussion of economic contract counter-scheduling.) Therefore, market prices are unaffected by contracts. Indeed, this is the reason why many simulation models, such as Aurora, can and do ignore contracts.[44]

Contracts do make a large difference to the parties of the contract, of course. The difference is financial, however, and Appendix L shows it is in the economic self-interest of the supplier to re-dispatch units whenever physical constraints are binding or plant economics are out of merit order.

The workbook modeling for this sensitivity study is involved. First, contract sales, both on- and off-peak, increase by the combined seasonal output of the IPP plants. The original level of sales, however, remains in the workbook for the economic costing calculation. The net cost of contracts will be the net position times the prevailing market price. (The study assumes contract cost of energy is fixed and embedded. This would be the case for a forward contract. The *net* value of contracts is the difference between this fixed cost and the value of the energy.) The seasonal IPP capacity is in Figure P-126.

	Fall	Winter	Spring	Summer
IPP cap	3259	3469	2547	2939

source: IPPs Removed.xls

Figure P-126: Seasonal Distribution of IPP Capacity

Second, the seasonal capacities reduce the import values for contracts. The model uses the original values, however, for the cost calculations described above. Figure P-127 illustrates the original and adjusted values for contracts (MWa) on- and off-peak. The on-peak values appear in rows {{83 and 84}} and the off-peak are in rows {{87 and 88}}. The differences between the original and adjusted values are the numbers in Figure P-126. (The difference cycles among the seasonal values throughout the study.) The on-peak energy values are negative, representing net sales out of the region. The original off-peak values are positive, representing net imports, and the adjusted off-peak values are negative, representing net sales.

Figure P-127 also shows the on-peak energy (MWh) and cost ($M 2004) calculations in rows {{367 and 368}}, respectively. The energy calculation uses the adjusted values; the costs use the original. In the formula for cell {{U368}},

= -1152*U83*U204/1000000

[44] A handful of models, such as the Henwood's PROSYM® model, do model contracts because they need to capture pre-dispatch commitment costs due to reliability provisions in transmission and capacity contracts, or because their results will be used for production costing, where financial arrangements are important.

the reference to {{U204}} is the on-peak price for energy, and 1152 is the number of on-peak hours in the hydro season. (Appendix L provides a more complete description of this formula and conventions.)

U368 = = -1152*U83*U204/1000000

	Q	R	S	T	U
82					
83	Original On-Peak Contracts	-996.63	-205.46	-657.54	-1147.22
84	Adjusted On-Peak Contracts	-4,255.74	-3,674.89	-3,204.04	-4,086.42
85					
86					
87	Original Off-Peak Contracts	247.50	759.93	372.54	309.84
88	Adjusted Off-Peak Contracts	-3,011.61	-2,709.50	-2,173.96	-2,629.36
366					
367	Fixed Energy ID: Reg Contracts 8 SubPer_003	-4902614.1	-4233478.4	-3691053.4	-4707557.2
368	Cost ($M)	54.3	8.4	32.0	42.8

Figure P-127: Contract Energy and Value

The formulas for energy balance and for total study cost point to rows {{367 and 368}}, just as before.

Reduced Discretionary Conservation

Many questions about the representation and assumptions regarding conservation arose during the studies that led up to the final plan. These questions included:

- How are decisions about conservation programs made?
- How can the model capture program diversity with a simple supply curve? It is not economic to develop only the least expensive conservation programs, as the supply curve approach assumes. Typically, a utility or customer implements a variety of programs when an opportunity arises to do so. Instead, these programs have a mix of different cost-effectiveness profiles.
- How should the model represent the fact that not all of the energy efficiency programs are mature and that they will mature at different times?
- How does the efficient frontier change as a function of the premium paid for conservation over the "myopic" cost-effectiveness standard?
- Is there value in sustained orderly development, and if so, how do we estimate that value?
- To what extent does the rate of deployment affect the cost of a measure?
- What is a reasonable rate of deployment for discretionary conservation, where large amounts of discretionary conservation are cost-effective?

Circumstances forbid sharing all of these studies. The last study, however, is especially important to the final plan. The Council incorporated the results of this study into the base case.

By definition, the region can pursue discretionary conservation at any time. The Council has estimated the amount that exists in the region, and much of it is cost effective today. In early studies, the regional model controlled the rate of deployment, and the model would choose thousands of average megawatts of this conservation in the first periods of the study. This is unrealistic behavior for several reasons, not least among them being the limited resources utilities have to pursue conservation.

Chapters 3 and 7 of the final plan describe the issues that the Council faced in deciding how rapidly the model could pursue discretionary conservation. Ultimately, this is an educated guess. This section presents some of the quantitative information the Council used to arrive at its conclusions.

The associated studies [**47**] examined three rates of discretionary conservation development: 10 MWa, 20 MWa, and 30MWa per quarter. The Council began examining the effect of these levels over six months before issuing the draft plan, and then checked the results again with the model used to prepare the final plan. The results here are from the final plan studies. By the time the Council had released the draft plan, the base case adopted the discretionary conservation deployment rate of 30MWa per quarter rate.

Figure P-128 shows how the feasibility space changes as the rate of acquisition moved to 20 MWa per quarter (L28b "Mildly Restricted Conservation") and to 10 MWa per quarter (L28c "Least Conservation"), from 30MWa per quarter. The policy of pursuing 30MWa per quarter appears to facilitate plans that are both less risky *and* less costly.

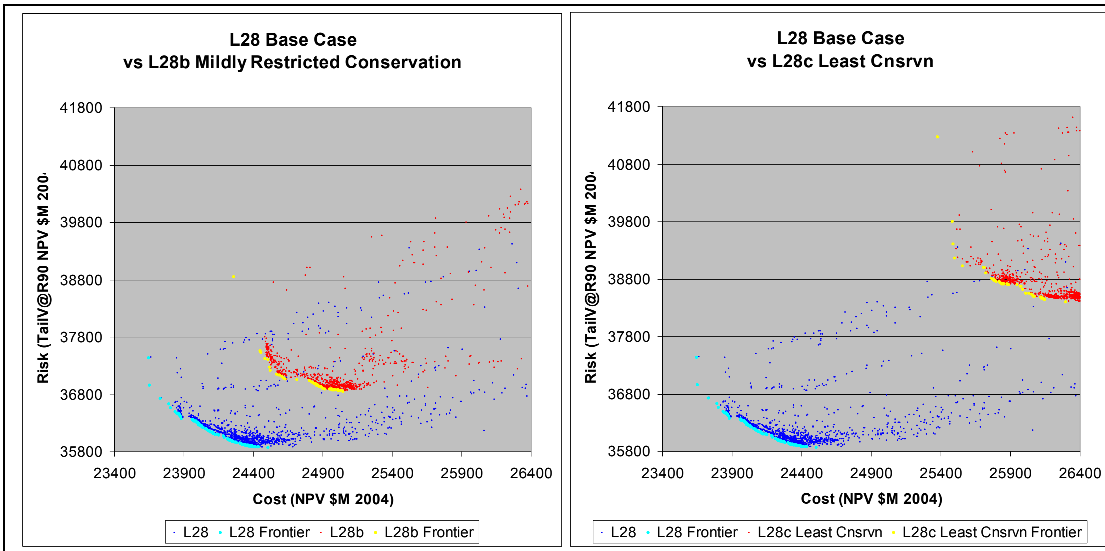

Figure P-128: Discretionary Conservation Ramp Rate Sensitivity Cases

To implement this sensitivity in the workbook, the study modified a supply curve modeling parameter. The cell {{J386}} controls the quarterly ramp rate. Figure P-129 shows the set up for the 20MWa per quarter case. A description of how the model represents discretionary conservation with a supply curve is in Appendix L.

	E	F	G	H	I	J	K	L
381								
382	DHF(0=Dis	Fixed Ener	Fixed Cost	Fuel Set (ID	Heatrate (E	Planning Flexibility ID (	Capacity ID (II	Cap_De
383	3	One	(none)	(none)	0	Conservation Annual_	Consv New C	0,3000,5
384								
385		CurveType	Treatment	Upper Pric	Lower Pric	Max Ramp	Init Inc En	Init Curr
386		0	1	600	0	20	0	
387								
388								

Figure P-129: Supply Curve Ramp Rate Specification

Value of Demand Response

Like the situation for discretionary conservation, early studies suggested that the regional model would take unrealistically large amounts of demand response (DR) immediately, given the assumptions in the model. The Council therefore chose to constrain DR to levels of development it deemed reasonable. Because constraining DR to these levels essentially fixed deployment of DR at these levels, the Council eventually decided to fix the DR-deployment pattern. Using a fixed pattern saves time by relieving the regional model's plan optimizer from examining plans known to be inferior.

One issue that interested Council staff, however, was the value of DR. In the discussion of IPP valuation, simply adding the IPP energy value and variable costs to the region's budget did not permit the Council to accurately capture the value of the IPPs to the region. The reason is that such an approach ignores IPP acquisition cost to the region. With DR, however, the model includes an acquisition cost. Because the Council fixes the DR-deployment pattern, the study finesses the question of whether a given acquisition cost would affect the deployment decision.

DR appears in regional studies as a simple dispatchable resource. It has low capital cost, and a fuel/dispatch cost corresponding to the payment for which the Council assumed loads might voluntarily remove themselves. The model represents DR as a combustion turbine with a fixed $150/MWh dispatch cost. The capital costs, however, are low: about $2.26 per kw-year real levelized [**48**]. (See Appendix L.) The Council has an action item in the final plan to study and refine its cost and availability information about DR potential in the region. Eventually, modeling will mature into a supply curve approach that reflects the short- and long-term diversity of costs among options.

One study [**49**] of DR evaluated the impact of removing DR entirely from the study. The change in feasibility space suggests that, in contrast with many other resources, DR retains its value at the least-cost end of the efficient frontier. (See Figure P-130.)

Most resources provide little value at the least-cost end of the efficient frontier because building new plants is no better than relying on the market. (Recall from the discussion for electricity price uncertainty that electricity market price is the same as the fully allocated cost of power plants in equilibrium. Appendix L, in the chapter on "General Paradoxes," and Chapter 6 elaborate on this principle.) If the region plans to build fewer resources, however, electricity prices become more volatile. This is precisely when DR becomes most valuable.

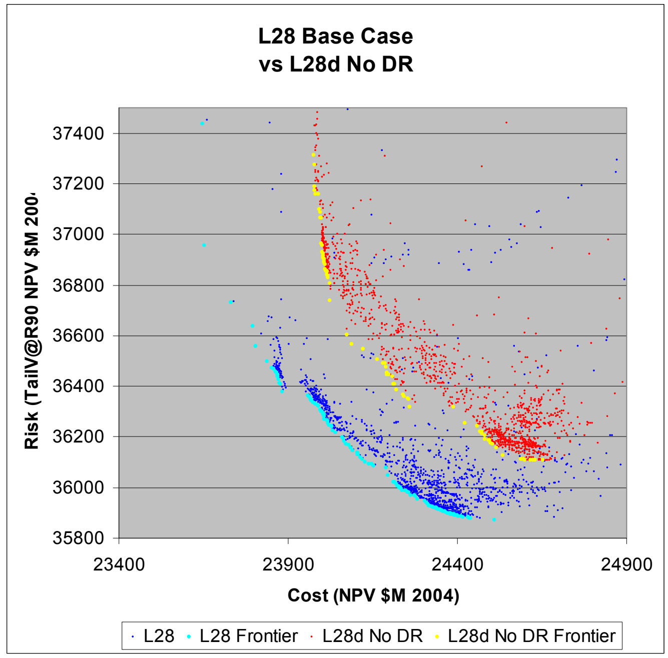

Figure P-130: Removing Demand Response from the Base Case

We can refine the valuation of the DR by comparing this sensitivity case against one where the study holds constant the level of DR across all years. For the purposes of valuing DR, the base case model is poor because the amount of DR is increasing over time. In a study [**50**] where DR is fixed at 500MW, the reader will find a similar pattern as before. Figure P-131 plots the horizontal shift in the efficient frontier as a function of the risk level [**51**]. Again, at the least-cost end (right end of graph) the value of DR increases. Over most levels of risk, the benefit is between $150M NPV and $200M NPV (2004$). This corresponds to $300 to $400 per kilowatt of benefit, net of program costs. The data is rather noisy at this level of resolution, so a fit polynomial in Figure P-131 reinforces the pattern.

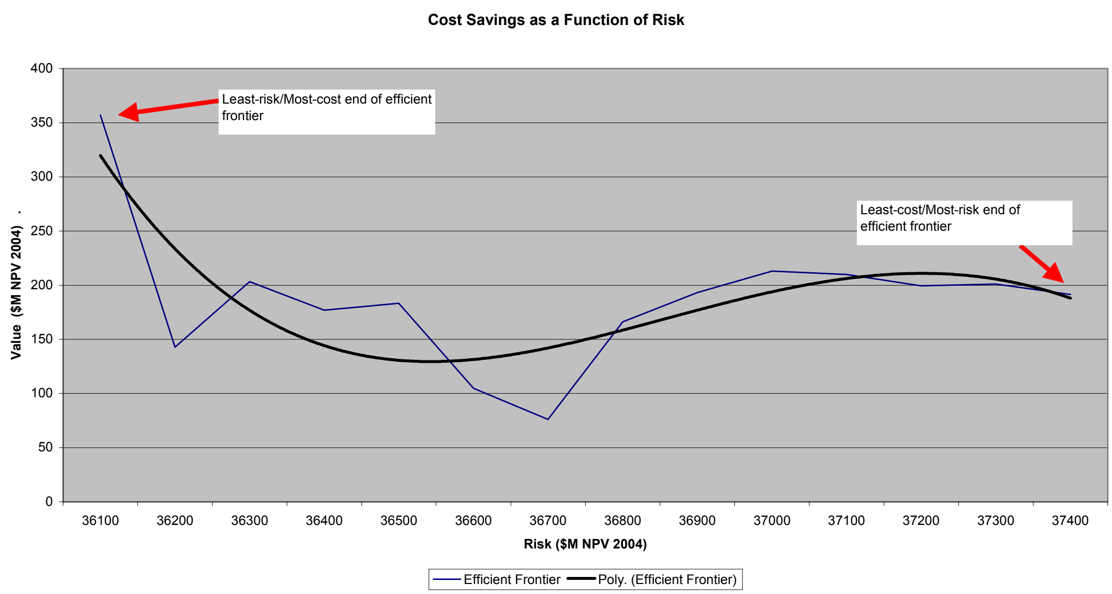

Figure P-131: Cost Savings as a Function of Risk Level (500MW DR Case)

Creating the workbook models to perform these studies was simple. As described in Appendix L, Crystal Ball decision cells determine the capacity for new resource candidates. The model considers DR a new resource, and the decision cells appear on row {{7}}, labeled "PRD" for *price responsive demand*. Figure P-132 shows the situation for the base case. In the base case, DR increases over time, and the values in the decision cells indicate the cumulative number of megawatts of capacity for the new resource option.

In the base case and in the sensitivity cases, the model removes control of the decision cells for DR from Crystal Ball. For this reason, the cells do not have the yellow background that other decision cells have.

PlnCap_0 = 0

	N	R	S	AH	AI	AJ	AP	AQ	AR	AX	AY	AZ	BF	BG	BH
1		Sep-04	Dec-0	Sep-08	Dec-08	Mar-	Sep-10	Dec-10	Mar-1	Sep-12	Dec-12	Mar-1	Sep-14	Dec-14	Mar-
2	Capacity Data ID														
4	CCCT Capacity	0.00			0.00			610.00			610.00			610.00	
5	SCCT Capacity	0.00			0.00			0.00			0.00			0.00	
6	Coal Capacity	0.00			0.00			400.00			400.00			400.00	
7	PRD	0.00			500.00			750.00			1,000.00			1,250.00	
8	Wind 1	0.00			0.00			1,200.00			1,200.00			1,200.00	
9	Wind 2	0.00			0.00			0.00			0.00			0.00	
10															

Figure P-132: Decision Cells for Base Case

To simulate the situation without DR, it is necessary only to set the cumulative capacity of DR to zero in all periods. This is illustrated in Figure P-133.

PlnCap_0 = 0

	N	R	S	AH	AI	AJ	AP	AQ	AR	AX	AY	AZ	BF	BG	BH
1		Sep-04	Dec-0	Sep-08	Dec-08	Mar-0	Sep-10	Dec-10	Mar-1	Sep-12	Dec-12	Mar-1	Sep-14	Dec-14	Mar-
2	Capacity Data ID														
4	CCCT Capacity	0.00			0.00			610.00			610.00			610.00	
5	SCCT Capacity	0.00			0.00			0.00			0.00			0.00	
6	Coal Capacity	0.00			0.00			400.00			400.00			400.00	
7	PRD	0.00			0.00			0.00			0.00			0.00	
8	Wind 1	0.00			0.00			1,200.00			1,200.00			1,200.00	
9	Wind 2	0.00			0.00			0.00			0.00			0.00	
10															

Figure P-133: Decision Cells for Case Without DR

Finally, the case where DR is fixed at 500MW in all years requires only that the cumulative capacity be set to that value and held across the study. See Figure P-134, below.

PlnCap_0 = 0

	N	R	S	AH	AI	AJ	AP	AQ	AR	AX	AY	AZ	BF	BG	BH
1		Sep-04	Dec-0	Sep-08	Dec-08	Mar-0	Sep-10	Dec-10	Mar-1	Sep-12	Dec-12	Mar-1	Sep-14	Dec-14	Mar-
2	Capacity Data ID														
4	CCCT Capacity	0.00			0.00			610.00			610.00			610.00	
5	SCCT Capacity	0.00			0.00			0.00			0.00			0.00	
6	Coal Capacity	0.00			0.00			400.00			400.00			400.00	
7	PRD	0.00			500.00			500.00			500.00			500.00	
8	Wind 1	0.00			0.00			1,200.00			1,200.00			1,200.00	
9	Wind 2	0.00			0.00			0.00			0.00			0.00	
10															

Figure P-134: Decision Cells for Case With 500MW DR

With these modifications, the model creates the three feasibility spaces described above.

Wind

Two prominent themes for wind generation studies dealt with the assumption of declining capital cost and with the opportunity cost for not pursuing wind.

Non-Decreasing Wind Generation Cost

Chapter 5 of the final plan and Appendix I describe key generation cost assumptions. The plan assumes wind construction costs decline at 1.6 percent per year. To understand the extent to which this declining-cost assumption might be driving the results of the model, a study [**52**] assumed that the wind costs did not decline from today's levels.

As expected, the overall system costs increased dramatically (see Figure P-135), and there was some reduction of wind along the efficient frontier, but wind still appeared in 2013 and develops to its full potential (5,000 MW) by the end of the study. Coal developed in somewhat more plans near the least-risk end of the efficient frontier, but never by more than 400MW. Conservation commanded more of a premium closer to the least-risk end of the efficient frontier.

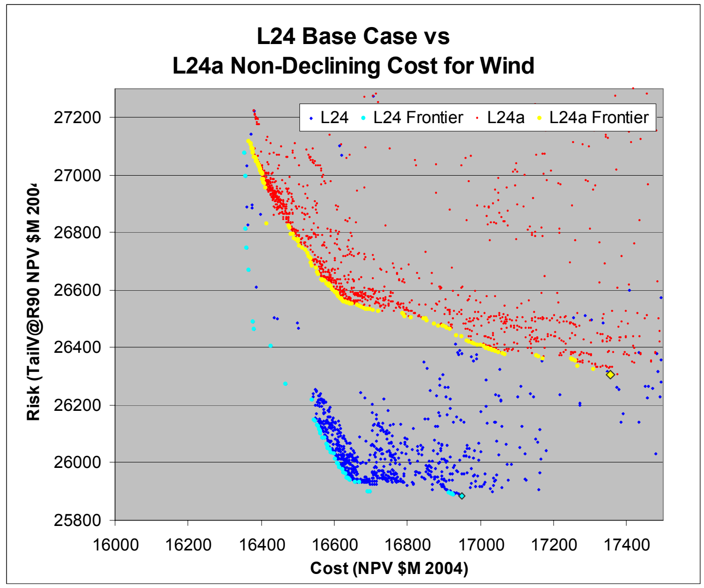

Figure P-135: Feasibility Spaces for Non-Declining Wind Cost Study

The rate of construction cost escalation is a parameter specified in the workbook. The cell {{K509}} of the base case stipulates that the quarterly escalation rate is -0.408 percent. The Council performed this sensitivity study merely by setting this value to zero, as illustrated in Figure P-136. Note that the row containing data labels has a modified format in this figure to make the labels easier to read.

	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
508	Criterion_Set_ID	Planning_Periods	Optional_Construction_Periods	Committed_Construction_Periods	Planning_Costs (RL $M/MWPeriod	Mothball_Costs (RL $M/MWPeriod	Cancellation_Costs (RL $M/MWPeriod	Construction_Costs (RL $M/MWPeriod	CancelThreshold (	Const Cost Escl (.01=1%/period)	ResourceLife (periods	OptionLife (periods	PermitMarketAdds (T/F	PlannedPlanning_Costs	Index
509	Wind	0	2	2	0	0.0007426	0.029704	0.007426	-99999	0.00%	80	80	FALSE	0	4

Figure P-136: Modified Cost Escalation for Wind

The Value of Wind

One study [**53**] examined the opportunity cost of ignoring wind as a capacity expansion option. This study removed wind generation as a candidate for system expansion from the base case.

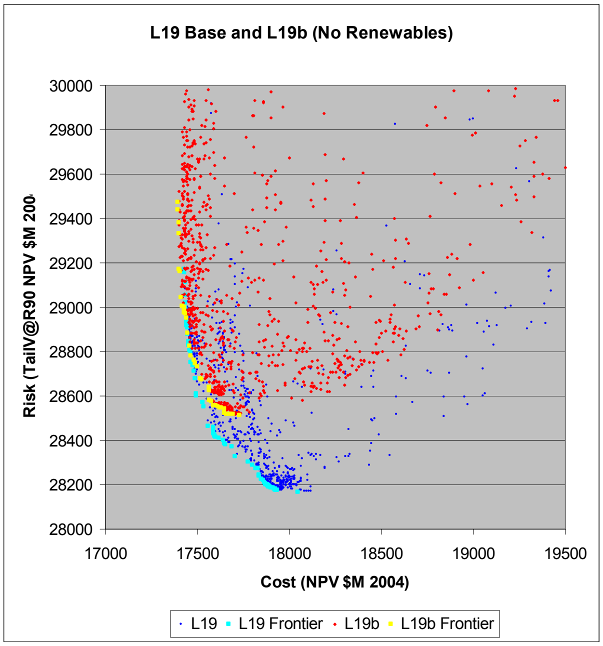

Figure P-137: The Feasibility Space, Removing Wind As An Option

As the section on the value of DR suggests, *value* – in terms of cost reduction – typically *depends on the level of risk the region is willing to assume*. To understand value, therefore, we must consider the efficient frontier. At each level of risk the efficient frontier may shift different amounts to the left, or not at all. Figure P-137 illustrates this principle, with negligible cost shift near the least-cost end of the efficient frontier and significant variation in the least-risk plans.

We should not be surprised to see little or no value at the least-cost end. There is little, if any, wind chosen along this part of the efficient frontier. After all, wind generation is expensive today relative to expected, long-term equilibrium market prices. If the region were content to "ride the market," the right answer would be to build little or no wind – or any thermal resource for that matter. After all, the equilibrium price for wholesale electricity is the same as the fully allocated cost of a CCCT. Why build when you can buy? This is the argument that the "Gas Price" and "Electricity Price" sections of this Appendix explore, and some would claim it is the fundamental assumption that led to the 2000-2001 energy crisis.

As risk mitigation becomes a consideration, however, the value grows. Moving to lower-risk plans, the difference in least-costs plans grows to about $200 million. Beyond the level of risk mitigation that maximizes the difference, however, the value is impossible to determine. Why is the value impossible to determine? Beyond that point, there are no plans *without wind* at any cost that provide the level of risk mitigation that plans *with wind* generation provide!

To create the feasibility space without wind generation, this study eliminated the optimizer's decision cells and constraints pertaining to wind. In the workbook, the values in the decision cells associated with wind are zero across the study. Because the optimizer cannot modify the decision cells, those zero values never change. The situation for the decision variables appears in Figure P-138; the absence of wind capacity constraints is evident in Figure P-139. For more detail about decision cells and how the optimizer modifies them to define a plan, see Appendix L, in particular the section “OptQuest Stochastic Optimization.”

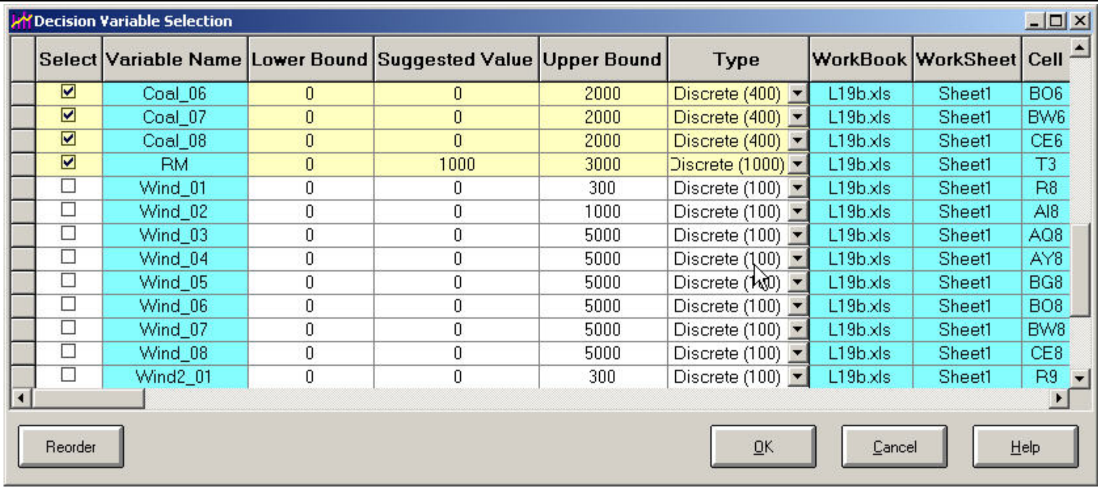

Figure P-138: Optimizer Decision Variable for No Wind Generation

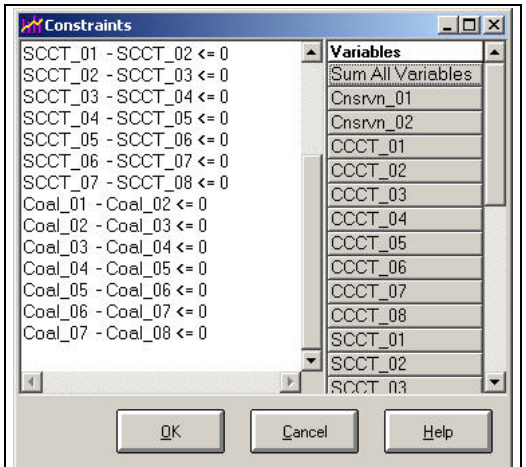

Figure P-139: Optimizer Constraints for No Wind Generation

Two other studies presented in this Appendix bear on the question of wind generation value. The section on CO_2 taxes looks at the issue of how CO_2 tax, green tags, and production tax credits affect the value of wind. As expected, new wind generation was constructed in less quantity and much later along the efficient frontier. Nevertheless, wind did appear in the least-risk plans. Despite its cost, low availability factor, and disadvantage with respect to dispatchable generation, it still provides a hedge against fuel cost excursions and has planning flexibility advantages, like short lead-time and modularity.

The second sensitivity study that bears on the value of wind generation is one that examines the role of planning flexibility for conventional coal-fired generation. In that study, the CO_2 tax, green tags, and production tax credits again are zero, but coal is given a shorter construction cycle. Coal then becomes competitive with wind. This study is the topic of the next section.

Conventional Coal

Conventional coal faired poorly in most studies, entering as a construction option only in the most risk-averse plans and then only in fairly small amounts, typically 400 MW or less. Various studies indicated that the problems with coal were associated with CO_2 taxes, long construction lead times, and to some extent, PTC and green tag programs that make wind more competitive. If this is true, removing these factors should cause coal to appear on the efficient frontier. All regional model studies assumed, after all, that coal had several benefits, primarily stable and low fuel price.

One study [**54**] assumed no carbon tax, no green tag credit or PTC for wind, and a construction cycle for coal that matches that for a CCCT. The construction cycle, after siting and licensing, is two years. Total overnight cost of coal, however, is the same as in the base case. Cash flow is compressed into the shorter construction interval.

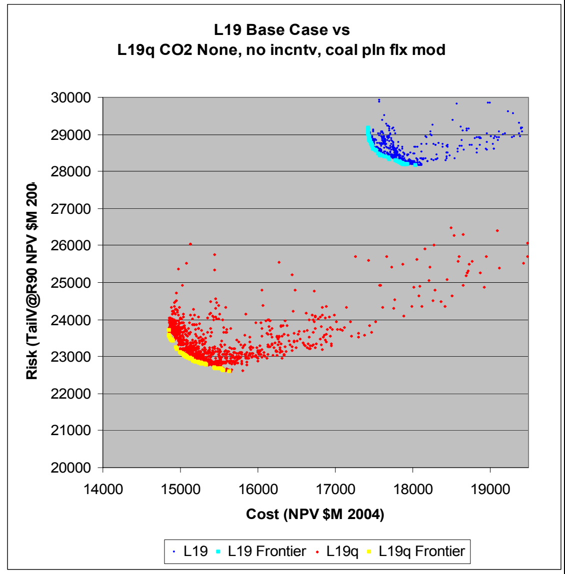

Figure P-140: Removing the Barriers to Coal Development

The study appears to corroborate the view that the perceived disadvantages drive the results. The feasibility space (Figure P-140) is generally less risky and less costly due to the elimination of the CO_2 tax and reduction of the coal plants' construction cycle. Moreover, coal plants appear in almost all of the efficient frontier's plans, being absent in only the risk-indifferent, least-cost plans (upper left hand extreme of the trade-off curve). This stands to reason, as few resources except some inexpensive conservation appear among these plans. At the other extreme of the curve, least-risk plans have substantial amounts of coal, adding up to 400 MW of coal-fired generation by 2012 and up to 2,000 MW by 2015. (In this particular study, coal was constrained at 2,000 MW from 2015 until the end of the study, so it is not possible to determine whether or how much additional coal the model might have added otherwise.) CCCT capacity and conservation develop, too. Coal displaces primarily wind capacity development.

To perform this study, {{rows 74 (CO_2 tax), 81 (PTC), and 83 (green-tag value)}} are hard-wired to zero. The study accelerates the rate at which cost accumulates during construction to achieve the same overnight cost in the shorter construction cycle [**55**]. The resulting values modify the construction cost information in {{row 483}}, as shown in Figure P-141. (Appendix L, section "Parameters Describing Each Technology," provides an interpretation of these parameters.)

	A	B	C	D	E	F	G	H	I
482		Criterion_Set_ID	Planning_Periods	Optional_Construction_Periods	Committed_Construction_Periods	Planning_Costs (RL $M/MW/Period^2)	Mothball_Costs (RL $M/MW/Period^2)	Cancellation_Costs (RL $M/MW/Period^2)	Construction_Costs (RL $M/MW/Period^2)
483		Coal Criter	0	4	4	0	0.000593	0.023708	0.005927
484									
485	Original	Coal Criter	0	5	9	0	0.000339	0.013547	0.003387

Figure P-141: Cost Data Cells in the Workbook

A separate study shows that while eliminating CO_2 tax, PTC, and green tags alone does result in some coal construction, the construction levels are relatively low. (See the sensitivity study "No CO_2 Tax or Incentives for Wind.") Those results, combined with the subject study, suggest that the relative lack of ***planning and construction flexibility*** associated with coal plants is the most significant source of risk and cost.

Larger Sample of Futures

As explained in the first chapter of this Appendix, Monte Carlo simulation provides many advantages for modeling uncertainty. One of the disadvantages, however, is that one must estimate the number of games necessary to guarantee a given level of estimate accuracy. Both the statistics that the regional model uses, mean cost and $TailVaR_{90}$, are averages and therefore have well-understood statistical properties. Because the regional model used 750 futures, the estimate of the mean cost was relatively precise: where the standard deviation of costs associated with a given plan is on the order of $6 billion, the standard deviation of the mean estimate is about $220 million. While the standard deviation of the tail is smaller, however, the sample of the tail has only 75 games, so the precision in the $TailVaR_{90}$ statistic is not much better.

Given the uncertainty associated with these statistical samples, the Council took several steps to assure that the results are representative. For example, staff examined plans that lie off the efficient frontier. The section "Portfolio Model Reports And Utilities" of Appendix L, for example, explains how reports are marked to reveal not only the plans lying on the efficient frontier, but also those lying within $250 million NPV cost and risk from the efficient frontier. In particular, staff studied these plans, searching for patterns or strategies that differed from those on the efficient frontier.

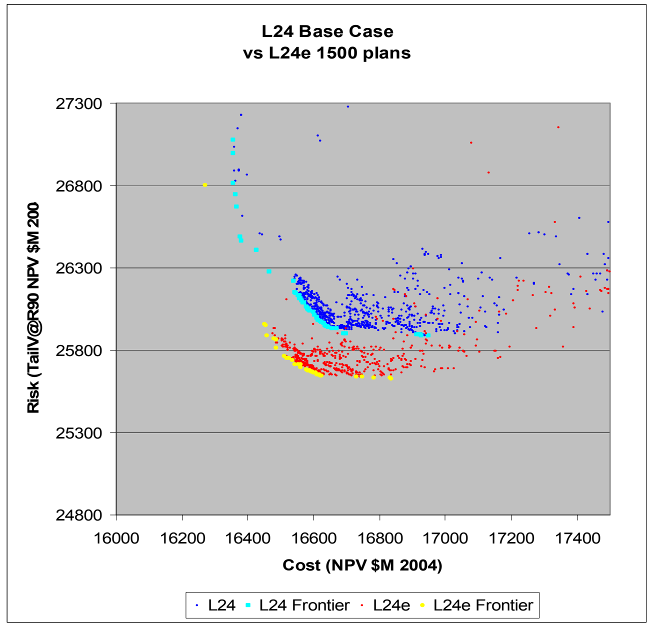

Figure P-142: Increased Number of Futures

Staff also reproduced the final study using 1,500 futures [**56**]. To our surprise, there did appear to be some differences between the two approaches. First, both cost and risk appeared to improve. (See Figure P-142.) The magnitude of the improvement, however, is consistent with sample variation. For example, Figure P-143 shows the average of N random values drawn from a normal distribution with mean 100 and standard deviation of 100 [**57**]. The value on the horizontal axis is N. At around 750, the estimate of the average is off about two percent. If the standard deviation of the costs associated with plans is about $6 billion, two percent corresponds to about $120 million. Figure P-144 demonstrates that this is about the effect on the 192 plans that both the base case and the sensitivity case evaluated [**58**]. More important, perhaps, is that the position of plans relative to the efficient frontier are, by and large, unchanged. In particular, of the 50 plans on or within $250 million of the efficient frontier, the range of costs differences for 49 plans is $99 million to $108 million ($9 million wide) and the range of risk differences is $206 million to $280 million ($74 million wide). Thus, the shifts are very regular. (The one plan in 50 that fell outside these ranges was associated with a high-risk plan.) Figure P-144 identifies the 50 plans lying near the efficient frontier with larger, pink points.

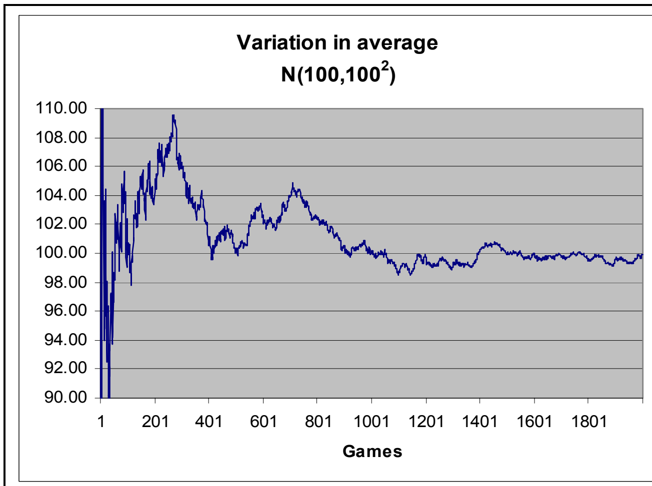

Figure P-143: Variation in Mean

A second difference observed in the sensitivity study was a minor difference in the make-up of plans on the efficient frontier. While coal-fired power plants still appeared in plans near the efficient frontier, none of the plans on the efficient frontier had this resource. More lost-opportunity conservation also appears in the plans on the efficient frontier, and it merits an additional 10-mill/kWh premium. The reason for these differences in plans on the efficient frontier may simply have been that the sensitivity study had fewer plans than did the final base case (827 vs 1010) and the optimizer had not yet found the best strategies. Because doubling the number of futures increased the study time proportionally, however, the sensitivity case had to be ended prematurely.

While the results are somewhat different than expected, the study did not contradict the results suggested by the base case. The plan recommended by the Council would appear very close to, if not on, the efficient frontier near the least-risk end of the efficient frontier.

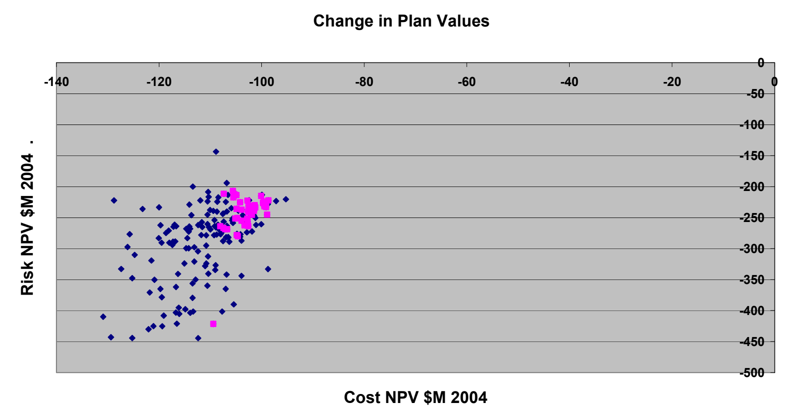

Figure P-144: Consistent Effect on Cost and Risk of Plans

Glossary and Abbreviations

ARMA – autoregressive, moving average technique for time-series analysis
B/C – benefit cost ratio
BMH – Bench Mark Heuristics, consultants
BPA – Bonneville Power Association
CCCT – combined cycle combustion turbine
coherent risk measures – See chapter "Risk Measures"
deciles – data values determining the 10^{th} percent, 20^{th} percent, …, 90^{th} percent, 100^{th} percent level
flat electricity prices – prices averaged across on-peak and off-peak periods, weighted by the respective number of hours in each
FERC – U.S. Federal Energy Regulatory Commission
FOR – forced outage rate
GBM – geometric Brownian motion (see section "Stochastic Process Theory" in the chapter "Uncertainties")
GENESYS – Council's computer model for performing reliability studies
innovation – the value of a random variable
IPP – independent power producer
IRP – a utility's integrated resource plan
IRR – internal rate of return
lambda – the short run marginal cost of a power production system
LOLP – loss of load probability
Mid-C – Mid-Columbia power trading hub
Monte Carlo simulation – see the "Introduction" chapter
myopic cost-effectiveness standard – any cost-effectiveness standard that relies on the assumption of perfect foresight
NOAA – U.S. National Oceanic and Atmospheric Administration
NPV – net present value
PNCA – Pacific Northwest Coordination Agreement
PNUCC – Pacific Northwest Utilities Conference Committee
PRD – price responsive demand
PTC – production tax credit
risk metric – a measure of bad outcomes
RRP – resource-responsive price
SAAC – the Council's System Analysis Advisory Committee
spinner graphs – graphs used to view futures and their effect
TailVaR90 – mean of the 10 percent worst outcomes (see "Risk Measures" chapter)
TLZ – a range of the portfolio model worksheet that performs iterative search
uncertainty – a factor relevant to a decision, over which a decision maker has little control or cannot foresee
V@R or VaR – Value at Risk (see "Risk Measures" chapter)
VBA – Visual Basic for Application
UDF – Excel workbook user-defined function

Figures

Index

References

1 Hoffman 1998, p. 239
2 WOP Forecast History_MJS.xls
3 Regional NonDSI Loads.xls
4 Terry Morlan, Ph.D., Council Staff
5 Regional NonDSI Loads_correction_040623.xls
6 Regional NonDSI Loads_correction_040623.xls, worksheet "Genesys."
7 Total D2.xls, worksheet "Genesys."
8 NG Prices 040903.xls, especially worksheet "Monthlies & Quarterlies –MJS"
9 QuickCheck.xls, worksheet "Sheet6"
10 Terry H. Morlan, Slide 15, Northwest Power Planning Council, Power Committee Briefing, July 16, 2003, NG Update 7_03.ppt
11 MS\Hydro General\HydroGen AddIn\BPAREGU.OUT
12 HydroGen_EastWest.xls or one of its successors.
13 HydroGen AddIn.doc
14 sustained_peak_MJS.xls
15 HydroGen_EastWest.xls or one of its successors.
16 vfuncHydro4x2.xls
17 Statement_of_work_Marty_Howard.doc, for work between January 29, 2003 and July 31, 2003, and follow-on revisions. Key work products are the CD-ROMs entitled "NW Energy Stats – Factors, et. al.," 7/10/03, "Final Data Version," and a revision CD entitled, "Re: Stats," 9/19/03. A good work summary document is "DocumentAnalysis.doc," dated 9/19/03, included on the later CD. These CDs are currently housed in the binder "Olivia, Vol.2" under tab 35.
18 See Dependence on local variables.xls. The change in electricity price is taken as the first-order change in electricity price due to a one-standard deviation change in the electricity price ARMA error term relative to the first-order change in electricity price due the sum of one-standard deviation changes in the local terms.
19 See Notes_L14.doc, notes from conversation, 6/5/05, and Schilmoeller email to Watson, 7/7/04 10:19AM, "Re: Bias in 'textbook' definition of elasticity."
20 See the workbook Averages.xls
21 The values in L28 are from Mo Price_extracted.xls, which in turn stems from the "Mo Price" worksheet of Jeff King's report PLOT R5B10 RvDmd 102104.xls...\Portfolio Work\Olivia\Data Development\Regional Electricity\New Electricity Prices 041028\Mo Price_extracted.xls - 'Comments and Instructions'!B4
22 See the workbook FOR_040502.xls
23 See workbook Aluminum Prices LME 1989-02.xls
24 The computation of the standard deviation is in Aluminum Prices.xls.
25 Terry Morlan, Ph.D., Council Staff, recommendation for L13.
26 Notes from carbon tax experts' meeting: Questions for CO2 Experts.doc
27 Worksheet "PacifiCorp's CO2 Tax" of workbook "CO2_Tax.xls"
28 Worksheet "Wind B1&2 112304 -- Viability" from Production_Tax_Credit for wind.xls
29 Jeff King, email to Michael Schilmoeller, May 31, 2004, Future cost of wind.txt.
30 Jeff King, email to Michael Schilmoeller, August 17, 2004, Value of the production tax credit.txt
31 Jeff King, conversation, 11/22/04. See also Stephano's Green Tag study 10/2004.
32 Coherence of TailVaR90.doc
33 Most of the illustrations in this section are from Comparison of Risk Measures.xls
34 Illustrations are in L27CostVolatility.xls. Other references in Portfolio Analysis Update083094b.ppt
35 Import-Export Analysis.xls
36 Scenarios A, B, C, and D.xls
37 Q:\MS\Plan 5\Portfolio Work\Olivia\Calibration and Verification\Genesys Validations\L13\LOLP_L13_LC\LOLP_L13_LC.zip\Durat.vbi

38 Chart 2 in Comparison of Risk Measures.xls
39 L19g High Gas Price
40 L24f Lower Electricity Price Volatility
41 L19i CO2 None and no incentives
42 L19t CO2 McCain-Lieberman
43 See, for example, the Base Case for L19
44 L19e CO2 50% & tiered, L19f CO2 50% & tiered + NG inc $1.50, L19n CO2 25% & tiered, and L19p CO2 25% & tiered + NG inc $1.50. The tiered approach was adopted in the Final Plan.
45 L28g Value of IPPs and notes.
46 L24i Contracting the IPP out of the region and notes.
47 L28 basecase, L28b Mildly Restricted Conservation and notes, and L28c Severely Restricted Conservation and notes.
48 ConvertingOvernightToPeriodCosts_v06 for DR only.xls
49 L28d No DR and notes.
50 L28j 500MW DR and notes.
51 Graph is developed in DR Study.xls
52 L24a Non-declining costs for wind
53 L19b Renewables=0 and notes
54 L19q No CO2 Tax or Incentives for Wind, Shorter Construction Cycle for Coal
55 ConvertingOvernightToPeriodCosts_v03A.xls
56 L24e 1500 futures and notes
57 L24e 1500 futures\Effect of Skew on average.xls
58 L24e 1500 futures\Looking for identical Plans in_Analysis of Optimization Run_L24e.xls

c:\backups\appendix risk and uncertainty\appp_060928.doc (Michael Schilmoeller)