

110519 SAAC minutes 110926.docx

Monday, September 26, 2011 version

May 19, 2011 SAAC meeting minutes

The agenda and list of attendees is attached.

Michael opened the meeting with review of the minutes from the last meeting. Several participants had contributed detailed corrections and clarifications. Michael asked for a discussion of the minutes and invited participants to share their views. The minutes of the February SAAC meeting, with the corrections noted, were unanimously adopted.

Review of Prior Material

Michael has started each meeting with a 90-minute review of previous meeting presentation material. Several participants who had not had the opportunity to participate in earlier meetings asked questions about the nature of plans along the efficient frontier and the means by which the model selected these plans. They questioned how the optimizer became involved at the end of 750 games or futures. They wondered how the model would go about choosing a different plan if optimization were taking place at that point.

The confusion arises from the use of the term "optimizer." Perhaps more descriptive terms would be "selector" or "tester." The model does not have enough information from a single plan under 750 futures to perform a global optimization. In fact, only linear models are capable of identifying a global minimum cost or risk, and the RPM model and behavior are very likely nonlinear and in any case cannot be shown to conform to the requirements of a linear model. Consequently, all that the optimizer or tester is doing at the end of each simulation over 750 futures is gathering two pieces in of information from the resulting distribution, the average cost and the risk. The tester then tries a different plan, selecting the size and timing for siting and licensing (or earliest construction) of resources and of conservation premiums that differs from those of the first selection. This process is largely random for the first 800 to 1000 resource portfolios. After that, the optimizer begins to learn enough information about how the 76 decision cells that define the plan can be modified to improve cost and risk. A total of about 3500 plans typically provide a high degree of assurance that an optimum plan has been identified.

Michael showed the members the yellow cells at the top and of the RPM workbook that define a plan. The plan consists of the earliest construction start dates for each power plant, the premium for conservation, and schedules for development of demand response. The optimizer tests different plans by changing the values of particular cells in the RPM workbook model. These cells, which have a distinct yellow color in the workbook, are called "decision cells," a Crystal Ball™ term. (See Figure 1 below.) The choices the tester can try are subject to any constraints

that the modeler provides. (See Figure 2.) The constraints limit the amount of searching the tester must perform. The modeler, however, must take care to verify that the constraints are not binding on the solution. If they were binding, this would indicate that the constraints should be moved to provide the optimizer an opportunity to find a better solution.

"Perhaps the plan should be called something else, like a strategy." One member felt that the word "strategy" better described a plan that provides optionality for decision makers to modify their behavior depending on the circumstances in which they find themselves. The terms "plan" and even "resource portfolio" suggests that a much more constrained schedule of construction to many utility planners.

PlnCap_0		= 0													
N		R	S	AH	AI	AJ	AP	AQ	AR	AX	AY	AZ	BF	BG	BH
		Sep-04	Dec-04	Sep-08	Dec-08	Mar-09	Sep-10	Dec-10	Mar-11	Sep-12	Dec-12	Mar-13	Sep-14	Dec-14	Mar-15
1	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	Wind1	0.00			0.00			1,200.00			1,200.00				1,200.00
9	Wind2	0.00			0.00			0.00			0.00				0.00
10															

Figure 1: Decision cells in the RPM

How does the premium over wholesale market price enter into the model's plan? Michael showed the participants two cells that the optimizer changes to test different levels of premium for lost opportunity and discretionary conservation.

Participants appeared to appreciate seeing the specific cells that define a plan. This made the idea of a plan much more concrete.

"Do the futures change from plan to plan?" Michael showed the participants the list of uncertainties that make up a future and explained that each plan sees exactly the same set of futures, with the single exception of the future for electricity price. Electricity price must be different for different plans, because the amount of energy generation at a given electricity price will be greater or less with different plans. Energy requirements must be balanced, however, with generation, imports, and exports. This would not be possible if electricity price were fixed irrespective of our resource plan.

"It would helpful to again see a simple example of a plan, now that we understand how they are represented in the worksheet. It also would also be helpful to see how the uncertainties that comprise futures are represented in the worksheet."

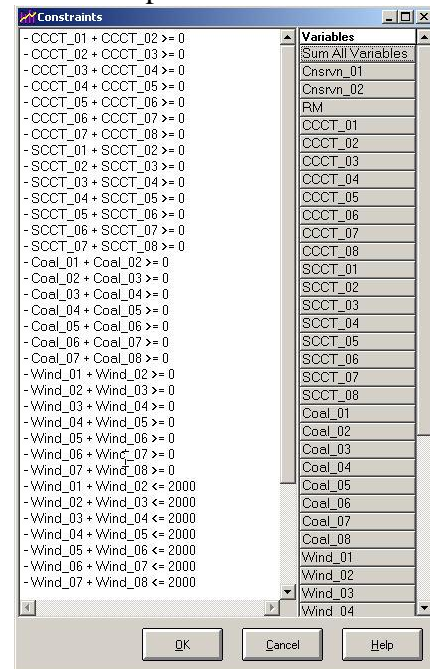


Figure 2: Decision cell constraints

One participant mentioned that his utility uses an optimizer to find the least-cost plan immediately after performing a

single simulation across uncertainties. There was some confusion over why the process required so many simulations. Schilmoeller explained that with nonlinear problems, one is never guaranteed that they have found a global minimum. That can only be done for linear programming problems, and it only works if one has perfect foresight about the future.

Michael discussed the representation of hourly data in the RPM. These data are captured with statistical distributions. Special techniques use the statistical distribution to calculate costs and quantities, rather than hourly simulations.

There was discussion of the correlation of natural gas price, electricity prices, hydrogeneration, and loads. Michael explained that the production of the electricity price futures begins with a set of futures much like those of natural gas prices. That is, futures are independent from each other and they are independent from those of other uncertainties. These prices are then correlated with natural gas price, hydrogeneration, and loads.

The resulting price is much more weakly correlated with these factors than some would expect. Michael said that in their early study of these influences less than 40% of the variation in electricity prices could be explained by the combined effect of these factors. A significant consideration is use of local electricity prices, local gas prices, and local loads. One significant cause for this weak relationship is the rest of the Western interconnect. Load variation and economic activity outside of the region have an effect on Western prices and, consequently, regional prices in many hours. Other influences are the lack of perfect information about the rest of the system, the risk-averse behavior of market participants, and operational details beyond the scope of our knowledge and certainly beyond that of production cost models.

The resulting electricity prices may not provide a feasible solution to the energy balance problem. For this reason, they are subject to further adjustment, an iterative search process, to bring about energy balance. This must be done to for both on- and off-peak energy, respecting at all times the constraint that on-peak prices must be higher than off-peak prices.

The issue of decision criteria drew several questions from participants. The participants understood that the model makes internal forecasts a commodity prices and evaluates the economic feasibility of projects on a forecast basis. They also understood that there is a forecast of regional energy adequacy, which is another reason to complete construction of power plants. There was confusion, however, about how these processes interacted with the optimizer. Michael stated that the optimizer tests the maximum amount of resource that can be constructed at a given point in time. These choices, this plan, are fixed in a given simulation over all futures. The decision rules only decide whether the plan's selected capacity for a resource in a particular period "should be" constructed, given the circumstances of the future.

One participant asked if there was a way of controlling the correlation between gas and electricity prices. Schilmoeller indicated that that is certainly possible in the model. He questioned whether it would have much effect given the rather weak correlation that already exists. Some participants were interested in pursuing this further.

Michael returned to the review. Unit aggregation using cluster analysis was another way to increase the computational efficiency of the model. Michael also spoke about the performance and precision of the model. He reminded the members of the reason why the resolution of the model is so much greater than the accuracy of the TailVaR90 estimate. Over 70 of the 75 futures in the TailVaR90 estimate show up among the worst 10 percent futures in all of the plans along the efficient frontier. Consequently, much of the variation could be described as simple sensitivity analysis.

The choice of Excel as a platform stems from a desire to make the model as transparent and accessible as possible. However, Excel VBA also permits Olivia, the meta-model that wrote the original RPM, to modify *the code* and logic of the RPM. This means that the model has only the logic that needs to perform the tasks desired. This makes the RPM as efficient as possible. Michael concluded the review and began the first presentation of new material for this meeting.

The Nature of Risky Futures

The first presentation of new material returned to prior meetings' discussion of the spinner graph's future 750¹. It features a combined cycle combustion turbine constructed midway through the study but used very little. There have been several questions raised and discussions about this particular future.

One of the questions about future 750 is, why is the cost of that future low relative to the average cost for this plan across all futures? Wasn't it in fact the case that overbuilding should increase costs?

Michael presented comparisons of the electricity price, natural gas price, carbon penalty, power cost per kilowatt hour, and load forecast between this future and the average. One of the things that's evident is that the load forecast in this future is lower than average, among other differences. Consequently it is more useful to compare future 750 with combined cycle combustion turbines against *the same future* without those turbines. It also useful to look at how the total cost affects the ratepayers in a low load future.

When we compare the costs of the future with and without the combined cycle combustion turbines, we see the future with the combined cycle combustion turbine is about \$6 billion, or 8 percent, more expensive. A comparison of the cost per kilowatt hour shows that without the combined cycle combustion turbines, the unit cost would be lower than average. With the combined cycle turbines, unit cost *is higher than* the average across the 750 futures.

Michael also pointed out that sometimes the model needs to get this wrong. The purpose of the exercise is to value the option of constructing the turbines and to determine a reasonable amount of capacity to pursue. The answer to that question turns on how expensive those turbines may be if they are eventually found not to be necessary for the region. This information is as important as the cost if we do not build the turbines and wind up needing them.

¹ The spinner graph is available from the Council's website:
http://www.nwcouncil.org/dropbox/Olivia_and_Portfolio_Model/Spinner_Graphs_6th_Pwr_Pln/Spinner_091220_2157_L813_2990_LR.zip

Finally, the primary reason for siting the nine 415 MW CCCTs in the Sixth Power Plan resource portfolio is carbon penalty risk. With a significant carbon penalty, the region stands to lose 5000 MWA of coal plants. The CCCTs would replace almost 70 percent of this.

A participant asked if there is a limit on how many plants can be built. Michael explained that there was, and that part of the modeling process is providing tunnel constraints for siting and licensing, that is, the least and maximum number of plans the can be built in each period.

Does the staff attempt to model constraints on natural gas delivery? If more gas-fired generation is needed, for example, do we raise the price of natural gas? No, we do not attempt to model supply and demand for natural gas, explained Michael.

There was significant discussion around discretionary wind power plant construction and the modeling of regional portfolio standards (RPS). We could have modeled uncertainty around the success of RPS statutes, that is, the likelihood that they will remain in force and produces the level of renewables currently expected. This uncertainty was not represented in the Sixth Power Plan. The Council elected to model RPS statutes and renewable power construction according to current law, in order to simplify the studies.

Sources of Regional Power Generation Cost Risk

Michael introduced with the second topic of the morning, the source of risk according to the RPM. He explained that Appendix J in the Sixth Power Plan includes a regression study identifying uncertainties that were driving the cost and risk of the model. System cost is a dependent variable, and the selected independent variables were electricity price, CO2 penalty, natural gas price, "position", and "market." Observations were taken both on and off peak for each period across the 750 futures. The "position" was of measure of the energy difference between load and those resources largely insensitive to electricity price. These sources include hydrogeneration, conservation, and RPS wind, must run resources, and firm contracts. Note that position will be strongly associated with load level. The "market" variable was an interaction term between the position and electricity prices.

Appendix J's regression analysis indicated that about 95% of the variation in on-peak system cost could be explained by these factors. The off-peak cost model is somewhat weaker, with an R^2 of 89%. Very little over-specification occurs in the model. The electricity price, position, and carbon penalty appear to be among the most significant factors. Natural gas price and position were relatively weaker, but still strong (p-value < 0.005).

Michael said that after finishing the appendix, he had searched for a more intuitive description of risk factors. He hit on the idea of sorting the futures in the spinner graph according to the future's NPV cost. He presented this first 75 (most expensive) futures and invited participants to examine variations within each future and among the futures. One of the things that Michael observed was the prevalence of high load futures. Other factors were high carbon penalty, high electricity prices, and extended periods of poor water condition.

A Per kWh Cost Measure

Michael asked how the group felt about selecting the least-risk strategy exclusively from futures with high loads. Was it conceivable that a low- or medium-load future might have bad consequences for regional ratepayers? What about the early 1980s, when ratepayers began paying for recently completed capital intensive resources and loads fell due to their cost and due to economic down-turn? Should we exclude this outcome from consideration if fuel prices were low enough that *total cost* did not reach excessive levels?

The issue prompted Michael to look at an alternative way to characterize costs. The problem in the early 1980s was high cost per kWh, which resulted in doubling of power rates for BPA customers. The high rate was in large part due to low loads, that is, a small number of kWh across which costs were spread. We need to avoid a rate measure per se, however, which ignores the service provided by conservation measures and reflects only utility revenues from energy charges.

If conservation is to be placed on par with other supply side resources, the alternative measure must have a unit measure of total service provided to customers, irrespective of source. In fact, the RPM has calculated such a value since the Fifth Power Plan. The denominator uses a “frozen efficiency load” that contains no new conservation² subsequent to the study’s beginning date. If time value of money is important, however, these values would need to be present-valued.

Several of the SAAC members found this proposal flawed or incorrect. A participant said that we need a cost metric anyway. He asked whether this proposed metric would replace the cost metric or supplement it. Michael replied that it would not replace the cost metric. “If the TailVaR90 measure is reflecting higher load situation, isn't that really a higher risk?” asked a member. Participants agreed to make a presentation outlining their concerns with the unit cost metric.

Another participant liked the new metric because it can correct for overdeveloping conservation. It will correct for overbuilding of all resources, agreed Michael.

Uses and Abuses of the Efficient Frontier

The first topic of the afternoon was "The Uses and Abuses of the Efficient Frontier". The efficient frontier is a technique for selecting from plans with multiple outcome attributes. It avoids, however, the problems of weighting the attributes that utilities very often run into. Weights are often assigned after the utility has performed its study. Some outside spectators will interpret this practice as disingenuous. It also puts the utility in the difficult position of defending trade-offs among very dissimilar attributes.

The efficient frontier avoids these difficulties by making comparisons only among plans that are better or worse in every aspect or attribute of the decision. While this is an efficient means of

² This is not true, strictly speaking. The frozen efficiency loads include new conservation that is the result of legislation or other causes known to guarantee subsequent conservation. “New” conservation here should be interpreted as conservation that we have not yet evaluated.

selecting from among a large number of plans, as in the case of Council studies, it probably is less useful when a utility is looking at a small number of plans. In the right circumstances, however, it can be an effective communication tool.

Stefan Brown, formerly of the Oregon PUC currently with Portland General Electric, once shared the truism with Michael that, “The efficient frontier does not tell us what to do; it tells us what *not* to do.” Michael finds this a concise way to express the limitations of this method.

Michael went on to discuss how the efficient frontier can be useful, despite it not telling us how to pick a particular plan. We can look at other risk measures to help us further localize sets of plans with strategies that are advantageous. It is important to look at which strategies are successful, in the sense that they lie close to the efficient frontier, and which strategies are risky and live further from the efficient frontier. Another question is how strategies along the efficient frontier change and how they are similar.

There are other considerations. How do details within particular futures differ? Are some plans more or less acceptable to other institutions? Do we really have to make a choice? What cost elements can you control?

Michael explained that the question of whether or not we have to make a choice has been pivotal in the last two Council Plans. In both of those plans, no money had to be spent during the Action Plan time period on conventional supply-side resources. This meant that we could adopt the least risk plan without incurring any real cost. It also meant that we preserved the option of later adopting any other plan along the efficient frontier because they typically require action later than the least-risk plan. The only resource affected by the choice of the least-risk plan in the early years is lost-opportunity conservation. Discretionary conservation is effectively controlled by our conservation “ramp-rate” assumptions.

The issue of identifying controllable costs requires further explanation. We return to it shortly in the context of misuse of the efficient frontier.

Errors in the Use of the Efficient Frontier

Michael illustrated some of the mistakes that are made with the efficient frontier. These include using the geometry of the frontier to select a plan, confusing controllable and uncontrollable costs, and viewing average cost on the frontier in the same way that we would view average cost or return for a financial portfolio.

The geometry of the efficient frontier can be completely meaningless, as Michael illustrated with an efficient frontier illustrating the effectiveness and side effects of hypothetical vaccination regimes. By using an example other than the efficient frontier presented in Power Plans, it is more apparent that meaningful or “natural” trade-offs are normally not available.

There is a tendency to use the differences in average costs of plans on the frontier in the same way they use a premium for insurance. When we buy insurance, controllable costs (premiums) are traded off against the expected financial risks faced without the insurance. We typically do not or cannot foresee all the cost consequences – good and bad – of buying insurance. We also

typically do not include the other obligations that such a purchase would involve. This is not to say we should ignore them if we could anticipate them. We merely distinguish the average costs on the Sixth Power Plan's efficient frontier – which include all consequences as well as the cost of the “premium” of siting and licensing – from the premium cost of insurance.

Michael presented an analysis that compares the rates for various kinds of insurance with the cost of siting and licensing power plants in the Sixth Power Plan. He showed that the relative size of the premium to the expected risk was on the order of 0.5 to 1.5 percent, the same as for commercial insurance.

Finally, when we are thinking about a single outcome from among the many very dissimilar outcomes, an average is not very helpful. It may in fact be meaningless. Probabilities are a much more useful way of thinking about these decisions. The average is helpful only in that it gives us a sense of where the most probable outcomes may lie. The efficient frontier can therefore be thought of as providing plans with distinct preference for likely outcomes versus best (least-risk) outcomes. A decision maker that favors a least-cost plan really emphasizes the most likely outcome. The decision maker who chooses plans from the least-risk end of the efficient frontier pays more attention to the worst-case outcomes, and values a plan that mitigates these.

The divergence of values from the Council's forecast in the RPM futures is a key reason why averages are less useful in describing the costs and plans. The average return of a portfolio is useful, because we strongly believe the stock returns are mean-reverting – at least in the short term. The average hydrogeneration on the region's rivers is a useful statistic largely for the same reason. Commodity prices and load futures in the RPM, however, do not generally revert to any one forecast over the 20-year study period. This is the value of a scenario analysis.

The Conservation Premium

The last presentation of the day was on the cost-effectiveness premium for conservation. Michael described how conservation supply curves are constructed, the sources of the cost-effectiveness premium over wholesale market price, and findings from the Fifth and Sixth Power Plans. Michael described the approach that he used to identify sources of value associated with premiums for conservation. Council studies identified a reduction in both cost and risk with the increased premium, up to a point, for conservation. To attempt to monetize the savings, Michael found plans on separate efficient frontiers – one with premiums, and one without any premium – that had the same risk level. He also tried to find two plans that differed as little from one another as possible in order to simplify the analysis. The plan that he found had SCCTs which were deferred by several years. There was also a reduction in market prices for electricity that appeared to reduce cost disproportionately to the value of the conservation alone.

Michael described the reasons why the model would prefer conservation over SCCTs for risk mitigation. The selection initially seemed counterintuitive. Low capital cost resources are usually considered preferred for risk mitigation. This only makes sense, however, if the expected capacity factor for the power plant is low.

In fact, in high-risk futures where the resources turn out to be valuable, high heat rate SCCTs will run very hard often because of requirements for energy and high market prices. In these futures, such resources operating at very high capacity factors are quite expensive.

Alternatively, the futures that are not high-risk futures may in fact have very low wholesale market prices. This is particularly true of low-risk resource portfolios. Low-risk resource portfolios typically have more resources, and these additional resources help to suppress wholesale power price excursions. This means SCCT have less opportunity to recover a portion of their capital costs. The price of electricity has to get up to the dispatch price of the unit before a SCCT begins to have value. Conservation, because it is a zero variable cost resource, has a value at every electricity price.

In short, in a low-risk plan, SCCTs are suboptimal in the two situations we are most interested. They are suboptimal in the high-risk futures, because they are exercised too many hours; they are suboptimal in many of the remaining futures because they are exercised too few hours to recover their fixed costs.

This also provides another example of how we can get into trouble with averages. Just as there is no Roulette die with 3.5 dots (the average of 1 through 6), there may be no future where the SCCT runs at its ideal, low capacity factor. If the SCCT is added strictly for mitigation of unforeseen risks, it is likely it will be run hard or – more likely – never run at all. Such is the case with insurance.

Michael summarized the sources of value for conservation. These can be classified under the headings of capacity deferral, protection from fuel and electricity price excursions, short-term price reduction, purchases at below average prices due to a "dollar cost averaging" effect, and the opportunity to develop and resell conservation energy. Electricity price excursions, moreover, are often associated with carbon penalties.

Finally, Michael presented some findings from the Sixth Power Plan. These show the amount of conservation developed when specific effects are controlled. These effects include carbon penalty, stochastic volatility electricity and fuel prices, and premium.

One participant asked why there is a premium for conservation and is not a premium for other resources. Michael responded that, in fact there is a premium for other resources, but the model does not value at the same way. Other resources, such as wind generation, are typically developed at costs that cannot be supported by market prices in most futures. This cost above market can be interpreted as a premium. Because the conservation and power generation resources are developed by very distinct means, the same approaches cannot be used for both conservation and conventional supply-side resources. Conservation does not have the long lead time and the large, uneven commitments of capital that conventional resources have.

A member felt that the approach to the valuing conservation purchased at above market prices may not be realistic because some of that value was ascribed to deferral of SCCTs. He wondered whether deferral is realistic when those turbines might be needed for integrating wind generation. Michael pointed out that wind generation already has costs included that are intended to reflect

resources necessary for integration and balancing. He suggested that assuming additional SCCTs are necessary for this task would effectively be double counting the cost. These SCCT should therefore be viewed as constructed for non-wind integration purposes.

The meeting concluded at 3:30 p.m.

Agenda
May 19, 2011

- Discussion and adoption of the February 2 meeting minutes
- Review of concepts
- The nature of the risky futures
- Primary sources of risk and a new risk metric
- Uses and abuses of the efficient frontier
- The risk premium for conservation

Attendance

Online

From: Black Burrell, Jo-Ann
Sent: Friday, May 20, 2011 3:06 PM
To: Schilmoeller, Michael
Subject: SAAC mtg. phone list

Jo-Ann Black-Burrell	jblackburrell@nwcouncil.org	8:52 AM - 4:44 PM
Rick Sterling	rick.sterling@puc.idaho.gov	8:52 AM - 3:40 PM
cathy carruthers	ccarruthers@cityoftacoma.org	8:57 AM - 4:04 PM
Steve Weiss	SDWeiss@bpa.gov	9:00 AM - 3:01 PM
Lauren Gage	lsmgage@bpa.gov	9:00 AM - 4:44 PM
Catherine Bailey	catherine.bailey@pgn.com	9:02 AM - 3:11 PM
Brendan ODonnell	bodonnell@rmi.org	9:02 AM - 12:04 PM 1:06 PM - 4:30 PM
clint kalich	clint.kalich@avistacorp.com	9:12 AM - 12:50 PM 12:59 PM - 3:39 PM
Shirley Lindstrom	slindstrom@nwcouncil.org	9:13 AM - 4:44 PM
Dave LeVee	dave@pwrcast.com	9:13 AM - 3:40 PM
Howard Schwartz	howard.schwartz@commerce.wa.gov	9:22 AM - 3:42 PM

Sign-in (next page)

