
2012 Integrated Resource Plan



Appendix

5A

Risk Framework and Tools

Table of Contents

1	Introduction	1
2	Eliciting Subjective Probability Estimates	2
3	Creation of Probability Distributions by Combining Estimates	4
4	Portfolio Effects When Aggregating Estimates	5
5	Deriving Discrete Distributions from Continuous Ones.....	7
6	Combining Scenarios using Probability Trees	9
7	Combining Scenarios and Decisions Using Decision Trees	13
8	Showing Variables Used for IRP Portfolio Modelling.....	14
9	Conclusions.....	15
9.1	Reference Section.....	16

List of Figures

Figure 1	Examples of Tail and Cut-off Values	8
Figure 2	Simple Probability Tree	9
Figure 3	Combining Load Growth and DSM Uncertainty.....	10
Figure 4	Net Demand – Three Point and Five Point Distributions	12
Figure 5	Decision Tree Example	14
Figure 6	Bubble Diagram Showing Modelling Variables of Interest.....	15

1 Introduction

Long-term planning in the energy sector is an inherently uncertain exercise.

BC Hydro has been working to incorporate risk and uncertainty into its long-term resource planning and analysis through the development and application of a Risk Framework. BC Hydro's IRP Risk Framework describes the method used in this IRP to characterize multiple, uncertain impacts in a way that helps compare resource options and that leads to a balanced, transparent decisions.

As explained in Chapters 5 and 6 of the IRP, the Risk Framework consists of a number of key elements, including:

- **Multiple objectives for resource planning** are addressed in section 5.2.1 with respect to how they were identified, defined, measured and used.
- **Capturing uncertainty** is a significant component of the Risk Framework. The key uncertainties considered are described in section 5.2.2. The methods for capturing these uncertainties in modelling and comparing options are discussed in section 5.2.3.
- **Portfolio analysis** is used to combine the above information and test how the various resource options address the key IRP questions. The technical approach, assumptions and models used are outlined in section 5.3.
- **Comparing resource options**, based on the above quantitative aspects, and the qualitative aspects not captured through modelling, enables conclusions to be drawn that can be translated into actions (i.e., a Base Resource Plan). This is addressed in Chapter 6.
- **Identifying residual risks and key warning signs** and providing strategies (contingency resource plans), trigger points, and off-ramps to avoid or mitigate these risks is discussed in section 6.10 and section 9.3.

This Appendix focuses on the quantitative aspect of the Risk Framework, exploring in more detail the ways in which this IRP has attempted to address risk and uncertainty and bring difficult to quantify elements into a rigorous and consistent analytical structure.

The following sections discuss six tools that were used in the risk and uncertainty analysis:

- Eliciting subjective probability estimates;
- Creating probability distributions by combining estimates;
- Accounting for portfolio effects when aggregating estimates;
- Deriving discrete scenarios from continuous ones;
- Combining scenarios into probability trees; and
- Combining scenarios and decisions using decision trees.

Each section will offer detailed explanation of the approach being discussed and will reference where this has been used in the IRP. References to academic literature, limitations of these approaches, and best practices in this area will be made in each section where pertinent. The Appendix concludes with a brief discussion of tools used to communicate the risk and uncertainty analysis within the IRP and also some overall conclusions.

2 Eliciting Subjective Probability Estimates

In cases where it was difficult to quantify the uncertainty of certain long-term planning variables, subjective judgement was used to derive probability estimates. The use of subjective judgement to put probability estimates to discrete scenarios or to generate a probability distribution around a forecasted point estimate is a key element to the IRP's Risk Framework. However, such judgements are not done without difficulty. Subjective probability judgements are hard to do and are subject to well documented perils; experts are often 'wrong' in the sense that they cannot

perfectly forecast future uncertainty and they tend to be biased, and tend to be overconfident in the estimates that they make.

Researchers in this field of decision science have developed protocols to address these issues and reduce these errors; assessments used in the IRP's Risk Framework were carried out with guidance from these advances. However, in the end it needs to be recognized that subjective probability assessment is an inexact science and the outcomes should be taken as rough indicators of relative likelihood and not highly precise measures.

In general, the approach for assigning probabilities to scenarios or to ranges of possible outcomes in this IRP followed the steps listed below.

- Gather subject matter experts;
- Decompose problem into a manageable set of key drivers of uncertainty;
- Pull out the key events underlying these drivers of uncertainties;
- Use this information to rank the relative likelihood of these events (from most likely to least likely);
- With ranking as a starting point, get an 'order of magnitude' feel and qualitative likelihoods (e.g. "very probable", "almost impossible"); and then
- Use the structure of the problem (e.g. probabilities must sum to 100 per cent) to find probabilities that the experts feel match the descriptions above.
- Review results with experts to confirm results. Revise if needed.

The method laid out above is consistent with the "textbook" approach to expert elicitation (Clemen (1996)). In their seminal work on risk and uncertainty in decision analysis, Granger and Henrion (1992) emphasize that the process of probability elicitation from experts is challenging, but it is the "only game in town" when trying to incorporate uncertain estimates into a framework of analysis. By sending out some of the key judgement tasks in advance, collecting them, and then using that as a

basis for discussion, this approach resembles the modified Delphi approach for reaching consensus among experts on parameter estimates (Burgman, 2005).

Where was this used in the IRP? – Probability assessment with subject matter experts was used in several key areas in the IRP, including:

- DSM energy savings – from rates, codes and standards, and programs;
- DSM capacity factors;
- Market Scenarios; and
- IPP attrition.

3 Creation of Probability Distributions by Combining Estimates

As noted in the previous section, it is often useful to decompose an assessment task into several sub-tasks where the expert is more comfortable about estimating probabilities. Once this is done, then these assessments can be recombined into a probability assessment for the overall task at hand. If this is done over a large number of sub-tasks, then this large number of random variables will roll up into a continuous, bell-shaped distribution.

For this IRP, BC Hydro used an Excel-based software package from Palisade Corporation called “@Risk” to perform these functions. In general, adding together random variables to create an overall distribution of the resultant sum can be done through a Monte Carlo distribution. If each probability sub-task is a distribution of possible outcomes, then a Monte Carlo simulation would take a random draw from each sub-task, add them together to get a sum, and store this sum. When done a number of times, the stored results form a distribution of possible outcomes for the sum of the random subtasks. From this distribution, different statistics of interest can be calculated such as central tendency (mean, median, mode), dispersion (standard deviation) or downside risk (tenth percentile, fifth percentile, etc.).

Where was this used in this IRP? Monte Carlo simulations were used to estimate the dispersion of outcomes around a sum of uncertain variables in a number of instances for this IRP, including:

- DSM savings,
- Load forecast uncertainty,
- IPP attrition; and
- Effective load carrying capability (**ELCC**) values for wind.

4 Portfolio Effects When Aggregating Estimates

When examining the spread of uncertainty around a sum of uncertain variables, some care must be taken. The spread of the uncertainty around the mean can be affected significantly by the relationship among the uncertain variables being summed.

Formally, given two variables that are estimated with uncertainty (A, B), the variance of their sum can be found in the following way:

$$\text{Variance}(A+B) = \text{Variance}(A) + \text{Variance}(B) + 2 * \text{Covariance}(A, B),$$

where the last term refers to the relationship between these two variables. If these two variables are independent, then their covariance is zero and the spread of uncertainty of their sum is equal to the sum of each of their individual variances. However, if a relationship exists among these variables, then ignoring the relationship can lead to misleading results about the spread of uncertainty.

As an example, if A tends to be high when B is high (i.e., they have a positive covariance), then seeing extreme results where both these variables are high will tend to be more likely than if they were independent. Ignoring this interrelationship will give a total distribution that is too narrow in range, and will under-represent more extreme outcomes.

This result can be thought of as a “portfolio effect” – the spread of uncertainty around a sum of uncertain variables depends on the type of relationships among these variables. For example, if the total amount of energy savings is the sum of many individual projects, then the estimated spread of uncertainty of the estimated energy savings will be influenced by whether or not these projects’ outcomes are independent. If these projects truly are uncertain but independent, then seeing a total sum that is very low will be unlikely; while one or two projects may not work out, it will be highly unlikely that all the projects will perform worse than expected (and some may even perform better than expected). This “portfolio effect” means that the spread of uncertainty for an aggregated sum will be narrow when the individual components of that sum are not positively correlated.

However, the portfolio effect may be an artefact. If there is a common driver that forces individual projects’ outcomes to move together, then the total spread of uncertainty could be large since projects would tend to underperform or over perform together. Ignoring these interrelationships will lead to an underestimate of the total spread of uncertainty.

While ignoring the interrelationships among uncertain variables can lead to misleading results, estimating these uncertain relationships is not easy, particularly when it is done through the elicitation of these probabilities from experts. Morgan and Henrion (1990) note that humans are notoriously poor at understanding the concept of correlations, spotting them in real life, and estimating them. While the use of structured conversations, visual aids, and a knowledge of these weaknesses can help, subjective judgements are not precise.

The IRP’s Risk Framework relied on experts’ judgements of correlations amongst variables. Visual examples of levels of correlated data, examples of correlations drawn from everyday life, and explicit discussion of these judgements were used to try to improve the quality of these assessments.

Where was this used in the IRP? – Correlations amongst uncertain variables were used in estimating DSM savings.

5 Deriving Discrete Distributions from Continuous Ones

The continuous distribution of outcomes around an expected value is an excellent visual guide to uncertainty in estimates. However, the portfolio modelling in the IRP required that parameters be assigned specific values so that they could be used as inputs into portfolio modelling.¹ This required that the specific outcomes be selected out of the whole distribution of possible outcomes. In order to be useful in the quantitative part of the Risk Framework, these values also need a likelihood attached to them. This section explains how this was achieved.

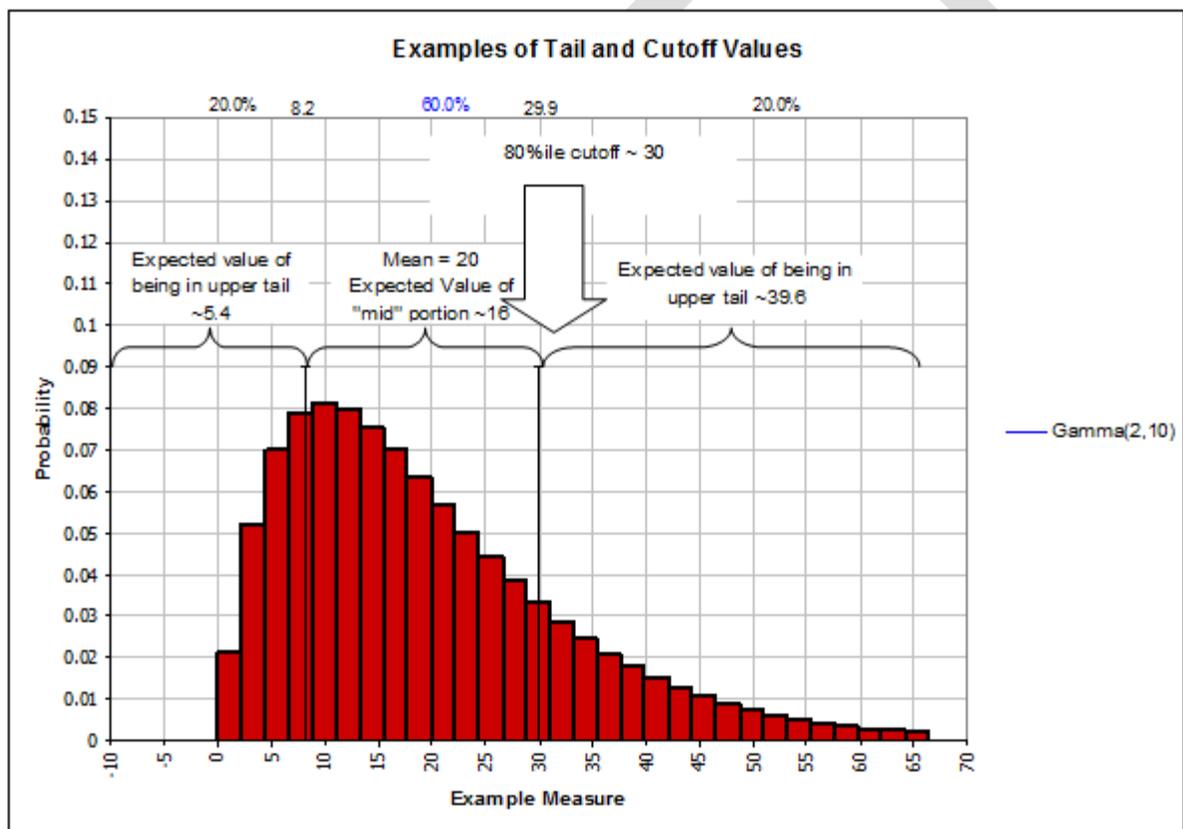
[Figure 1](#) demonstrates this using an asymmetric distribution as an example. In this case, distribution is a continuous one with a mean of 20 and a standard deviation of 14, but is noticeably skewed to the left. A way to reduce this to specific outcomes and their associated probabilities is to divide the continuous distribution into three regions (hi, mid, and low), assign the probability of being in each of these regions, and then derive the value of being in each of these regions.

In this example, the upper part of the three-part distribution was created by taking the upper 20 per cent of the distribution, which extends from roughly 30 (on the x axis) onwards. The value 30 is the 80th percentile cutoff as it has 80 per cent of the curve to the left of this point and 20 per cent of the curve to the right. The value assigned to this upper portion is the expected value of curve beyond the 80th percentile cutoff. This works out to roughly 39.6. The key point here is that the 80th percentile cutoff value of 30 is well below the value assigned to the upper tail. Following a similar process for the lower tail yields a 20th percentile cutoff value

¹ A key difference between BC Hydro's IRP analysis and that of other utilities is the IRP's use of a relatively small number of portfolios. A common approach used elsewhere is to run thousands of portfolios by repeatedly sampling from continuous distributions using a Monte Carlo process. Given the time and resources required to set up and run HYSIM and the System Optimizer, BC Hydro's analysis was limited to less than 200 portfolios in total – not nearly enough to generate an efficient frontier in risk/return space for each question examined.

of 8.2, and the expected value of the lower tail being 5.4. Again, the cutoff value and the value assigned to the tail are not the same. Finally, given that each of the extreme options was designed to have a 20 per cent probability, the middle portion of this curve will have the remaining 60 per cent of the probability assigned to it. The expected value of being in this middle portion in this example is roughly 16. Note that this is different from the mean of the overall distribution, which is 20. In general, the value of the middle section and the mean will diverge when the distribution is asymmetric.

Figure 1 Examples of Tail and Cut-off Values



It is important to note that in reducing a continuous distribution to a discrete, three-point distribution, some information is lost. In particular, the mean and the standard deviation of the two distributions will likely not be the same. The benefits, however, are that the use of a “hi/mid/low” display of uncertainty will allow the role of

uncertainty and the probability of different outcomes to be made explicit in the portfolio modelling.

Where was this used in the IRP? – Continuous distributions were reduced to discrete distributions for the DSM savings estimates and load growth forecasts.

How these discrete distributions were then used is covered in the following section.

6 Combining Scenarios using Probability Trees

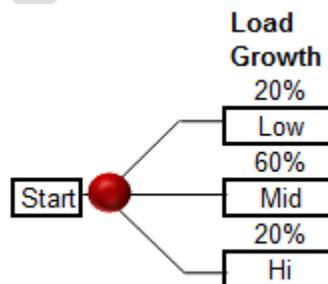
In a number of instances in this IRP, scenarios were created by combining discrete probability distributions (e.g., High, Mid, Low). These discrete probability distributions might have been discrete forecasts, or they might have been derived by segmenting a continuous distribution (as discussed above in section 5). This section will explore how the forecast DSM savings and the load forecast were used to create net demand, a key step in deriving the High, Mid and Low Gap scenarios.

A useful way to depict the scenarios derived from section 5 is in a probability tree.

[Figure 2](#) below shows how a simple probability tree can be constructed to show three discrete outcomes for load growth. By construction, the continuous distribution from the Monte Carlo analysis has been segmented into a discrete, three part distribution with the probabilities shown in [Figure 2](#).

Figure 2

Simple Probability Tree

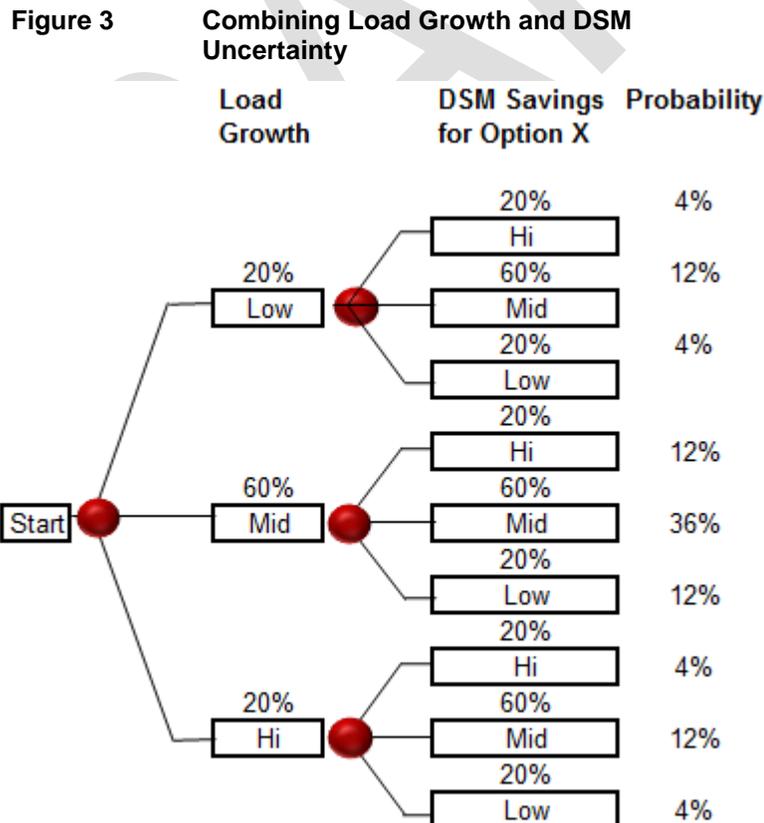


A probability tree is read from left to right. The (red) circle is a chance node where proceeding down one path (to the exclusion of the others) is treated as a chance

event. The probabilities of going down one path or another are given in the figure. Here, the chance of seeing Low load growth is 20 per cent.

The advantage to using probability trees is that they can be used to combine any number of scenarios together, as long as probabilities are assigned to the scenarios. And in combining scenarios, a probability tree allows the probability assessments to be aggregated in a logical and consistent way that is transparent to the reader. There is no real alternative in analysis when dealing with multiple scenarios. Not using a probability tree would mean trying to incorporate multiple scenarios with no guidance or structure to combining the various combinations of possibilities.

To build on the example, DSM success can then also be considered. This is shown in the following diagram where DSM forecasts have been segmented also into high, mid, and low outcomes.



Again, [Figure 3](#) is read from left to right. Tracing along the top path, there is a 20 per cent chance of seeing Low growth, and 20 per cent chance of seeing High DSM savings, giving a joint probability of 4 per cent of seeing both Low growth *and* High DSM savings.

There are some caveats required in the use of probability trees. One difficulty is that the size of the trees grows exponentially as more layers of uncertainty are added. As a result, analysis with many elements of uncertainty may either be too large to be viewed at once, or they may be simplified to focus on the key issues of interest. For instance, the 9 branches of the tree in the above figure could be simplified to Hi/Mid/Low cases of Net Demand by picking the upper-most, lower-most, and middle branches. This keeps the analysis simple, easy to explain, and captures both the range and the mid-point of the spread of outcomes. Scaling the relative likelihoods up to 100 per cent also allows these to be applied to the reduced tree. This was done in a number of instances in the IRP, using a three point distribution for Net Demand, and (where finer discrimination was needed for DSM savings uncertainty) a five point Net Demand distribution. Refer to [Figure 4](#) below for this method and the associated probabilities.

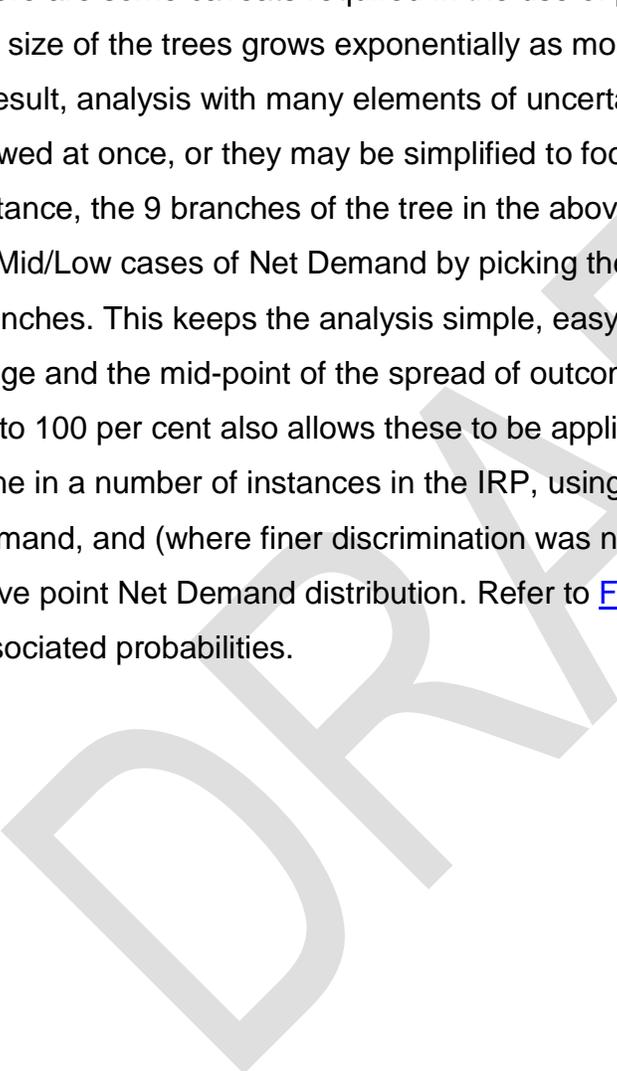
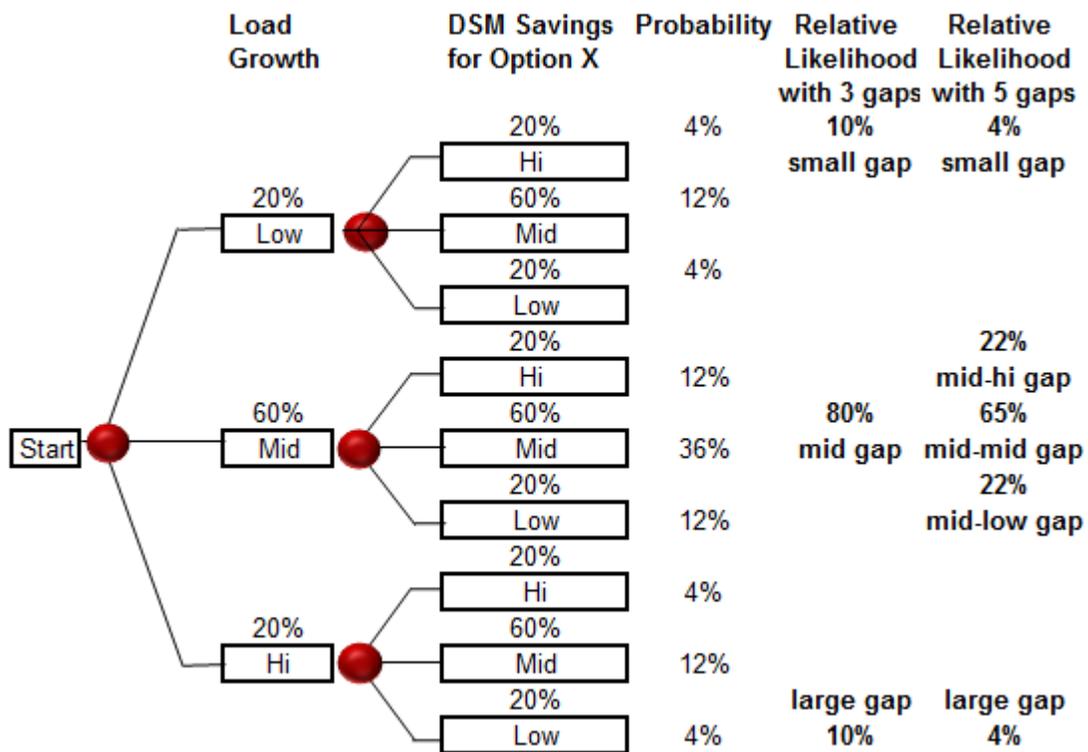


Figure 4 Net Demand – Three Point and Five Point Distributions²



A final caveat regarding probability trees is in the logic of their structure. The example used here assumes that Load Growth and DSM Savings are independent. However, it may be the case that variables included in the probability tree are related in some way. If this is true and the model is not adjusted in some way, then the results will mis-state the range of uncertainty. This happens because the model mis-estimates the probability of seeing the extreme cases (e.g. High Growth, Low DSM savings).

Testing for interdependence is a key step for model building of this type. If interdependence is suspected, then it is possible to adjust the model by estimating conditional probabilities. This might mean asking experts the probability of seeing high DSM savings, *given that High Load growth has occurred*. Such an exercise is

² Load – DSM yields Net Demand. Load – DSM – Supply yields Net Gap. The term Net Gap is used here to draw this link explicitly to Chapter 6, although additional step of subtracting supply is not shown.

perhaps even more difficult than a simple elicitation of probabilities, and an analyst must weigh off the relative merits of pursuing this approach against the implications of not capturing this interdependence in the models. Since every case is different, the best solution will differ from situation to situation.

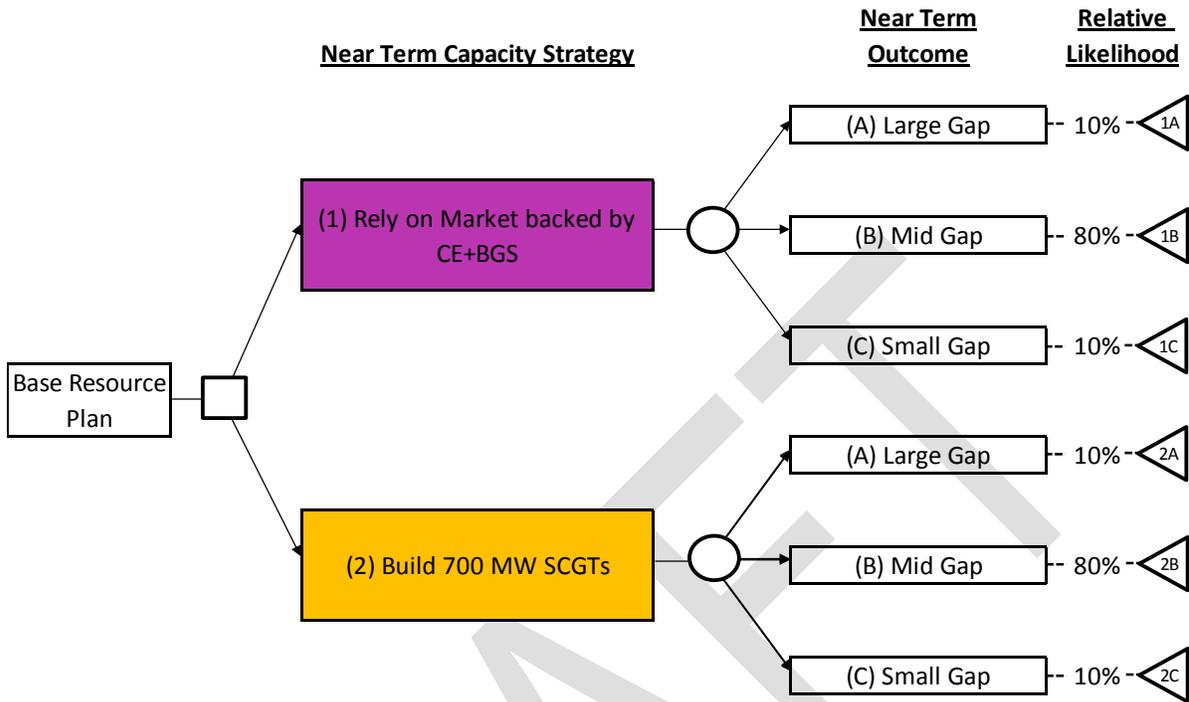
Where was this used in the IRP? – Probability trees were used to arrive at: Net Demand and Net Gap.

7 Combining Scenarios and Decisions Using Decision Trees

The tree-figures shown in section 6 are a tool to be explicit about how uncertainty is being captured in an analysis. A natural extension of this tool is to incorporate both decisions and uncertainty into a diagram. This allows clarity around what decisions are made, what uncertainties are being taken into account, and the sequence with which decisions and learning about uncertain outcomes occur.

A decision tree was used in section 6.10.5 regarding near term capacity requirements. As in the previous section, the diagram is read from left to right. The circles in the decision tree represent chance nodes where an uncertainty is revealed. Choices are represented by square nodes. In this case, a decision is made first, and then a key uncertainty is resolved. Finally, the triangles at the far right represent payoff nodes. That is, they are the end of the tree where the final outcome is recorded for that particular sequence of decisions and uncertain outcomes.

Figure 5 Decision Tree Example



8 Showing Variables Used for IRP Portfolio Modelling

Probability trees and decision trees are an excellent modelling tool to track and show explicitly how different variables fit together. One of their downsides, however, is that they grow exponentially with each new layer of uncertainty incorporated. This makes them difficult to use as a communication tool.

For this IRP, “bubble diagrams” as shown below in [Figure 6](#) were used to highlight what key variables were used for each set of portfolio runs. Here, the different Gap sizes (taken from the [Figure 4](#) above) are summarized in the top row. In this example, all five Net Gap levels (an uncertainty) were being used in testing the five DSM options. As well, the choice variable regarding the level of thermal generation to use is set at “No Additional Thermal”.

While this visual loses some information about how the combinations are made and their sequence, it is a compact way to portray which key variables are at play in each set of portfolio analyses.

Figure 6 Bubble Diagram Showing Modelling Variables of Interest



9 Conclusions

Understanding uncertainty and bringing this topic explicitly into the analysis is a key part of long-term resource planning. A considerable amount of effort and discussion within the IRP has been devoted to this topic due to its importance and also due to the unique nature of tools used to address it. However, it is important to remember that a number of key uncertainties are not easily quantified and so qualitative information is also an important part of considering uncertainty and risk. Ultimately, BC Hydro will apply its professional judgement to balance the quantitative and qualitative information amongst the competing, multiple objectives in long-term integrated resource planning.

9.1 Reference Section

Burgman, M., (2005) *Risks and Decisions for Conservation and Environmental Management*. Cambridge University Press.

Clemen, R., (1996) *Making Hard Decisions: An Introduction to Decision Analysis*. Brooks/Cole Publishing Company

Granger, M., and M. Henrion and M. Small (1992) *A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press

Gregory, R., M. Harstone, L. Failing (2012) *Structured Decision Making; a Practical Guide to Environmental Management Choices*. John Wiley and Sons

DRAFT