2008 Long Term Acquisition Plan

BC hydro

APPENDIX F14

Risk Framework – Explanation and Applications

Table of Contents

1.1 Introduction 1
1.1.1 Eliciting Subjective Probability Estimates 2
1.1.2 Creation of probability distributions by combining estimates 3
1.1.3 Portfolio Effects When Aggregating Estimates 4
1.1.4 Deriving Discrete Distributions from Continuous Ones 5
1.1.5 Combining Scenarios into Probability Trees 8
1.1.6 Applications - Probability Assessments for DSM Savings Estimates 12
1.1.6.1 Energy Savings From Rate Structures 13
1.1.6.2 Energy Savings from DSM Programs 17
1.1.6.3 Energy Savings from Codes and Standards 19
1.1.7 Total DSM Savings 20
1.1.7.1 Use of DSM Energy Savings Forecast in 2008 LTAP Analysis 22
1.1.8 DSM Programs and Capacity Savings 23
1.1.9 Limitations to DSM Probability Assessment 25
1.2 References 27

1.1 Introduction

Long term planning in the energy sector is about reflecting uncertainty and managing the risks that result from making decisions in the face of this uncertainty. Rather than spending effort developing and then justifying single point estimates, BC Hydro has been working to add more rigour to analyzing risk and uncertainty in its planning processes. To this end, BC Hydro has employed a Risk Framework in the 2008 Long-Term Acquisition Plan (**LTAP**) to increase the rigour and consistency of its treatment of uncertainty and risk in its long term planning. As explained in Chapter 5, this framework consists of qualitative applications, consisting of independent professional opinions, internal research and the application of professional judgement, as well as the quantification of subjective probability estimates. This appendix focuses on the quantitative aspect of the Risk Framework, exploring the ways in which this LTAP has attempted to bring difficult to quantify elements into a rigorous and consistent analytical structure.

This Appendix follows in six parts:

- Eliciting subjective probability estimates;
- Creation of probability distributions by combining estimates;
- Portfolio Effects When Summing Uncertain Variables;
- Deriving Discrete Scenarios from Continuous Ones;
- Combining Scenarios into Probability Trees; and
- Applications of Probability Assessments – Demand Side Management (**DSM**) Savings Uncertainty.

Each section will offer detailed explanations of the approach being discussed and will reference where this has been used in the LTAP. References, limitations of these approaches, and best practices in this area will be mentioned where pertinent. This appendix will conclude with detailed examples drawn from the DSM work to illustrate both how these approaches were used throughout this LTAP and how to interpret the findings

1.1.1 Eliciting Subjective Probability Estimates

For many variables, uncertainty can be dealt with through the use of quantitative forecasting based on historical data bases. Some judgement is needed to determine the relevance of the data sources and the nature of the forecasting model, but this also has a significant "mechanistic" approach. Many parameters, however, do not have an historical data base through which to build up a range of possible future values. For these, the range of possible outcomes is generated through professional judgement. This section will describe how this professional judgement can, in some cases, yield itself to quantified, subjective probability measures.

The use of professional judgement to put probability estimates to discrete scenarios or to generate a probability distribution around a forecast point estimate is a key element to the LTAP Risk Framework. Such judgements are not done without difficulty, subjective probability judgements are hard to do and are subject to well documented perils (c.f. Burgman, Ch 1). However, researchers in this field of decision science have developed tools and protocols to improve performance on these difficult tasks; assessments used in the LTAP 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 LTAP 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 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 probability elicitation (cf Clemen, R (1996)). In their work on risk and uncertainty in decision analysis, Granger and Henrion emphasize that while the process of probability elicitation from professionals have shortcomings, it is the "only game in town" (Morgan and Henrion, 1990) when trying to incorporate uncertain estimates into a framework of analysis.

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

- DSM Energy Savings – from rates, codes and standards, and programs;
- DSM Capacity Factors;
- Greenhouse Gas Offset Cost Scenarios;
- Natural Gas Price Forecast Scenarios; and
- Independent Power Producer (**IPP**) Attrition.

1.1.2 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 professional 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 the 2008 LTAP, 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 (all Monte Carlo simulations for the Risk Framework were done 10,000 times), the stored results form a distribution of possible outcomes for the sum of the random sub-tasks. 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 the LTAP? Monte Carlo simulations were used to estimate the dispersion of outcomes around a sum of uncertain variables in a number of instances for this LTAP, including: DSM savings, IPP Attrition, and Load Growth.

1.1.3 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 the covariance will add to the variance.

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.

Estimating these uncertain relationships can be assisted through the structured elicitation of probabilities from participants. Structured conversations, visual aids, and a knowledge of the strengths and weaknesses of these elicitations can help. However, it needs to be kept in mind that these judgements are not precise.

The LTAP Risk Framework relied on professionals' 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 LTAP? – Correlations amongst uncertain variables were used in estimating DSM savings. Details are provided in this appendix, Section 1.1.6.1 and 1.1.6.2.

1.1.4 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 LTAP required that parameters be assigned specific values so that they could be used as inputs into the analysis. This required that the specific outcomes be selected out of the whole distribution of possible outcomes. 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 F14-1 demonstrates this using a Gamma (2, 10) 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. As this figure shows, a skewed distribution will have the most likely outcome (i.e. the mode), the expected outcome (i.e. the mean) and the middle point of the distribution (i.e. the median) falling at different values.

Figure F14-1 – Example of Skewed Continuous Distribution

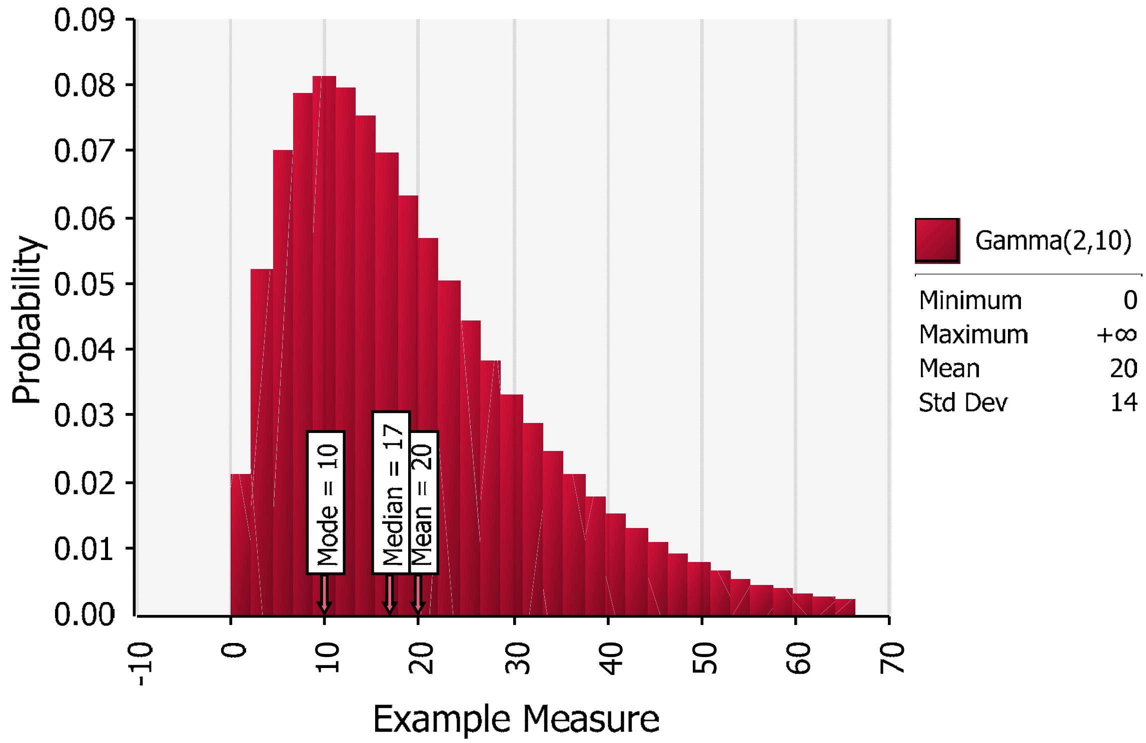

A way to reduce this continuous distribution of values to specific outcomes and their associated probabilities is to divide the continuous distribution into regions (e.g. hi, mid, and low), assign the probability of being in each of these regions, and then derive the expected value of being in each of these regions.

In this example, the upper part of the three part distribution was created by taking the upper 20% of the distribution, which extends from roughly 30 onwards. The value 30 is the 80%ile cutoff as it has 80% of the curve to the left of this point and 20% of the curve to the right. The value assigned to this upper portion is the expected value of curve beyond the 80%ile cutoff. This works out to roughly 39.6. The key point here is that the 80%ile cutoff value of 30 is well below the value assigned to the upper tail. Following a similar process for the lower tail yields a 20%ile cutoff value 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% probability, the middle portion of this curve will have the remaining 60% 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 three statistics describing the central tendency shown in Figure F14-1 - the value of the middle section and these statistics will diverge when the distribution is asymmetric.

In general, this approach can be done through approximating the areas under the curve for any given distribution. However, for the load forecast, the assumption of normality allowed an exact calculation of the areas under the tails to be carried out.

Figure F14-2 - Examples of Tail and Cutoff Values

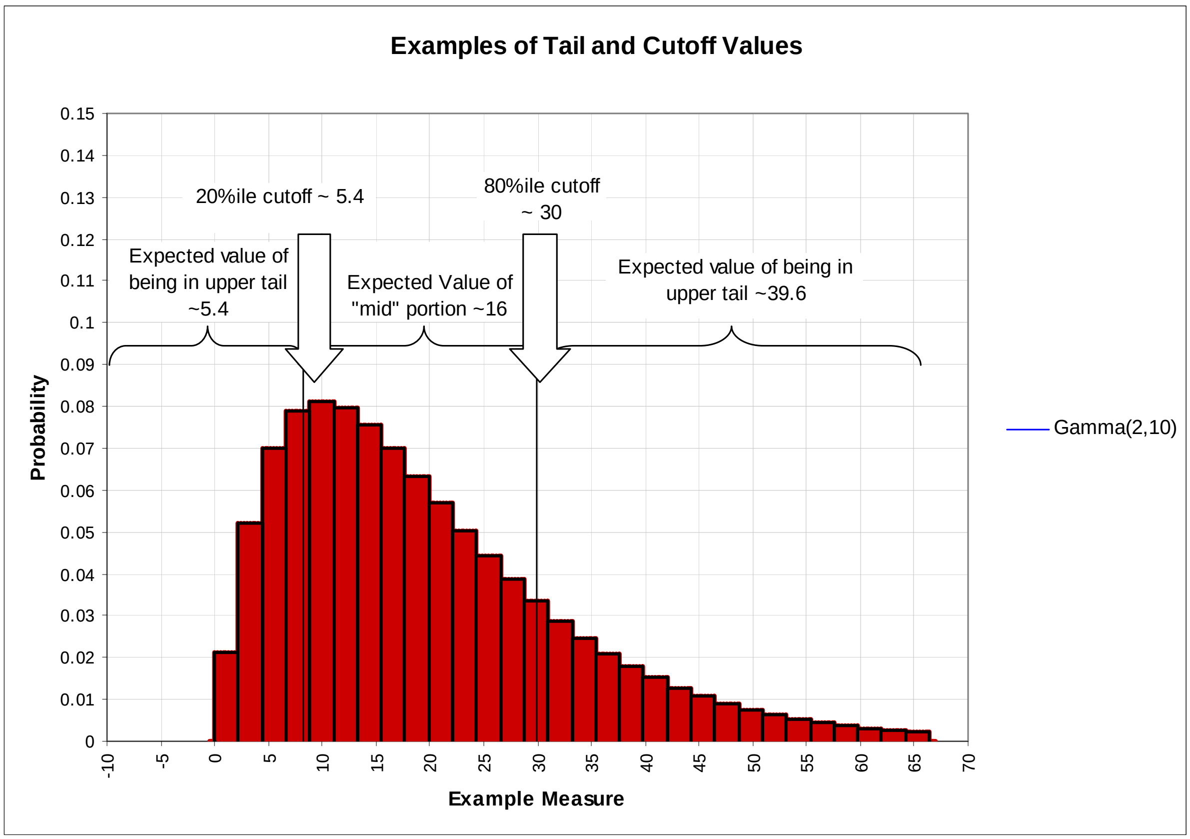

It is important to note that in reducing a continuous distribution to a discrete, three-point distribution, some information is lost. In particular, the mean and the standard deviation of the two distributions will likely not be the same. The benefits, however, are that the use of a "hi/mid/low" display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

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

1.1.5 Combining Scenarios into Probability Trees

The probability elicitation process for this LTAP resulted in scenarios that captured hi and low cases as well as a mid case, each with a probability attached to them. It is important to note that these probabilities are best viewed as relative likelihoods to each other. For instance, imagine that the future price of gasoline is captured in the following way:

- High gasoline case, price of $6.00/l in 2020 with probability 10%;
- Mid gasoline case, price of $4.00/l in 2020 with probability 70%; and
- Low gasoline case, price of $1.00/l in 2020 with probability of 20%.

A correct way to view this is that the estimated chance of seeing gas prices in the neighbourhood of $6.00/l in 2020 is much smaller (about 1/7) than the chance of seeing gas in the price range of $4.00/l. And so while these scenarios do not include a price forecast of $4.50/l in 2020, it does not mean that this forecast is saying this has a 0% chance of occurring. Rather, the simplified approach has picked three possible values to represent hi, mid and low points on a broader range, and has attached probabilities of being in those portions of this range.

A useful way to depict the scenarios described above is in a probability tree. Figure F14-3 below shows how a simple probability tree can be constructed for the gasoline example developed in the previous section.

Figure F14-3 – Probability Tree for Gasoline Price Uncertainty

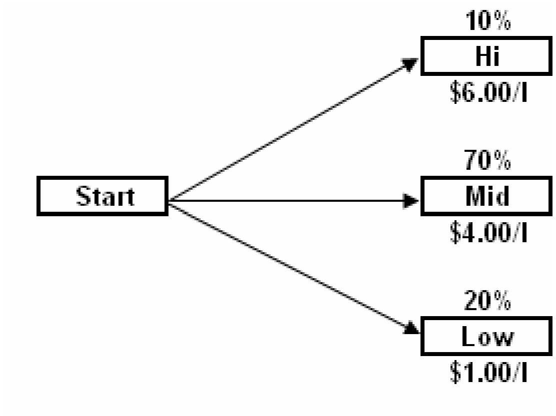

The diagram is read from left to right. The box labelled “start” is a chance “node”, where proceeding down one path (to the exclusion of the others) is treated as a chance event. The probabilities of going down one path or another are given to the right. So the chance of going down the “Hi” path is 10%. The different paths are labelled here as “Hi”, “Mid”, and “Low” to represent the three scenarios being described. As well, the prices associated with each scenario are shown.

The advantage of using probability trees is that they allow different scenarios to be combined in a logical and consistent way. For instance, imagine that problem at hand is the cost of operating a car, and the two key drivers of uncertainty are the price of gasoline in 2020 and the cost of parking in 2020. Also, imagine that three scenarios regarding the price of parking are as follows:

- High parking costs, price of $100/day in 2020 with probability 20%;
- Mid parking costs, price of $30.00/day in 2020 with probability 50%; and
- Low parking costs, price of $20.00dayl in 2020 with probability of 30%.

These separate scenarios can be put together in a probability tree, as shown in Figure F14-4 below.

Again, this tree is read from left to right, with labels denoting the probability and consequence of being on a particular branch of the tree. Tracing along the upper branch of the tree, there is a 10% chance of being in world of $6.00/l gas in 2020 and a 20% chance of being in a world of $100/day parking. These two events can be combined into a hi gas, hi parking cost world. The probability tree suggests that there is a 2% chance of this happening. Recall that this is not a forecast but rather a relative likelihood; it is about 1/3 as likely as being in a low gas, low parking cost world (the lowest branch on the tree). Note as well that there are no decisions in this tree; the probability tree outlines the cost risks to a car owner of future gas and parking prices but there is no decision making on the part of the car owner depicted in this figure.

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.

Figure F14-4 – Probability Tree for Gas and Parking Price Uncertainty

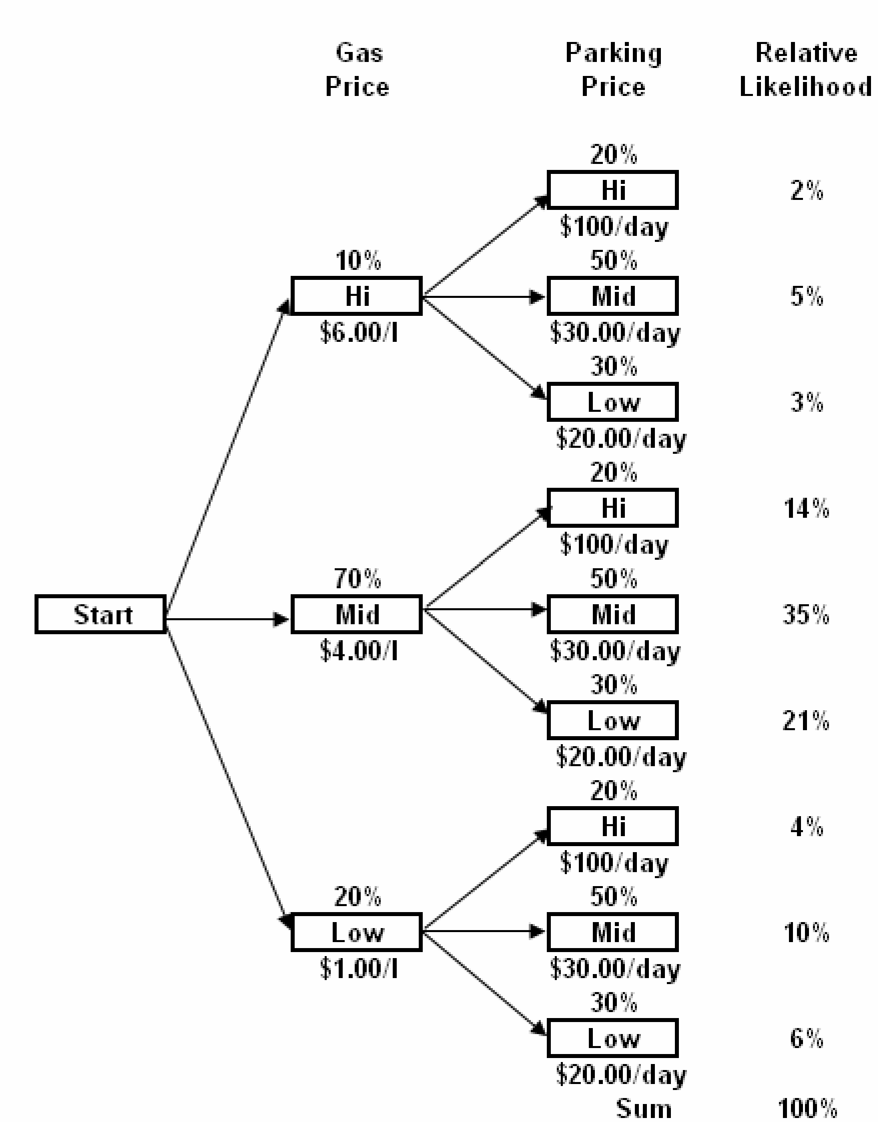

There are some caveats required in the use of probability trees. One difficulty is that the size of the trees grows exponentially as more layers of uncertainty are added. As a result, analysis with many elements of uncertainty may either be too large to be viewed at once. To address this, the 9 branches of the tree in the above figure could be simplified to Hi/Mid/Low cases of "Cost of Car Use" 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% also allows these to be applied to the reduced tree. However, this simplification does discard some intermediate cases that might be of interest.

A final caveat regarding probability trees is in the logic of their structure. The example used here assumes that the cost of gas and the cost of parking are independent. However, it may be the case that variables included in the probability tree are related in some way. For example, it may be the case that hi gasoline costs and high parking costs are positively correlated (i.e. parking costs tend to be high when gas costs are high, and *vice versa*). If this is true and the model is not adjusted in some way, then the results will tend to underestimate the true range of uncertainty. This happens because the model underestimates the probability of seeing the extreme cases (e.g. high gas costs and high parking costs).

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 knowledgeable individuals the probability of seeing high parking costs, *given that the high gas price world has occurred.* Such an exercise is 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 LTAP? – Probability trees were used to arrive at: Net Demand, Cost of Thermal Operations, and the Basic Scenarios (11-branch, 9-branch, and 5-branch). Interdependence of load growth and DSM deliverability was identified as a potential issue. In the analysis for the 2008 LTAP, there was not enough evidence of a strong net bias in one direction or another, relative to the precision of elicited conditional probabilities, to suggest that it would be beneficial to build such interrelationships into the probability tree.

1.1.6 Applications - Probability Assessments for DSM Savings Estimates

The previous sections of this appendix gave a generic description of some of the approaches used in quantifying subjective probability estimates. This final section of the appendix will demonstrate how these were applied in the case of DSM savings estimates. As stated at the start of this appendix, fully incorporating uncertainty into the analysis requires a mixture of qualitative assessment, quantitative analysis and professional judgement. This application of probability assessments is focusing on quantitative methods of probability assessment as it is an application of new set of tools for the LTAP. This final section will conclude with some of the limitations of the probability estimates used in the DSM analysis.

The main elements and details of DSM (Option A and B) are laid out in Chapter 3 and Appendix F-17. However, the savings arising from DSM are estimates that are subject to considerable uncertainty. One of the goals of the LTAP risk framework was to avoid relying on point estimates – rather the focus was on developing a range of possible outcomes with the associated probability of being within this range.

This section will focus on the three DSM tools:

- Rate Structures;
- Codes and Standards;
- DSM Programs.

The key sources of uncertainty underlying the energy savings estimates will be highlighted for each of these DSM tools, and the range of possible outcomes will be reported for each. This section will conclude with the overall spread of possible energy savings outcomes from DSM.

1.1.6.1 Energy Savings From Rate Structures

The focus of the DSM savings arising from rate changes was on the implementation of a tiered or stepped rate system. It is believed that when customers are faced with a higher marginal cost then they will be motivated to reduce their consumption levels. For this LTAP, two-step inclining block rate structures were the basis of analysis for all customer groups (residential, commercial, and industrial).

There are a large number of factors influencing the success of conservation rate structures. BC Hydro highlighted three key drivers of uncertainty that could significantly impact the amount of savings arising from rate design changes. They were:

- The elasticity of demand;
- The timing of the rate change; and
- The magnitude of the 2nd block price.

For each of these issues, a range of values was adopted that was broad enough to capture the potential range of what might be possible, but narrow enough so that the upper and lower bounds considered were still within the realm of possibility. Table F14-1 shows the range and probabilities associated with elasticity of demand assumptions. As these were covering a wide range of possible outcomes, no attempt was made to further differentiate these parameter estimates among customer classes.

Table F14-1 – Elasticity of Demand Estimates

	Hi	Mid	Low
Elasticity of Demand Coefficient	**-0.15**	**-0.10**	**-0.05**
Probability	**30%**	**60%**	**10%**

Ranges of possible outcomes were also developed for the timing and magnitude of the rate changes. To simplify the estimation process and to reduce the number of variations considered, these were combined to capture a high case, a low case, and two mid cases. Table F14-2 below shows the result of this analysis.

Table F14-2 – Timing and Size of Tier 2 Price Changes

	Hi	Upper Mid	Lower Mid	Low
Timing and Tier 2 Cap	**2008, $0.15/kWh**	**2008, $0.12/kWh**	**2009, $0.12/kWh**	**2009, $0.10/kWh**
Probability	**10%**	**40%**	**40%**	**10%**

The goal of the risk framework was to generate a range of possible outcomes with their associated probabilities. Putting together the three elasticity values and the four pairs of timing and price levels gave twelve different elasticity/timing/price level combinations. Rate savings for each customer class were estimated for each of these combinations yielding a range of possible savings arising from rate changes with an associated probability for each outcome.

A simple additive model was assumed for the total rates savings across all customer classes; the total rates savings was calculated as the sum of rate savings from each rate class.

However, calculating the appropriate range of rate savings required more than just adding up the individual results. It was identified that the performance of the rate structures across rate classes was probably interrelated; capturing some aspects of these correlations was needed in order to get a better sense of the spread of uncertainty around the average.

Concerning the elasticity measures, BC Hydro adopted the view that the ability to change with respect to changing energy prices is probably similar amongst the commercial and industrial customers. However, these groups ability to change are probably not correlated to the ability of residential customers to change in the face of new rate structures.

The timing of the rate changes across sectors was seen to be roughly related. If rate changes are slow to be developed, passed through the regulatory process, or implemented for one set of customers, then it is likely that this delay will also occur with rate changes for other customer classes as well.

Finally, there is likely to be a positive relationship seen between the size of the Tier 2 cap across sectors as this is driven by a combination of common policies, considerations of fairness, and marginal cost of incremental energy from other sources.

Table F14-3 summarizes this in the following way. Savings arising from rate changes seen at the business level (commercial and industrial) were estimated to be strongly correlated (0.8), but only loosely correlated to savings seen from residential customers (0.3).

Table F14-3 – Interrelationships Among Rate Structure Savings, by Sector

	Residential	Commercial	Industrial (Small)
Residential	**1**		
Commercial	**0.3**	**1**	
Industrial	**0.3**	**0.8**	**1**

Total rate structure savings was taken as the sum of rate structure savings from each customer class.

Figure F14-5 – DSM Rate Structure Savings

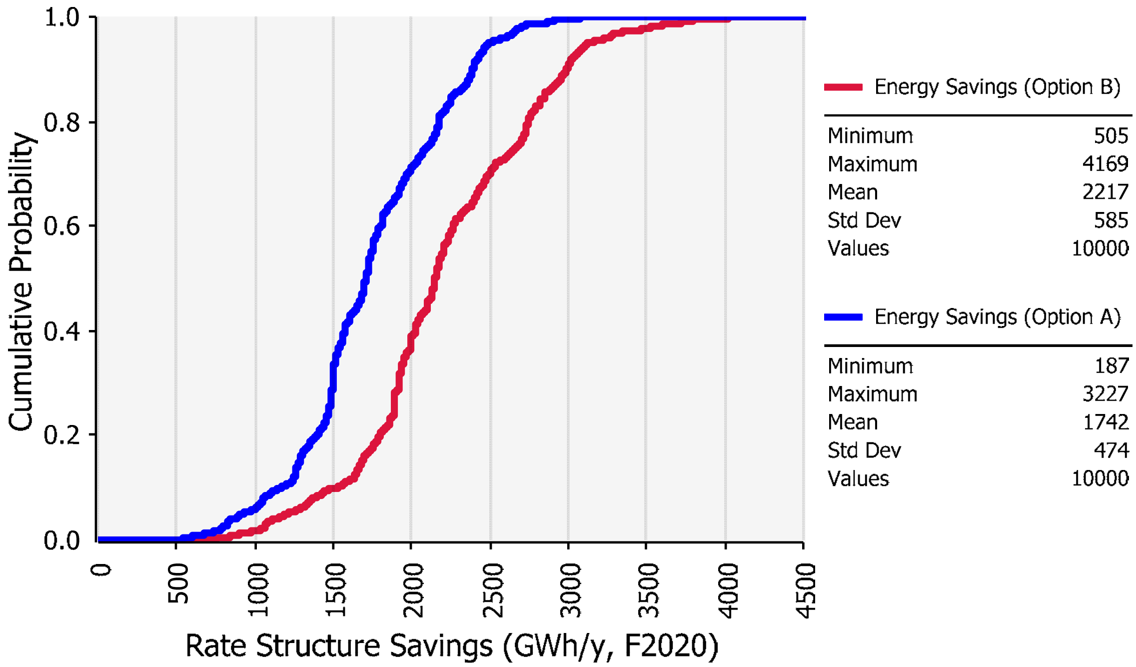

Savings from each customer class was a random variable depending on which of the twelve different elasticity/timing/price level combinations were used. Taking interrelationships into account, a Monte Carlo simulation analysis was carried out using 10,000 random draws to calculate the total savings across all customer classes. The results are shown in Figure F14-5.

These results, when combined with energy savings from codes and standards and from DSM programs, give the total level of energy savings available.

1.1.6.2 Energy Savings from DSM Programs

DSM programs consist of roughly twenty-one individual programs spread across the residential, commercial and industrial customer groups. These include programs from residential lighting and appliance replacement to manufacturing plant design and optimization with large industrial customers. While BC Hydro has extensive experience working with its customer groups to encourage energy conservation and efficiency, the size and scope of these new DSM programs and their interdependence on rate changes and codes and standards suggested that forecasts around energy savings were subject to significant uncertainty. As a result, the delivery risk of these energy savings was made a key part of the LTAP Risk Framework’s quantitative analysis.

While each individual DSM program was planned to achieve a certain level of energy savings by 2020, this level of energy savings per program was subject to two key drivers of uncertainty;

- The participation rate of customers in that program; and
- The energy savings per participant.

Total savings from DSM programs was estimated as simply the sum of the twenty-one individual DSM programs. However, calculating the spread of uncertainty was more involved as DSM program managers identified some significant interdependencies amongst some of the key drivers of uncertainty; failing to capture these interrelationships would significantly understate the spread of uncertainty around DSM Program savings.

The first key interrelationship estimated was participation rates. It was felt that BC Hydro’s DSM promotion and marketing strategy and the creation of a “conservation culture” through raised awareness, advertising or accessible energy saving information could be seen as a common influence across all sectors. As a result, BC Hydro estimated a moderate correlation of 0.5 amongst participation rates across all programs and across all customer classes.

A second set of important relationships was that between the participation rate and the savings per participant for each project. It was felt that a DSM program that delivered more savings per participant would draw in more participants, but that those that did not deliver high savings per participant would see fewer participants involved. BC Hydro estimated that there would be a correlation of 0.3 for this relationship. This means that these variables would tend to move together, but that it would not be uncommon to find higher than planned participation rates even when savings for that program were lower than expected and lower than planned participation rates for some programs even when savings per participant were higher than expected.

Total DSM program savings were taken as a sum of the savings from the twenty-one individual programs. The savings from each program was a random variable, depending on the product of the participation and the savings per participant. Taking into account interrelationships amongst these random variables, a Monte Carlo simulation analysis was carried out using 10,000 random draws to calculate the total savings across all customer classes. The results are shown in Figure F14-6.

Figure F14-6 – Energy Savings From DSM Programs

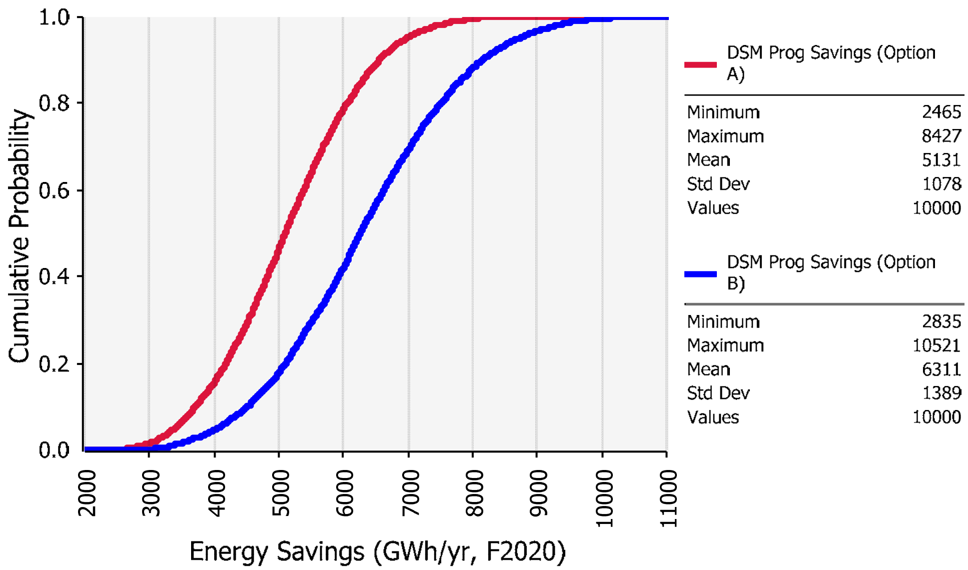

1.1.6.3 Energy Savings from Codes and Standards

BC Hydro included roughly 30 potential changes to Codes and Standards in its DSM options. These ranged from regulations for standby power sources to commercial building codes and covered both federal and provincial regulations.

The savings estimates from changes to Codes and Standards were subject to many forms of uncertainty including the timing and magnitude of any potential changes. Program designers were taken through several workshops to elicit ranges and probabilities around these energy savings estimates. These conversations followed similar guidelines to those used when discussing the potential success of DSM programs.

The total amount of energy savings attributable to changes Codes and Standards was treated as the sum of all of the individual impacts. However, the spread of uncertainty around this estimate required additional analysis as BC Hydro staff identified several ways in which the success (or lack of success) amongst the programs would probably be correlated.

Staff identified several common drivers of uncertainty for changes to Codes and Standards including continued support at the corporate level, the provincial level and the federal level of Government. This suggested that the success or lack of success with savings in this area might tend to move together. However, this correlation is not perfect as there are different levels of decision makers, and each code and standard is also subject to resistance or support from other stakeholders. The estimate was that this effect could be strong, but there was not enough information to differentiate amongst different levels of Government/ decision makers. As a result, the LTAP risk framework used a correlation of 0.6 amongst all Codes and Standards changes.

Savings from changes to Codes and Standards was calculated as the sum of the savings arising from each of the roughly twenty changes. The savings from each change was a random variable. Using the estimated energy savings, the probability distributions, and the assumed relationships among these efforts, a Monte Carlo simulation analysis was run using 10,000 random draws to estimate the expected level of energy savings and the spread of outcomes around this mean. These results are summed up in Figure F14-6.

Figure F14-6 – DSM Energy Savings from Codes and Standards

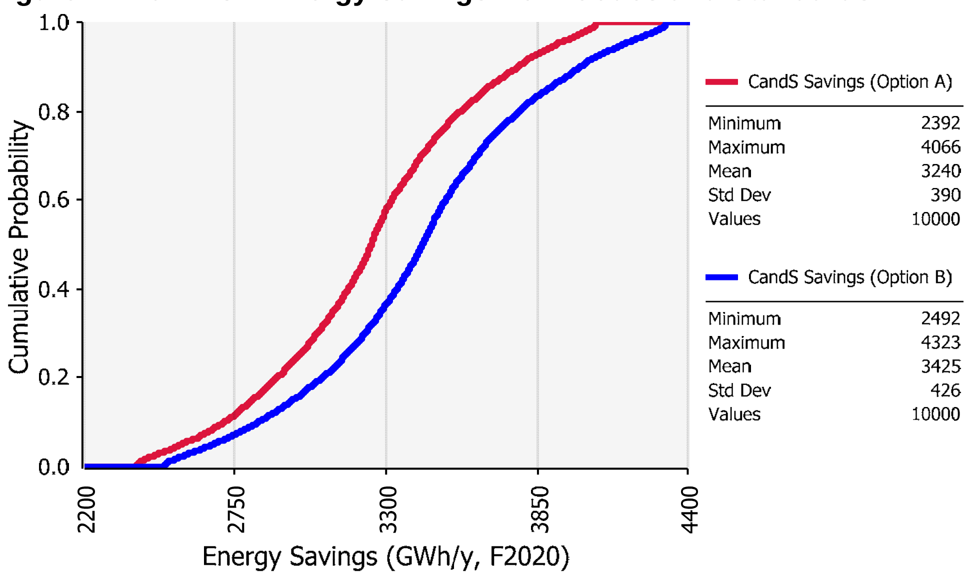

1.1.7 Total DSM Savings

The total energy savings from DSM was modelled as a sum of the energy savings across customer classes arising from DSM Programs, changes to Codes and Standards, and rate structure changes. It was assumed that the level of energy savings arising from each of these three areas (Program, Codes and Standards, and Rate Structures) was independent from each other. This assumption allowed the modelling and costing estimates to be tied to specific DSM activities levels, which in turn, allowed the incremental cost of DSM delivery to be examined across the Hi, Mid, and Low gaps. The limitations of this assumption are addressed in Section 1.1.9 below.

A Monte Carlo simulation analysis was carried out and 10,000 random draws were used to calculated the mean and spread of outcomes for all DSM savings. The results for this are shown in Figure F14-7.

Figure F14-7 – Total DSM Savings

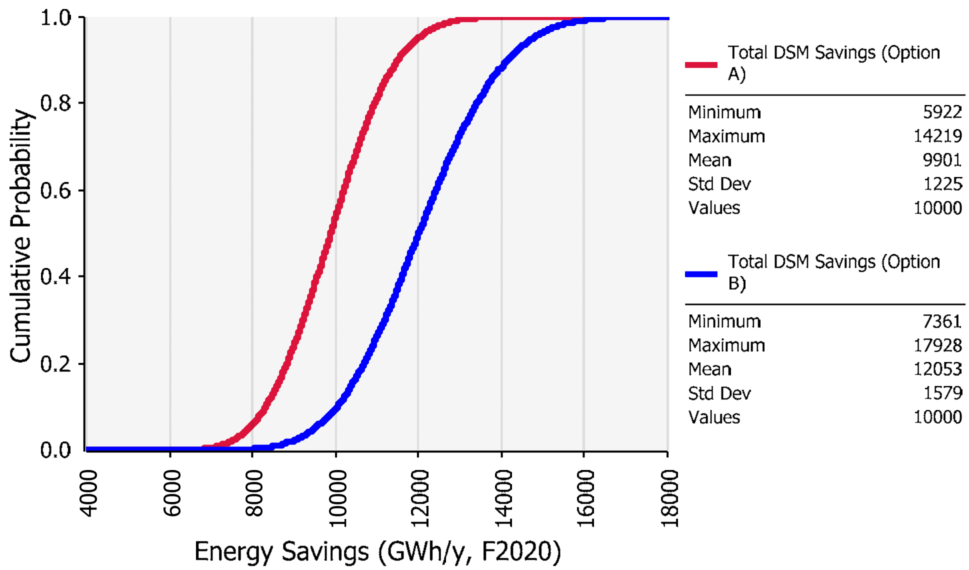

As this figure shows, moving from Option A to Option B increases the expected level of savings from roughly 10,000 GWh/y in 2020 to roughly 12,000 GWh/y in 2020. But this change also increased the spread of possible outcomes around the mean, indicating that more extreme upside and downside deviations from the expected outcomes are likely with the larger DSM effort.

There are many measures of downside risk. To pick an arbitrary benchmark, say that the risk of interest is being short of energy by 2000 GWh or more by F2020. This can be marked on the distributions of total DSM savings, as done in the following figure.

Figure F14-8 - Delivery Risk and DSM Savings

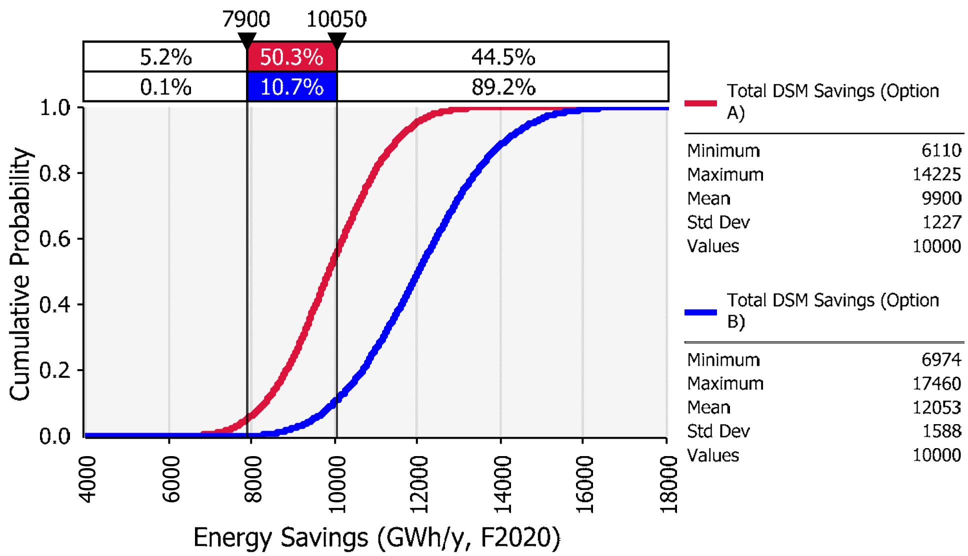

As this figure shows, there is roughly a 5% chance that Option A will be 2000 GWh or more short of energy by 2020, whereas Option B has more than a 10% chance of missing its expected level of performance by that magnitude by 2020.

1.1.7.1 Use of DSM Energy Savings Forecast in 2008 LTAP Analysis

Given the size of the estimated DSM energy savings and the spread of uncertainty around these estimates, delivery risk was highlighted as a key element of focus for the LTAP risk framework.

To incorporate the uncertainty estimates of DSM energy savings into the analysis, the spread of outcomes was reduced into a three point distribution: Hi, Mid, and Low. This accomplished two objectives. It allowed actual levels of DSM energy savings, costs and other impacts to be modelled explicitly for the portfolio analysis. Secondly, it allowed a simple representation of delivery risk to be used with the probabilities of the DSM savings.

The following table shows the values derived from the probability density curves above.

Table F14-4 – Low, Mid, and High Ranges for DSM Energy Savings

	Low	Mid	High
Probability	**20%**	**60%**	**20%**
		Moderate DSM	
GWh/y savings, 2020	**8,500**	**10,200**	**12,000**
		Aggressive DSM	
GWh/y savings, 2020	**9,900**	**12,000**	**14,300**

Recall that, as a distribution is reduced from a continuous one (as in Section 1.1.4 above), to a three-point distribution, some information is lost. One effect of this is that the asymmetry of the distributions makes the "mid" values different from the means reported in the distribution of DSM energy savings.

DSM energy savings, when paired with expected load growth, determines the size of the load/resource gap that needs to be filled over time with additional resources. Treating DSM savings and load growth as different processes was a simplification that allowed the separate impacts of load growth (driven by economic growth, population change, etc) and DSM savings (driven by elasticities of demand, government initiatives, participation rates, etc) to be varied. The LTAP resource planning analysis then considered different sizes of the load/resource gap along with their relative likelihoods of occurring. The spread of possible outcomes around DSM capacity savings were an important input for considering Contingency Resource Plans (**CRPs**).

1.1.8 DSM Programs and Capacity Savings

Chapter 5 and the CRPs in the LTAP both touch on the relationship between DSM energy savings and capacity savings. This section will show how the probability assessment in the Risk Framework was applied to the topic of capacity savings.

In addition to the uncertainty previously discussed regarding DSM energy savings, there is an additional level of uncertainty introduced when calculating the capacity savings associated with these energy savings.

BC Hydro staff identified three key drivers of uncertainty when estimating capacity factors:

- Measurements of peak load;
- The shape of energy savings applied to load; and
- Extrapolating results into the future (foreacasting).

Since these influences may vary, depending on customer class, the relative importance of each of these was considered for the residential, commercial and industrial customers. A range of capacity factors and a best estimate was derived for each customer class that took into account these factors. These ranges are reproduces in Table F14-5.

Table F14-5 – Capacity Factor Estimates, by Customer Class

	Customer Class Share of Energy	Capacity Factor Estimates (MW per GWh/y)		
		Lower Bound	Most Likely Estimate	Upper Bound
Industrial	32%	0.11	0.12	0.14
Commercial	22%	0.12	0.14	0.2
Residential	46%	0.16	0.2	0.3

Using these estimates, triangular distributions were used to capture the range of possible values and their associated probabilities for capacity factors A Monte Carlo simulation was undertaken using 10,000 draws to derive a distribution of possible capacity factor values across all customer classes. When multiplied by the energy savings in each customer class and summed, this gave an estimate of total capacity savings for DSM Option A. This is reproduced below in Table F14-6. No similar analysis was undertaken for Option B.

Figure F14-6 – Capacity Savings from DSM (Option A)

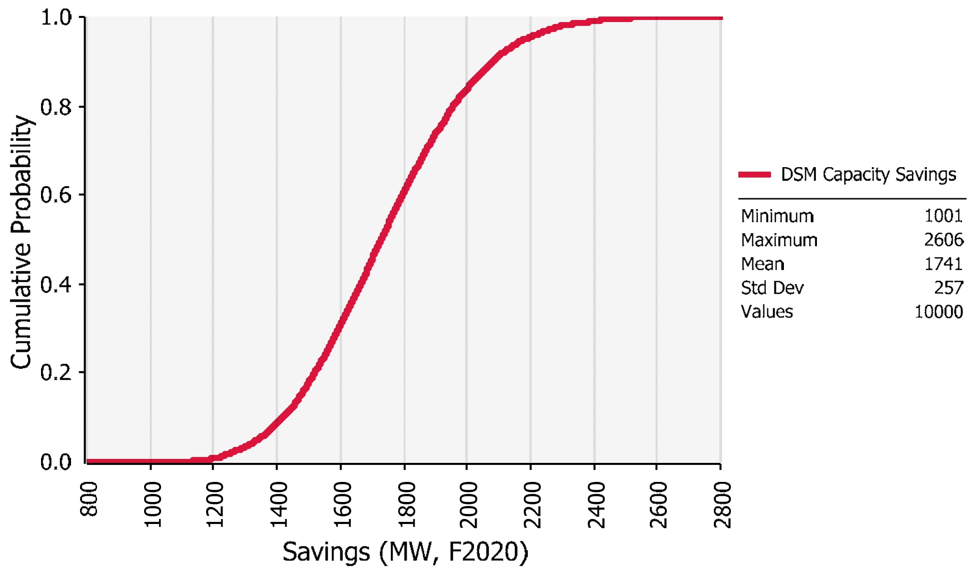

1.1.9 Limitations to DSM Probability Assessment

Estimating the range of uncertainty around future DSM savings was a new activity for BC Hydro that involved implementing a new set of analytical tools across a range of modelling activities that were not coordinated previously. Given that the DSM options involved new DSM tools and an unprecedented level of effort on the part of BC Hydro, it is important to be clear about the limitations of the uncertainty estimates regarding future DSM savings. The goal of this is to supplement the quantitative probability analysis with some additional, qualitative factors. This, in turn, will suggest additional work that may be valuable in the future as BC Hydro gains more experience with this increased level of DSM activity.

Timing of savings – the uncertainty analysis focused on the quantity of energy to be saved by F2020. The path of energy savings to that target was also uncertain, and could vary from the planned path. It was recognized that adding this topic would have introduced an additional layer of complexity that would stretch the capacity of the analysis. However, with DSM savings predicted to grow at over 100 MW per year up to 2020, a year's delay or a slower ramp up could have a significant impact on planning for the early stages of the planning period.

Delivery Risk – the probability analysis focused on the delivery of the DSM Plan, assuming that these efforts were put in place as designed and carried out as planned. Discussions did not explicitly address the risk that the programs may not be enacted at all, or carried out only in part due to internal or external reasons.

Model Specification – the probability assessment focused mainly on uncertainty around the input values (price elasticity of demand, savings per participant, etc). Some additional effort was made to address the relationships among these inputs but only *within* each DSM tool. More sophisticated approaches are possible, that capture explicitly the dynamic interplay *among* the DSM tools (e.g. how the performance of codes and standards and rates influence DSM programs), and among the DSM tools and other impacts on demand (e.g. DSM savings and load growth). As a first attempt, BC Hydro used a more simple approach at modelling DSM savings estimates.

Integration with Load Forecasts – BC Hydro’s load growth models are mostly driven by “top down” estimates, whereas DSM savings are mostly derived as “bottom up” calculations. Currently these two efforts are done separately and their integration is a challenge. Until more work is done to draw precise linkages between these models, additional caution, over and above the quantitative portion of the risk framework, is warranted when using the results for planning purposes.

Making predictions for the purpose of long term energy planning is a necessary but difficult task. The goal of applying the quantitative tools of the risk framework to this topic was to move analysis away from using point estimates and towards using ranges and probability distributions of outcomes. This appendix has pointed out some of these tools, their uses and limitations. It has highlighted the ongoing need to use a mixture of quantitative analysis and professional judgement in drawing from the analysis. Time and experience will allow model validation and continued improvement in the application of the Risk Framework to energy planning topics.

1.2 References

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

Clemen, R. (1996) *Making Hard Decisions,* Duxbury Press, Pacific Grove, USA

Golderberger, A (1991) *A Course in Econometrics,*. Harvard University Press, Boston, USA

Morgan, M. and M. Henrion (1991) *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*, Cambridge University Press, Cambridge, UK