ORNL/CON-399

OAK RIDGE NATIONAL LABORATORY

MARTIN MARIETTA

An Approach to Understanding, Representing, and Managing Uncertainty in Integrated Resource Planning

Bruce E. Tonn
Anthony J. Schaffhauser

MASTER

MANAGED BY
MARTIN MARIETTA ENERGY SYSTEMS, INC.
FOR THE UNITED STATES
DEPARTMENT OF ENERGY

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

ORNL/CON-399

ENERGY DIVISION

AN APPROACH TO UNDERSTANDING, REPRESENTING, AND MANAGING UNCERTAINTY IN INTEGRATED RESOURCE PLANNING

Bruce E. Tonn
Oak Ridge National Laboratory

Anthony J. Schaffhauser
University of Tennessee-Knoxville

December 1994

Sponsored by
Office of Energy Efficiency and Renewable Energy
U.S. Department of Energy

Oak Ridge National Laboratory
managed by
Martin Marietta Energy Systems, Inc.
under Contract No. DE-AC05-84OR21400
for the U.S. Department of Energy

Table of Contents

EXECUTIVE SUMMARY

This report addresses the issue of uncertainty in integrated resource planning (IRP). IRP is a process employed by electric utilities, often at the behest of Public Utility Commissions (PUCs), to evaluate the acquisition of resources to meet forecast energy demands and other criteria such as energy efficiency and fuel diversity. Uncertainty plagues the preparation of IRPs, from forecasting energy prices and electricity demand to estimating costs and benefits expected from new resource acquisitions. As a result, all those involved in the IRP process (including PUC commissioners and staff, utility decision makers, IRP analysts, and experts in particular areas) need to appreciate how uncertainty affects IRP and learn about available techniques to reduce its effects.

All parties can gain a shared understanding of uncertainty in key IRP estimates by separating uncertainties into the following three categories: irreducible, operational, and application. Irreducible uncertainty refers to limitations on what *can* be known due to inherent and insurmountable factors. Long range forecasts of aspects of large-scale, disorderly, and chaotic systems entail high levels of irreducible uncertainty, which time and money are powerless to overcome. Operational uncertainty refers to limitations on what *is* known due to flawed theory, data, and models. Application uncertainty refers to limitations on *using* what we know for a particular problem or decision. For example, over-generalizing results from one study to another context causes application uncertainty. The effect of each type of uncertainty can be minimized by increasing the flexibility of the integrated resource plans through the use of imprecise probability.

Imprecise probability allows the specification of lower and upper bounds on probabilities, as opposed to the single number required by precise (classical) probability. Imprecise probability is able to represent different levels of uncertainty with different ranges between lower and upper probabilities. In contrast, precise probability presents data in such a manner that suspect information is often indistinguishable from accurate information.

Methods are available to calculate upper and lower expected values using imprecise probability. These methods are straightforward for IRP analysts to use. Using imprecise probability also simplifies the IRP process because the bounds are firm. Second-guessing sensitivity analyses or "tweaking" of inputs due to uncertainty about their values are not necessary.

In addition, several methods are available to IRP analysts for synthesizing numerous, seemingly disparate pieces of information that could be of value to the IRP process. These include combination, consensus, and conditionalization methods. Consensus methods provide a means to synthesize expert judgment and/or model outputs for a particular estimate. Combination methods provide a means to synthesize pieces of evidence that can be used to develop an estimate or judgment. Conditionalization methods provide a means to update estimates in light of new evidence and to sharpen estimates by using ancillary statistical data. Each is compatible with imprecise probability.

The major uncertainty concepts and methods presented in the report are illustrated using analytical IRP problems presented in Portland General Electric's (PGE) 1992 Integrated Resource Plan.

IRP has a long way to go in the area of uncertainty. Many integrated resource plans do not address the issue of uncertainty at all. In this age of electric industry restructuring, every effort is needed to manage uncertainty. The ideas presented in this report should be valuable in this respect. However, additional research to illustrate the applications of these methods would also be useful. Lastly, PUCs often use the results of IRP processes as inputs for rate proceedings and often push for point estimates instead of ranges. It remains to be seen whether PUCs will present an institutional barrier to the use of imprecise probability.

LIST OF ACRONYMS

BPA	Bonneville Power Administration
CCDF	Classical cumulative distribution function
CDF	Cumulative distribution function
CF	Capacity factor
DSM	Demand-side management
DST	Dempster-Shafer Theory
IPP	Independent power producers
IRP	Integrated resource planning
LCDF	Lower cumulative distribution function
NPPC	Northwest Power Planning Council
NPV	Net present value
NRC	Nuclear Regulatory Commission
O&M	Operating and maintenance
PGE	Portland General Electric
PP	Purchased Power
PUC	Public Utility Commission
UCDF	Upper cumulative distribution function

1. INTRODUCTION

This report is about uncertainty and integrated resource planning (IRP). IRP is a planning and decision making process employed by electric utilities and public utility commissions (PUCs) to ensure that utilities provide cost effective and environmentally sound energy services to their customers. The key functions of IRP are to: (1) forecast electricity demand; (2) identify a broad range of energy resources potentially acquirable by the utility; (3) evaluate and rank order each resource according to predetermined criteria (e.g., costs and benefits); and (4) choose the best set of resources to meet forecast energy demand and other criteria. Given this flexible framework, IRP is notable for its ability to consider not only traditional electricity supply options but also demand-side management (DSM) options.

Uncertainty pervades utility resource planning. For example, integrated resource plans are based upon forecasts of energy prices and electricity demand. These forecasts entail significant uncertainties. Uncertainties also arise in estimating costs for: building new power plants adhering to environmental regulations; and purchasing power from independent power producers (IPPs). Integrated resource plans should explicitly represent and address uncertainty (Hirst 1992).

This report is written for the broad spectrum of people who are involved in IRP, including: PUC commissioners and staff; utility executives and managers; utility IRP analysts; and experts in particular areas such as oil price forecasting, environmental protection, renewables, demand-side management, and power plant construction. The purpose of the report is to help create a shared understanding of the concept of uncertainty. This is an important goal because all the parties involved in the IRP process deal with uncertainty in one form or another.

For example, IRP analysts typically interview experts to acquire their judgments about the values of key variables (e.g., future fuel prices). The experts need to communicate very clearly the reasons for uncertainty in their quantitative estimates of the magnitude of the uncertainty. The IRP analysts need to understand the qualitative aspects of uncertainty because it may be possible and desirable to attempt to reduce the uncertainty in a key variable. The analysts need quantitative measures of uncertainty as inputs into the resource selection process. Miscommunication about uncertainty between the experts and analysts will greatly reduce the quality of integrated resource plans.

At various points in the IRP process, IRP analysts must communicate the results of the analyses, including levels of uncertainty associated with the results, to utility executives and PUC commissioners and staff. The latter two groups may not have a great deal of expertise about uncertainty methods and may also be predisposed to expecting the analysts to present unambiguous results that will make their decision making easy. Indeed, because PUCs often use IRP results as inputs into rate proceedings, there is a strong tendency to ignore uncertainties in quantitative estimates. Because it is not the analysts' responsibility to make decisions and because it would be unethical to "cook" the results simply to make decision making and rate setting easier for these two groups, it would greatly benefit the latter two groups to have more than an intuitive grasp of uncertainty. Understanding qualitative reasons for uncertainty and what can and cannot be done to reduce uncertainty in key IRP variables would be a great help in reaching informed decisions. Also, having the ability to understand quantitative estimates of uncertainty would be helpful because this is the language of the methods that analysts use to assist the decision making process.

It is our belief that a sound understanding of uncertainty is an important prerequisite for sound decision making under uncertainty (another is a strong understanding of decision methods). It is also our belief that there are significant benefits to the IRP community in investigating the concept of uncertainty in some depth in a report such as this, not only to repackage existing ideas on the topic in new ways but also to communicate new ideas about uncertainty that are tailored explicitly for analysts and decision makers, instead of for statisticians and philosophers. Thus, this report focuses solely on the concept of uncertainty, leaving many other important issues for future consideration (e.g., decision making under uncertainty, uncertainty elicitation, risk communication and perception, and risk management).

Section 2 contains a qualitative discussion of the reasons for uncertainty in key IRP variables. This section provides the basic language for people to use when discussing uncertainty. The qualitative framework presented will allow all those involved in IRP to consider what, if anything, can and should be done to reduce uncertainty in the key variables.

Understanding reasons for uncertainty provides a sound conceptual framework for Section 3, which addresses how to represent uncertainty quantitatively. Two representations are presented, precise probability (i.e., classical probability) and imprecise probability (i.e., lower and upper probabilities). The topic of how to represent probability quantitatively has been very controversial for many years. From our perspective, one of the reasons for the controversy is that many practitioners have found classical probability to be limiting at best and unusable at worst. We believe that imprecise probability, which is a generalization of precise probability, offers flexibility to the practitioner who needs mathematical rigor and deals with wide variations of uncertainty in key variables. This section also contains discussions about how to calculate expected values and conduct sensitivity analysis using imprecise probability.

An often neglected subject pertains to how one can manipulate quantitative representations of uncertainty. Section 4 presents a qualitative discussion of three methods that IRP analysts can use in their endeavors and that others involved with IRP should know about: consensus, combination, and conditionalization. Basically, they are methods that one can use to merge opinions from various experts, synthesize pieces of evidence that shed light on the value of a variable, and revise estimates in light of new information, respectively.

This paper contains five formal examples that illustrate the concepts and methods discussed herein. All of the examples revolve around a central problem reported in Portland General Electric's (PGE's) 1992 Integrated Resource Plan. The problem faced by this utility was whether to shut down its Trojan nuclear plant immediately (in 1993), by 1996, or maintain operations until 2011. The integrated resource plan acknowledges that uncertainty was a major factor in the decision making process. In each of the examples, we contrast methods that were described in the integrated resource plan for dealing with uncertainty with methods proposed herein. This integrated resource plan was chosen because of its excellent quality and use of uncertainty methods. Comments on the integrated resource plan in this report are made to illustrate our methods, which are new to the IRP field, and are not made in any way to criticize PGE's fine analysis. In addition, because we did not have at our disposal PGE's models and data, our presentation and interpretation of PGE's work should not be considered authoritative. IRP is a continuous process and much work gets done that does not appear in publicly available reports. Thus, our presentation of the analytical problems tackled by PGE should be viewed as illustrations to facilitate the discussions of this report.

For the most part, this paper treats its subject qualitatively. Except for portions of Section 3, mathematics has been kept to a minimum even though sophisticated mathematics underlies the qualitative discussions of Sections 3 and 4. In addition, the report does not address many important but basically academic questions and issues on the topic. However, to address several technical questions about our ideas and methods that were provided by reviewers, we have included a short appendix. In addition, a companion piece is available. The interested reader is urged to see Tonn and Wagner (1995) for in-depth, mathematical presentations on these topics.

2. UNDERSTANDING UNCERTAINTY

This section addresses three categories of reasons underlying uncertainty in estimates used in IRP.

2.1 Irreducible Uncertainty

There are fundamental limits to our knowledge of the world. This is not a new idea. Heisenberg's famous Uncertainty Principle states that it is impossible to know both the velocity and location of a sub-atomic particle at a particular point in time. More recently, developments in chaos theory have shown that it is virtually impossible to predict the future states of highly complex, non-linear systems because of the extreme sensitivity of models of these systems to initial conditions. Thus, uncertainty is destined to always be a factor in science.

Unfortunately, uncertainty is also an unavoidable fact of life for those engaged in IRP. Numerous pieces of information central to IRP are fated to always be uncertain because they are inherently uncertain. As much as we may wish otherwise, *irreducible uncertainty* cannot be overcome through more work, time, and money. The best approach to handling irreducible uncertainty is recognizing one's limited knowledge in analyzing decisions. The worst approach is wishing away uncertainty and making decisions as though one possesses certain information.

In evaluating irreducible uncertainty in an estimate, three important questions should be asked (Tonn and Schaffhauser 1992):

(1) What does the estimate represent? (e.g., a long-range forecast, the status of the current world, an aspect of the past, a logical prediction);

(2) What kind of system is the estimate a part of? (e.g., a large-scale/disorderly system, a chaotic system, a small scale/orderly system, a logical system); and

(3) What is known about the values that the estimate can have? (e.g., is the range of values poorly understood or well understood).

Estimates that represent long-range forecasts of large-scale, disorderly or chaotic systems entail a great deal of irreducible uncertainty. Oil price forecasts fall into this category. Estimates that represent an aspect of the past are much less problematic, everything else being equal. Estimates of energy savings attributable to DSM programs are a good example of this case. Logical deductions from sound axiomatic systems entail no irreducible uncertainty, but unfortunately, IRP estimates are not typically derivable by logical deduction.

2.2 Operational Uncertainty

Unlike irreducible uncertainty, *operational uncertainty* can be overcome with more work, time, and money. Operational uncertainty is caused by numerous factors that fall into the category "The estimate would have been better if we had only done...." In evaluating operational uncertainty about an estimate, four important questions should be asked (Funtowitz and Ravitz 1990):

(1) How sound is the theory underlying the estimate? (e.g., ad hoc reasoning versus sound engineering principals);

(2) How much data were collected and used versus how much were required? (e.g., billing histories for 100 households were collected and used versus the need for data from 1000 households);

(3) Are the data high in quality? (e.g., billing histories had numerous missing months versus billing histories with no gaps); and

(4) How reasonable are the methods used to obtain the estimate? (e.g., the billing histories were not weather adjusted but needed to be versus the billing histories were properly weather adjusted in order to estimate energy savings due to residential retrofit).

It is often possible to collect more and better data, use better methods, and develop more theoretically sound approaches for solving problems. The important question is whether the benefits of pursuing any of these activities is worth the cost. This question is revisited in Section 2.4.

2.3 Application Uncertainty

How and for what purposes estimates are used may also cause uncertainty. For example, using estimates of energy saved due to residential retrofits in the winter in the coastal Pacific Northwest, which are characterized by low levels of irreducible and operational uncertainty, to estimate savings for a similar program in Wisconsin is not entirely appropriate because of potential differences in weather, housing construction, occupant life-styles, and stock of end-use technologies, among other factors. If these were the only estimates available, then one would need to carefully reflect on how much uncertainty is caused in the estimates due to their use in another context.

This type of uncertainty is labeled *application uncertainty*. In evaluating application uncertainty in an estimate, four important questions should be asked:

(1) Is the estimate informative for your context? (e.g., the estimate is too broad to be of much value to the analysis versus the estimate is tight and focused);

(2) Will the estimate's informativeness for your context decline over time? (e.g., the estimate for savings attributable to a DSM program from last year may not be a valid estimate for energy savings five years from now);

(3) Has the estimate been generalized from a more narrow study?

(4) Does the estimate measure what is needed for your context? (e.g., the estimate is a poor proxy for the variable that is really required versus the estimate is exactly what is needed).

In general, estimates that are informative, time insensitive, not generalized, and measure what is required will be more valuable for, and cause less uncertainty in, IRP than estimates that are uninformative, time sensitive, grossly generalized, and poor proxies for measures that are required. Similar to operational uncertainty, application uncertainty can be overcome by spending time and money to produce more informative, timely, and context relevant estimates.

2.4 Overall Uncertainty in Estimates

As depicted in Figure 2.1, all three types of uncertainty combine to cause an overall level of uncertainty in an estimate. IRP analysts desire perfect knowledge of all the values for all the variables in their analyses. As discussed above, true knowledge is rare. At best, one can hope to overcome all sources of operational and application uncertainty, which may still leave the analysts with bounds on their variables due to irreducible uncertainty. In most cases, operational and application uncertainty will also be present, which act to broaden the bounds of estimates even further.

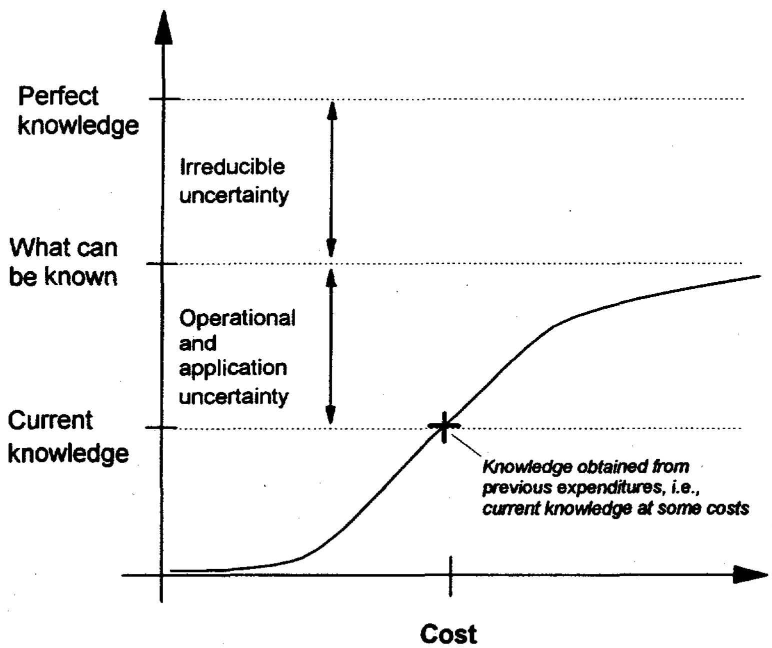

Fig. 2.1. Factors that underlie uncertainty in estimates.

By qualitatively assessing overall levels of and reasons for uncertainty in key IRP variables, one can begin to appreciate how best to expend scarce resources to reduce overall uncertainty in the IRP. The challenge for IRP analysts and decision makers is to assess the trade-offs associated with having more information for some cost, as depicted in Figure 2.2. At some point, increasing costs for overcoming operational and application uncertainty will overwhelm any incremental value of the additional information to the utility. Also, irreducible uncertainty places a ceiling on what can be known in any case. Analysis resources will be well spent on activities that can reduce the most operational and application uncertainty in estimates that are most important to IRP decisions.

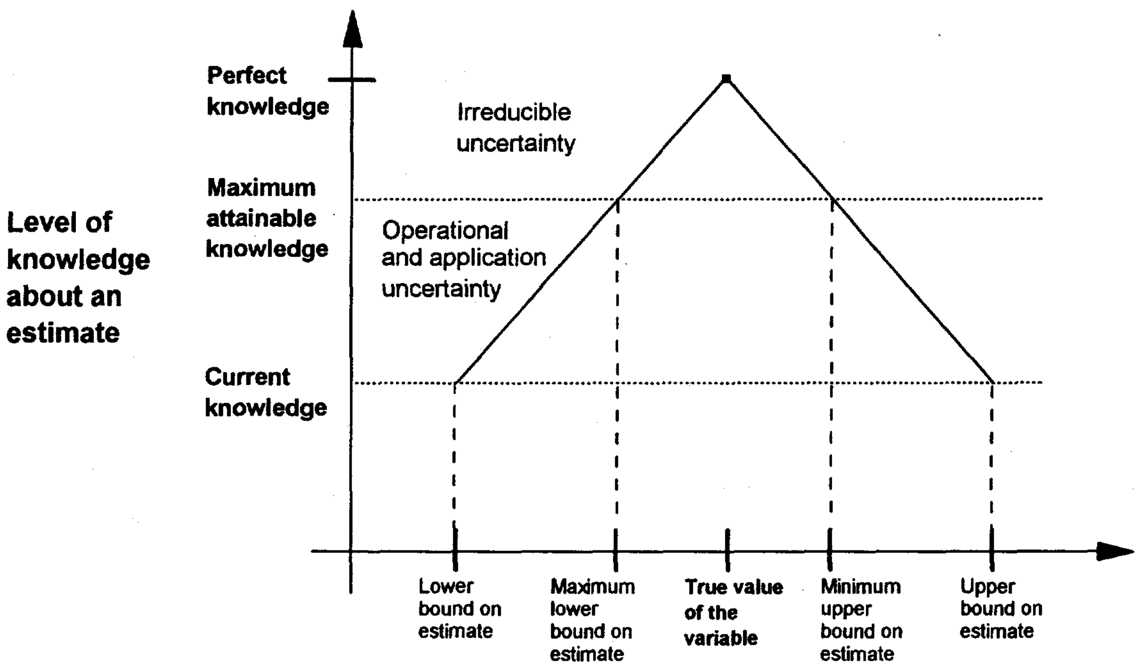

Fig. 2.2. Estimated value of the variable.

2.5 Example 1. Irreducible, Operational, and Application Uncertainty in PGE's IRP

The PGE 1992 integrated resource plan encompasses variables that exhibit different magnitudes of irreducible, operational, and application uncertainty.

As briefly mentioned in Section 1, the major question considered in the PGE IRP was whether to shut down the Trojan nuclear plant by 1993, by 1996 or operate the plant until 2011. Shutdown would require PGE to acquire replacement resources. The Trojan plant is a pressurized light water reactor that was 16 years old at the time of the integrated resource plan. To address this question, estimates for numerous key variables were required out to the year 2011 (e.g., average annual capital additions, capacity factor for steam generator replacement). As explained in 2.1, each forecasted estimate entails some degree of irreducible uncertainty commensurate with how far into the future forecasting is done. Thus, all these estimates entail some degree of irreducible uncertainty.

With respect to average annual capital additions, PGE's integrated resource plan states that "...future regulatory requirements are the major driver of future capital additions." Changes in regulatory requirements are hard to predict due mainly to the problem of identifying relationships between many possible political influences and their regulatory results. Thus, irreducible uncertainty emanates from the large-scale, disorderly (from a systems analysis perspective) system of political influences on regulations. The only factor that minimizes the impact of irreducible uncertainty in this estimate upon PGE's decision is that the range of possible values for annual capital additions is fairly well understood based on past experience with Trojan and other plants as regulatory requirements have varied.

Operational uncertainty characterizes the reported estimate of the costs for replacement power in the event that Trojan is shutdown immediately. PGE's integrated resource plan states that the earliest PGE could acquire ownership or contractual rights to new resources is 1996. Thus, PGE assumed an equal mix of the following three purchasable resources to replace Trojan's output: (1) purchased power from Bonneville Power Administration (BPA); (2) block firm power from other utilities' existing resources; and (3) short-term secondary power market (nonfirm) purchases or operation of a peaking unit, whichever is cheaper. PGE's integrated resource plan states that "We believe that secondary power is likely to be at a premium before new replacement resources come on line in 1996. The secondary power prices assumed for purchases to replace Trojan from 1993 to 1996 include a 5 mills/kWh premium." Analysts estimated replacement power costs following immediate shutdown would average $41/MWh.

This estimate is not totally accurate because it does not include an in-depth analysis of the cost and availability of secondary power, thus resulting in operational uncertainty. Additional work and time could have been spent in analyzing the cost and availability of this third resource to reduce the operational uncertainty of the entire estimate. As it turns out, in an update to PGE's integrated resource plan, further analysis of secondary power was indeed done and resulted in a best estimate of a range of replacement power costs of between $32 and $39/MWh, with $42/MWh being used to "...test the outcome under extreme conditions. PGE believes this case has a low probability." The lower estimate in the update was probably due to the reduction in uncertainty at the time of the update because Trojan had been shut down and the effect this had on the energy market could be observed. From our point of view, however, PGE could have attempted to place a bound on the original $41/MWh estimate, noting operational uncertainty due to the lack of analysis of secondary power purchases. It seems that such a bound should have at least captured part of the range used in the update.[1]

---

[1] PGE analysts did in fact place a bound on the $41/MWh estimate using a probability analysis. This range included the $31/MWh to $39/MWh range used in the update. However, the probability analysis range was not reported in the analysis documented in the integrated resource plan.

Estimating Trojan's annual fixed operating and maintenance (O&M) costs provides a good example of application uncertainty. PGE's integrated resource plan states that "...the most significant cause for increasing O&M at nuclear plants has been compliance with Nuclear Regulatory Commission (NRC) regulations" and that the major source of uncertainty in these estimates is "...the fact that the data set of all United States nuclear plants includes no comparable plants in the second half of their license lives." Therefore, the integrated resource plan relies on historical data from other kinds of nuclear plants and from the Trojan plant in the first half of its life, as well as on expert judgment from plant personnel, U.S. DOE, a consulting firm, and BPA. The application uncertainty arises because the estimates in the integrated resource plan:

(1) are generalized from studies of kinds of nuclear plants and for regulatory actions of the past which are subject to change;

(2) will decline in validity as more experience is gained about the Trojan plant and comparable plants in the second half of their license lives; and

(3) are less informative for comparing the options Operate Until 2011 and 1996 Phaseout than for comparing 1996 Phaseout to Immediate Shutdown because the ranges of values from all who projected O&M costs yield estimates within 30% of each other from 1993 to 1996, but only within 60% by 2010.

Comparing the estimates of Trojan's annual capital additions to replacement power costs provides a good example of how understanding uncertainty can be of great value in determining where to reduce overall uncertainty in an integrated resource plan. For annual capital additions, most of the uncertainty is irreducible, due to the nature of forecasts and the characteristics of the system that generates regulations. PGE received and reviewed estimates from three outside sources in addition to the projections of its own plant personnel. It appears that further study of this cost would not have been significantly helpful to PGE due to the irreducible uncertainty in the estimate. In other words, the studies reduced operational uncertainty to a minimum. In contrast, the additional study of the replacement power costs was helpful because the second estimate was able to overcome significant operational uncertainty.

In conclusion, framing the causes of uncertainty according to the concepts presented above may have led PGE executives, PUC officials, and stakeholders to require additional study of replacement power costs, and perhaps less study to project Trojan's post-1996 costs. This is because they would have realized that post-1996 costs are plagued with a great deal of irreducible uncertainty that no amount of study can reduce. If some of the resources spent to project costs in the post-1996 period had instead been used to analyze the cost of replacement power for the pre-1996 period, more useful information may have resulted.

2.6 Summary

This section presents a qualitative framework to understand reasons for uncertainty in quantitative estimates used in IRP. IRP analysts can use the framework to interrogate experts about uncertainty in their subjective and model derived estimates, and the experts should be prepared to answer such questions. PUC commissioners and staff, and utility executives and managers can also use the framework to interrogate IRP analysts to gain an understanding about the amount of uncertainty in the information they will use to make decisions. With this understanding, decisions can be made to put off decisions until more is known, if operational and application uncertainties can be overcome, and decisions can be analyzed in a way that accurately represents the uncertainty.

3. REPRESENTING UNCERTAINTY

3.1 A Brief History

People have been struggling with the concept of uncertainty for hundreds, if not thousands, of years. Over time, how people have thought about the topic has changed and become more sophisticated. The earliest thinking encompassed the concepts of chance and probability, which were seen as being unquestionably distinct; the former was associated with games whereas the latter was associated with matters of opinion (Hacking 1975). For example, an esteemed religious authority could argue that a proposition that is "probable is impossible," meaning that the proposition has a favored opinion but cannot be true.

During the 1600s, probability began to acquire its current connotation. The earliest reference is from Hobbes' *Humane Nature*, published in 1650, where he says that "...if the signs hit 20 times for one missing, a man may lay a wager of twenty to one of the event; but may not conclude it for a truth." The book *Port Royal Logic* (1662) contains the first use of the word "probability" in a quantitative sense. Jacob Bernoulli, in his manuscript *Ars Conjectandi* (1713), laid the groundwork for modern probability and statistics (Shafer 1978).

In this century, the concept of uncertainty has come to be for many completely synonymous with classical probability. That is, it is argued that when one needs to discuss and represent uncertainty, one ought to base such discussions and representations on the axiomatic foundation of classical probability (see 3.2). It is also argued by many that probability must be treated using the frequentistic viewpoint, as exemplified by the above quote from Hobbes. These assertions are generally noncontroversial among statisticians and physical scientists.

However, the second assertion has come under attack from people who wish to use classical probability to make decisions. IRP is not unusual in a decision making sense in that many resource decisions cannot be made on the basis of statistical information. Indeed, many decisions are one-of-a-kind and are not amenable to statistical analysis (Hirst and Schweitzer 1990). A debate has raged between *objectivists*, who argue for a completely frequentistic implementation of probability, and *subjectivists*, who argue that probabilities can also be developed through rational introspection.[2] The subjective viewpoint is clearly dominant within the decision sciences and is most inclusive for application to IRP because it can encompass both objective and subjective probability.

In recent decades, many practitioners, especially in the field of artificial intelligence, have come also to question the first assertion. Specifically, the axioms of classical probability, whether seen from an objective or subjective viewpoint, it is argued, are too limiting. For example, those

[2]Ramsey (1931), DeFinneti (1964), and Savage (1964) present frameworks for such rational thought exercises. It is our observation that this framework is difficult for people who are not decision theorists to use these frameworks. In practice, subjective probabilities are usually directly elicited from experts without the aid of these exercises. See Wallsten and Budescu (1983).

developing expert systems found that it is difficult to represent and synthesize pieces of diagnostic evidence using classical probability, even subjectively interpreted. Thus, research has focused on generalizing probability and on developing completely new methods for representing uncertainty. These new approaches include: fuzzy sets (Zadeh 1978); certainty factors (Buchanan and Shortliffe 1984); and Dempster-Shafer Theory (Shafer 1976).

The development of new approaches to uncertainty underscores two very important points that IRP practitioners need to understand. First, uncertainty is a concept that is broader than the language and methods of probability. As already illustrated in Section 2, one can engage in far ranging discussions about uncertainty in estimates without having to couch one's terminology within the paradigm of probability. One also need not be an expert in the mathematics of probability to discuss uncertainty in estimates that are key to making IRP decisions.

Second, it is also important to realize that methods for representing uncertainty are continuing to evolve and that the evolution is being driven not by the theoretical concerns of statisticians and physical scientists but by the practical concerns of decision makers and those who build decision support and artificial intelligence systems. Thus, if as a practitioner or decision maker you have had nagging doubts about the applicability of classical probability theory in your context, you are not alone. We believe that imprecise probability (see 3.2), which is a generalization of classical probability and has some roots in Dempster-Shafer Theory of the 1960s and 1970s and in Bernoulli's original work over 250 years ago, provides to IRP practitioners a significantly better framework within which to handle uncertainty.

3.2 Precise and Imprecise Probability

Axiomatic accounts of classical or precise probability theory always state, for the probability of event A, P(A), that $0 \leq P(A) \leq 1.0$, for all events A, with $P(\Omega)=1.0$, where Ω is the set of all possible events. In addition, additivity is always postulated and countable additivity is often postulated (always, among mathematicians).[3] As a consequence of these postulates, one always has $P(A) + P(\bar{A}) = 1.0$, where $\bar{A}$ means <u>not</u> A or the complement of A (i.e., the probability that the event that will become true is not event A).

One can calculate an objective precise probability from data or subjectively develop a precise probability. For example, if one knows one has a fair coin, one could calculate the probability of tossing heads as being 0.50. To subjectively estimate the probability of an event, such as the recurrence of an oil crisis, one could cognitively construct an estimate directly, or be assisted by tools (e.g., probability wheels) and betting frameworks. For continuous variables, probability density functions (e.g., normal curves) and cumulative distribution functions (CDF) are used to represent uncertainty in estimates.

---

[3]Additivity is expressed as $P(A \cap B) = \emptyset \rightarrow P(A \cup B) = P(A) + P(B)$. Countable additivity is expressed as $P(A_1 \cup A_2 \cup ...) = P(A_1) + P(A_2) + ...$, for every infinite sequence $A_1, A_2, ...$ of pairwise disjoint events.

It is our belief that precise probability is not expressive enough to capture large uncertainties associated with IRP. Basically, practitioners are limited to one probability estimate for a discrete event and one density function/cumulative distribution function to express uncertainty about a key IRP variable. Metaphorically, one is limited to one unit of measurement, to one gauge to express uncertainty. Imagine physicians being limited to one test to diagnose a disease, painters being limited to one color to express emotion, or consumers being given only one feature (e.g., price) with which to judge a product. As much as we sympathize with those who wish to keep the world as parsimonious as possible, there are cases where over simplification results in a loss of information and with respect to uncertainty in IRP, this is such a case.

A natural way to relax the demand for a single number expressing the probability of an event A is to allow assessment of uncertainty by an interval $[\underline{P}(A), \overline{P}(A)]$, where $0 \leq \underline{P} \leq \overline{P} \leq 1.0$. The two numbers $\underline{P}(A)$ and $\overline{P}(A)$, called respectively, the lower and upper probabilities of A, are chosen so that one is, given present evidence, reasonably sure that the probability of A is neither less than $\underline{P}(A)$ nor greater than $\overline{P}(A)$. Thus, for discrete events two numbers are used instead of one and for continuous variables, two CDFs, a lower and an upper, are used instead of one.

Figure 3.1 compares the expressiveness of precise and imprecise probability in four cases. Each icon within the figure represents a gauge or measurement device, where each gauge ranges from 0 to 1.0, i.e., empty to completely shaded, respectively. An empty gauge expresses a probability of zero, a completely full gauge expresses a probability of one. To assist this discussion,

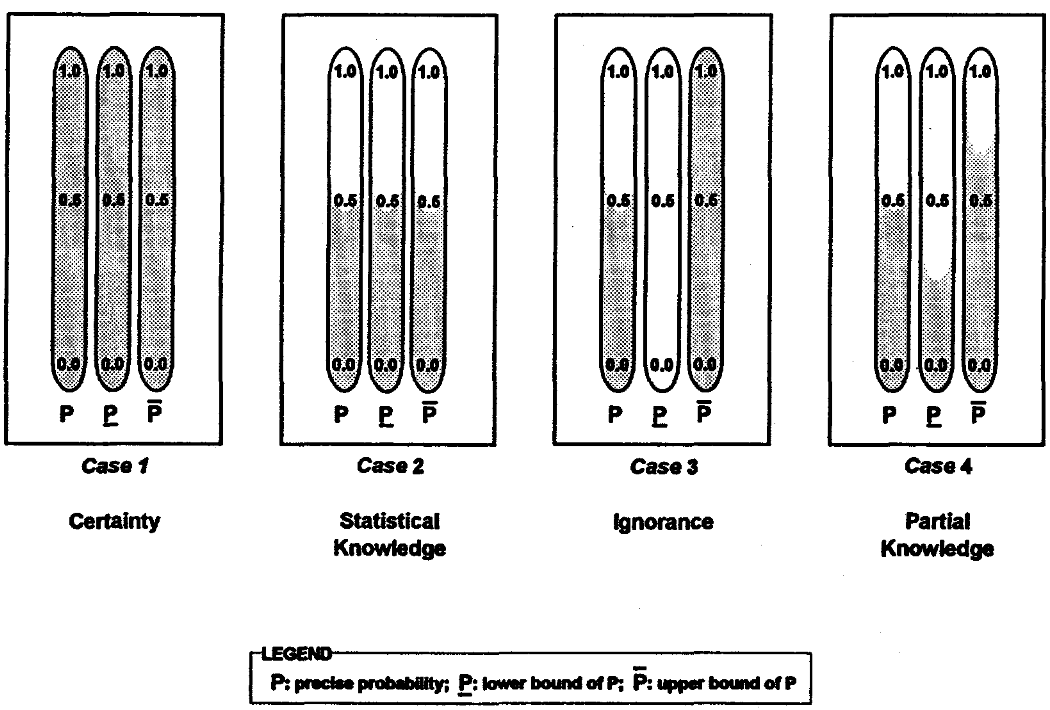

Fig. 3.1. Expressiveness of precise and imprecise probability.

imagine a container of 100 balls, which could be red or blue. The task is to estimate the probability of selecting a red ball from the container, P(r).

In the first case, the IRP analyst has examined all the balls in the container and knows for certain that they are all red. In this case, the analyst has "certain" knowledge that a red ball will be selected. As expected, the precise probability gauge, P, is completely shaded (i.e., P(r)= 1.0). The two imprecise probability gauges are also completely shaded, which indicates that when $\underline{P}(r) = \overline{P}(r) = 1.0$, certainty is achieved. Certainty would also be achieved when all the gauges are empty.

In the second case, the analyst has examined all the balls in the container and knows for certain that 50 balls are red and 50 balls are blue. The straightforward frequentistic interpretation is that P(r)= 0.5. It is also appropriate to assign $\underline{P}(r) = \overline{P}(r) = 0.5$. Thus, the gauges are half shaded in all three cases. It should be pointed out that at this stage, there is no difference between precise and imprecise probability. Certain knowledge and strong statistical knowledge are essentially represented by the two approaches in the same way, showing that the two approaches are consistent when one has strong evidence.

In the third case, the analyst has no knowledge about the balls in the container, other than that there are 100 of them. There could be 100 red balls and no blue balls or vice versa. There could be any other combination of red and blue balls that add to 100. In these cases of total ignorance, adherents of the classical probability paradigm apply the Principle of Insufficient Reason, which basically stipulates that probability mass be spread out evenly among the competing events. In our example, the uniform probability would be chosen, which in this case means P(r) = 0.5. Thus, when limited to only one number, a high degree of statistical knowledge is represented the same way as total ignorance!

Here is where the expressiveness of imprecise probability comes into play, when there is a lack of knowledge. Using imprecise probabilities, the lower probability gauge is empty and the upper probability gauge is full. In other words, ignorance is represented as $\underline{P}(r) = 0.0$ and $\overline{P}(r) = 1.0$.

In the fourth case, the analyst is able to see 80 of the 100 balls through the clear sides of the container, which are now uncovered, and now knows that the number of red balls is at least 40 and the number of blue balls is at least 40, meaning that the number of red balls is between 40 and 60. In this case, $\underline{P}(r) = 0.4$ and $\overline{P}(r) = 0.6$. From the figure, one can easily see how increased knowledge is straightforwardly exhibited in a narrowing of the range between the lower and upper probability gauges. Unfortunately, such is not the case with classical probability. In place of the Principle of Insufficient Reason, a maximum entropy heuristic would be applied to estimate the precise probabilities which maximize the famous entropy formula proposed by Shannon.[4] Applying such

[4] See Shannon (1948) for the original presentation of the entropy formula and Jaynes (1957) for the seminal discussion on maximizing entropy. From our perspective, the application of either the Principle of Insufficient Reasoning and maximizing entropy, also known as MAXENT, adds information to the problem that is not based on any evidence. These heuristics are incompatible with good decision making and with the whole notion of imprecise probability.

a heuristic would again produce P(r) = 0.5. Thus, precise probability does not allow one to express how much knowledge is being brought to bear on a decision or how knowledge may change over time with the addition of new information.

This is an exceedingly important observation because it indicates that classical probability is not well suited to employ the qualitative concepts discussed in Section 2.0. Using imprecise probability, one can evaluate how efforts to overcome operational and application uncertainty may reduce the range between lower and upper probabilities, and thus how decisions might change. With only one gauge, it is much more difficult to imagine how one could capture the benefits of such activities if one where limited to precise probability.[5]

Figure 3.2 presents a similar story for continuous variables. In the middle of the figure is a classical cumulative distribution function, (CCDF). Flanking the CCDF are an upper cumulative distribution function ($UCDF_2$) and a lower cumulative distribution function ($LCDF_2$). An expert providing these two functions would be communicating that there exists some lack of knowledge about the true CDF for this variable. A very large degree of ignorance is captured by $UCDF_1$ and $LCDF_1$.

Whether one is working with discrete or continuous variables, the amount of uncertainty expressed quantitatively needs to be consistent with the degree of uncertainty revealed through the qualitative framework presented in 2.0. For example, one would expect that, due to large irreducible uncertainties, experts would rather provide lower and upper CDFs for future oil prices rather than one CDF. On the other hand, costs of delivered electricity for an existing plant may be so well studied that a precise CDF may be sufficient and defensible (provided the plant were supplied by a long-term fuel contract, or were fueled by a predictable renewable source).

3.3 Example 2. Using Precise and Imprecise Probabilities to Represent Uncertainty in IRP Variables

Consider once again the Trojan decision example. To address uncertainty in many numerical estimates, such as those associated with operations and maintenance costs, a simulation approach was implemented. For each variable in the simulation model, a precise probability distribution functional form was chosen, such as a normal curve. To specify each function, the IRP analyst

---

[5]Imprecise probability has been criticized because lower and upper probabilities are not additive and that the loss of additivity means that people cannot make rational decisions, as defined within the context of Savage's axioms. We have three responses. First, we believe that strict additivity is unnecessarily restrictive. As Walley (1991) shows, meaningful statistics (i.e., frequentistic-based lower and upper probabilities) can be calculated with a relaxed additivity axiom. Second, imprecise probability still retains axioms related to additivity. For example, one can impose superadditivity and subadditivity (see Wagner 1994). Thus, mathematical rigor is retained. Third, Schmeidler (1986) has generalized Savage's axioms for lower and upper probabilities. Fishburn (1988) expounds more on this point. Thus, imprecise probability is completely consistent with the goals of subjectivists and decision analysts.

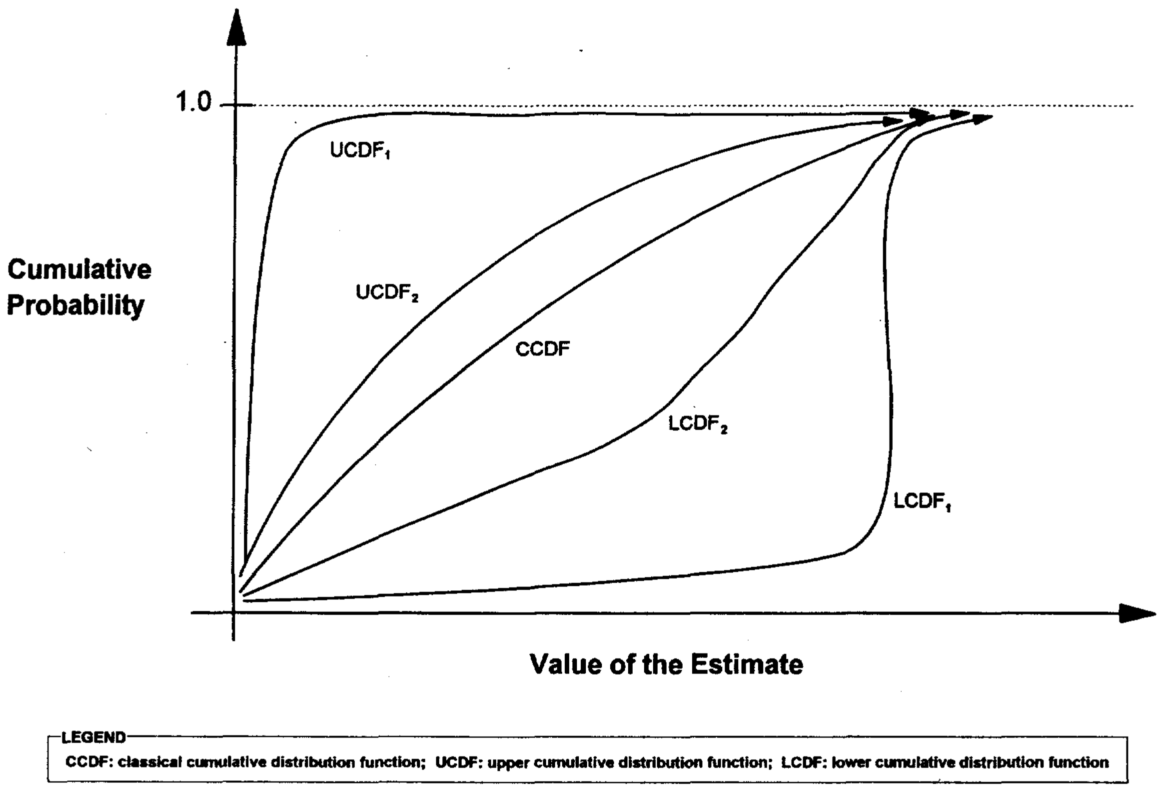

Fig. 3.2. Classical, lower, and upper cumulative distribution functions.

constructs a range of potential values for the variable and then judges with what precise probability the true value will lie in the range. The simulation code samples from the distributions to obtain values to input into the various cost calculations. Repeated sampling and calculation to create a distribution of cost outputs is commonly known as Monte Carlo analysis.

As noted in Example 1, all three types of uncertainty plague the estimates. Irreducible uncertainty is a factor because many of the estimates are long-range forecasts. Operational uncertainty is a factor because data are not available to support the choices of distribution functional form and standard deviations for many of the variables. Application uncertainty is a factor because operations and maintenance cost estimates for the current Trojan plant may not be valid descriptors of the future Trojan plant. The cumulation of these factors leads us to believe that imprecise probability could have better represented uncertainty in the simulation exercise.

We would recommend in this case, the following six-step process.

(1) Each of the questions presented in Section 2 should be asked for each of the key variables in the analysis.

(2) An overall qualitative understanding of the uncertainty in each key variable should be developed.

(3) Based on this qualitative understanding, LCDFs and UCDFs should be specified for each key variable. Without having data to support this step, the specification of these functions is admittedly subjective. Having answers to the questions of Section 2 should improve this step and analysts and experts are still free to use decision aides to assist elicitation. Referring to Figure 3.2, one would expect that LCDFs and UCDFs for near-term forecasts would resemble $LCDF_2$ and $UCDF_2$, whereas those for long-term forecasts would more resemble $LCDF_1$ and $UCDF_1$.

(4) The LCDFs and UCDFs can be input into the simulations to calculate lower and upper bounds on the outputs. Mathematically, they are not different from CDFs, so simulation codes would not have to be altered.

(5) At this point, expected values can be calculated, which is the subject of the next section, 3.4.

(6) Finally, sensitivity analysis can be performed, which is the subject of 3.6 and Example 4 in 3.7.

3.4 Calculating Expected Values Using Precise and Imprecise Probabilities

In many integrated resource plans, there is a requirement to calculate expected values for key decision variables. For example, it is often useful to calculate expected values for streams of costs and benefits associated with the acquisition of a resource. This section illustrates how to calculate expected values using precise and imprecise probability.

To begin, to calculate expected values for discrete variables using precise probability, one uses this formula

$$E(A) = \sum_{i=1}^{n} a_i \, P(A = a_i) \tag{1}$$

where E(A) is the expected value of variable A, a_i is one of n values that variable A can take on, and $P(A=a_i)$ is the precise probability that the variable A will take on the value a_i.

To calculate lower and upper expected values, one follows a two step process. First, the potential values that A can assume are arranged in ascending order and labeled as follows: $a_1 < a_2 < \ldots < a_n$. Second, the following equations are applied

$$\underline{E}(A) = a_1 + \sum_{i=2}^{n} (a_i - a_{i-1}) \, \underline{P}(A \geq a_i) \tag{2}$$

$$\overline{E}(A) = a_1 + \sum_{i=2}^{n} (a_i - a_{i-1}) \overline{P}(A \geq a_i) \qquad (3)$$

where, $\underline{E}(A)$ and $\overline{E}(A)$ are the lower and upper expected values of variable A, respectively; $(a_i - a_{i-1})$ is just the difference in values between two potential values of A, and $\underline{P}(A \geq a_i)$ and $\overline{P}(A \geq a_i)$ are the lower and upper probabilities that A equals or exceeds the value a_i. Technically, formulas (2) and (3) calculate what are called Choquet expected values for the lower and upper probabilities $\underline{P}$ and $\overline{P}$. There are extensions of (1), (2), and (3) for continuous variables. In addition, a method for using nonlinear optimization to calculate lower and upper expected values for a cascading set of variables, such as those found in energy supply and demand models, has been developed (Wagner 1994).

A quick inspection of (2) and (3) reveals that the lower and upper expected values will equal the expected value calculated by (1) when $\underline{P} = \overline{P}$ for all values of A. Thus, referring to Figure 3.1, assuming that choosing red and blue balls resulted in different monetary awards (e.g., $5 for a blue ball and $10 for a red ball), the calculated expected value for choosing a red ball will be the same using precise and imprecise probabilities in the first two cases (i.e., $10 in the first case and $7.50 in the second case). As expected, however, the expected value calculations are different in the final two cases. With respect to precise probabilities, the expected value holds at $7.50 regardless of how much information is at hand. With respect to imprecise probabilities, the range of expected values is $5 to $10 in the third case, and a narrower $7 to $8 in the fourth case. We argue that decision makers ought to be presented the bounds of expected values of key variables, no matter how large the bounds are, rather than unjustifiably specific values calculated using precise probabilities. If the bounds are not presented, then no distinction is made between cases where a lot is known and cases where less is known. Ignoring this distinction assumes more is known than what actually is. This false precision misrepresents the decision situation, and judgments based on a misrepresentation are flawed. (Appendix A, provides a more in-depth discussion of this issue.) The following example illustrates how one could make judgments when ranges of expected values are presented.

3.5 Example 3. Ranking Resources using Lower and Upper Expected Values

The goal of this example is to illustrate how upper and lower expected values can be used to rank order resource supply options. Table 3.1 presents a list of supply-side resource options PGE considered admissable in its integrated resource plan. (In addition to these so-called "discretionary" supply-side options, PGE selected approximately 570 Mwa {i.e., average megawatt output} of resources based on avoided cost comparisons. These so-called "base" resources include DSM, special contracts with industrial customers, anticipated PURPA purchases, and system efficiency improvements and upgrades.) This list includes all resources needed to replace Trojan in 1996.

Table 3.1 Supply side options considered by PGE.

Option description		$E(C)_{low}$	$E(C)_{up}$	$E(MWa)_{low}$	$E(MWa)_{up}$
1.	**CCCTs** (firm gas supply contracts)	246	283	563	597
2.	**Geothermal** (50% of electricity from flash cycle units, 50% from binary cycle units)	24	29	48	52
3.	**Pulverized Coal** (N.E. Oregon site, Montana, or Wyoming coal)	289	406	485	515
4.	**Wind Turbines** (Oregon, Washington, or Montana site)	23	29	47	53
5.	**SCCTs** (market-priced gas)	102	131	160	170
6.	**Combined Cycle Repowering** (new CTs and heat recovery steam generators [HRSGs] added to existing plant)[6]	38	45	97	103
7.	**Fuel Cells** (phosphoric acid)[7]	31	37	47	54
8.	**Coal Gasification CCCTs**	273	342	376	424
9.	**CCCTs** (market-priced gas)	110	151	243	258

The table includes, from left to right, lower expected costs (levelized real cost per year), upper expected costs, lower expected performance (average MW), and upper expected performance. (Note that lower and upper expected costs include lower and upper expected value for external costs,

[6]This option entails adding new CTs and HRSGs and using the existing steam turbine and other plant facilities. The existing CTs/HRSGs would be retained for meeting peak demands and to support combined-cycle operation if one of the new CT/HRSGs were unavailable.

[7]PGE screened out this resource because it "is still in the research and development stage, and not yet commercially available." However, given its relatively short lead time, estimated at three years, it could be deployed rather quickly if it did become commercially available. Thus, we feel it is a relevant resource to consider due to the fifteen-year planning period assumed for this example.

e.g., environmental externalities). These expected values can be assumed to have been constructed using equations (2) and (3) discussed in 3.4.[8]

There are numerous ways to choose resource supply options from the nine presented in Table 3.1. Let's assume that the lower expected resource need is 600 MWa and that the upper expected resource need is 1200MWa. Further, suppose that this resource need reflects projected demand over fifteen years. The most conservative approach would be to : (1) use the most risk averse cost and performance, which is calculated by dividing the upper expected cost $E(C)_{up}$ by the lower expected performance $E(MWa)_{low}$ times 8760 hours/yr; (2) rank order the options using this upper cost per MWh; and (3) choosing options such that the sum of the lower MWa supplied just exceeds upper expected resource need. If this approach were used, the least cost resources to be acquired, in order of ascending cost, would be as follows—6,1,2,4,9,7, and 5.

To balance this conservative approach, the most optimistic approach could also be explored, which would use the lower expected cost $E(C)_{low}$ combined with the upper expected performance $E(MWa)_{up}$ and choose options to just exceed the lower resource need. If this approach were used, options 6 and 1 would be indicated as the least cost resources to meet demand. The upper expected supply would be 700 MWa.

This example presents some interesting observations. Both the conservative and the optimistic approaches indicate options 6 and 1 as the resources of least cost. Furthermore, options 6 and 1 would meet lower expected resource needs even if they had lower expected performance. Their lower expected combined output is 660 MWa. Thus, the utility could first acquire options 6 and 1 in order to meet demand over the first part of the planning period and add additional resources respective of their lead times, while gaining additional information about demand, costs, and performance as time passes. Supposing that resources to meet demand greater than 660 to 700 MWa must be acquired soon, the utility might select options 2 and 4 which are ranked next in line by the conservative approach. The optimistic approach places option 9 ahead of options 2 and 4, but options 2 and 4 have the advantage of adding some diversity to a new resource portfolio that is all natural gas technologies at this point. The addition of options 2 and 4 would raise lower expected supply to 755 MWa and upper expected supply to 805 MWa. If project lead times and the timing of projected demand dictated the near-term acquisition of greater supply, then option 9 would be selected. It is the next resource in the conservative ranking and it is ranked third in the optimistic ranking. Adding option 9 provides 997 MWa by the conservative approach and 1063 by the optimistic. Finally, if near-term acquisition of even greater supply were needed, option 5 (160 MWa to 170 MWa) would be a good choice. Although option 7 supercedes option 5 in both the conservative and optimistic rankings, fuel cells would not be commercially available in the near

[8]The options and data are conceptually based upon real analyses contained in PGE's 1992 integrated resource plan. We stress, however, that the values in Table 1 are purely hypothetical. This example is similar to ORNL/CON-388, 3.6.4 Example 4, except that here we use E(C)low/E(MW)up for the optimistic viewpoint, and E(C)up/E(MW)low for conservative viewpoint.

term (see footnote 8). Furthermore, although option 3 supercedes option 5 by the optimistic ranking, it is much larger than needed (485 MWa to 515 MWa).

3.6 Sensitivity Analysis and Precise and Imprecise Probability

Many practitioners who employ precise probability in their analyses conduct sensitivity analysis to overcome perceived deficiencies in the validity or reliability of precise probabilities. For example, referring to Figure 3.1 again, one could conduct sensitivity analysis to examine a range of expected values associated with selecting a red ball in the third and fourth cases. To do this, the analyst may assume that the probability of choosing a red ball in the third case could be .6, or .35, or .1, or any other potential value. Using these values, a range of expected values could also be constructed using precise probabilities. Sensitivity analysis is conducted to investigate at what point a change in precise probabilities could lead to a change in a decision.

This type of sensitivity analysis is unnecessary when one uses imprecise probability. The lower and upper probabilities provide firm bounds on the probabilities and on the expected values. There should be no doubt about these bounds, given the evidence at hand and one's current knowledge of the situation. One could examine if decisions might change *in light of new knowledge*, if such knowledge could be gained, as provided by insights guided by the questions in 2.0 (i.e., if operational and application uncertainty can be overcome to some degree). This type of sensitivity analysis is quite different than the type described above.

We argue that for IRP, the second type of sensitivity analysis based on imprecise probabilities is superior to the first type based on precise probabilities. We offer the following five reasons.

One, traditional sensitivity analysis provides only a scattershot view of the entire range of potential expected values. Figure 3.3 presents a visual analogy of this problem. Let's assume there are three variables in the analysis. Using imprecise probabilities and knowledge about how the variables interact with each other, it is determined that the bounds that the variables could take on are represented by the shaded portion of the intersection of the three variables. Using lower and upper probabilities, lower and upper expected values can be calculated for each of the variables that in effect describe the bounds of the shaded area. Using precise probabilities, only a few points within the shaded area are examined and the true bounds of the expected values are never completely described. Imprecise probability provides a comprehensive view of the expected values.

Two, within the traditional sensitivity analysis approach, it is difficult to justify why certain precise probabilities are chosen over other equally justifiable probabilities (i.e., other probabilities that are not inconsistent with what is known). Again referring to Figure 3.3, what points one chooses in that space to investigate will significantly influence one's analysis. Since one can never be completely sure that one has chosen the best points for sensitivity analysis, we argue why submit oneself to the choice process in the first place. Imprecise probability eliminates the need for these arbitrary decisions. (See Appendix A, for more discussion of this issue.)

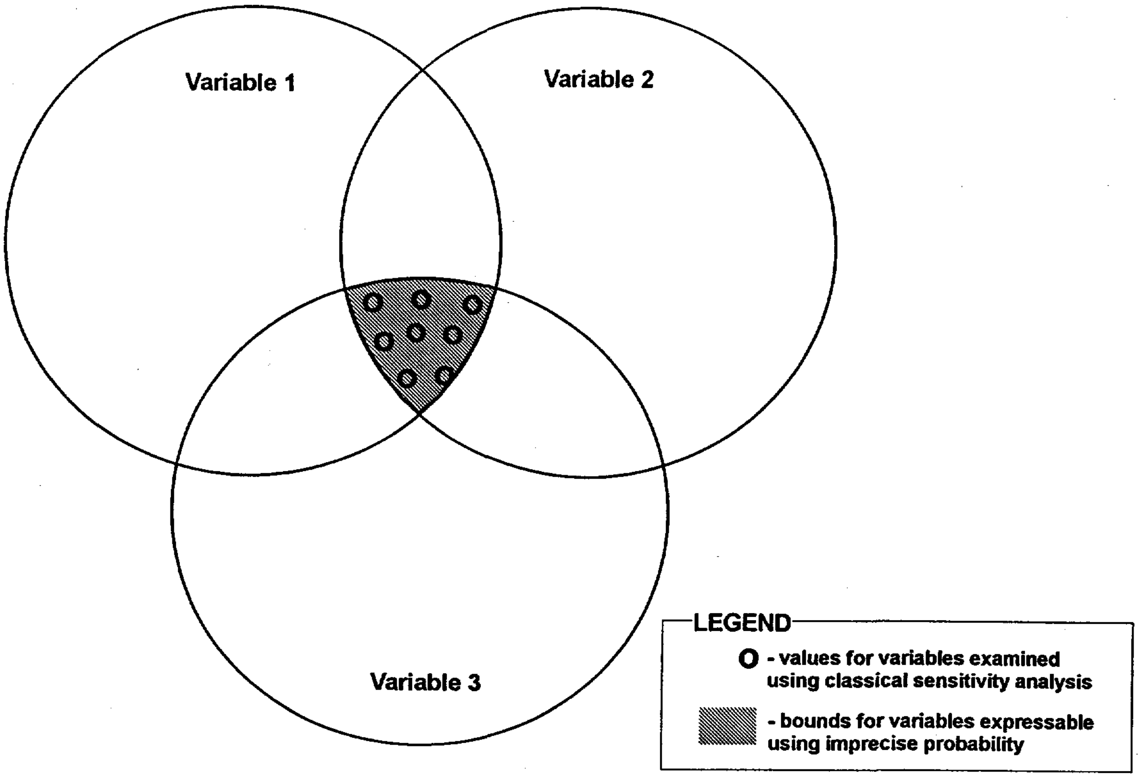

Fig. 3.3. Classical sensitivity analysis versus bounds on variables based on imprecise probability.

Three, for many contexts, it becomes a combinatorial problem to conduct traditional sensitivity analysis. For example, let's say that the analyst decides to try just four different precise probabilities for just four different variables. This results in 256 different sensitivity analysis results! We would be at a loss to decide what to do with so many results. With imprecise probability, one only has to deal with one set of results.

Four, traditional sensitivity analysis can lead to intractable overlapping expected values. If one really had 256 different results, it would be likely that no one decision alternative would dominate the others, (i.e., where the minimum expected value of one decision alternative would be greater than the maximum expected values of all the other potential decision alternatives). What does one do in this case? Conduct additional sensitivity analysis and produce even more conflict? It is true that imprecise probability will also result in overlapping expected values. The difference is that imprecise probability produces just one set of results, and decisions can be made using sound, well-defined, and prespecified decision rules (e.g., maxi-min, mini-max) whereas with respect to sensitivity analysis using precise probabilities, the decision set is arbitrarily defined, and grows as more values are changed. If the decision is not to be left ambiguous, decision rules can only be applied in an ad hoc fashion to finally close out the analysis. (Again, see Appendix A, for more discussion.)

Five, we object to conducting sensitivity analysis only with respect to its relative value of potentially impacting one decision at one point in time. In our view, statements of uncertainty about a variable represent statements of one's fundamental knowledge about the variable. Many other decisions could be based on this same variable and similar decisions could be made over time based on the same variable. Thus, from the perspective of a utility, it is much more preferable to examine the costs and benefits to all the utility's decisions of increasing knowledge about a variable (e.g., by narrowing the range of the lower and upper probabilities). The framework presented in Section 2 is an excellent guide for those types of analyses.

3.7 Example 4. Sensitivity Analysis in a Power Plant Shutdown Decision

This example contrasts sensitivity analysis based on precise probability versus imprecise probability. Again, this example concerns decisions pertaining to the operation of the Trojan Nuclear Plant. The three options PGE considered are listed below.

Continued Operation: Replace the steam generators and continue operation until the plant's license expires in 2010.

To continue to operate Trojan past 1996, the steam generators need to be replaced. 1996 replacement requires steam generators be ordered in the very near future.

1996 Phaseout: Operate Trojan until 1996, while acquiring replacement resources in the meantime.

PGE states that this option reflects the "shortest reasonable time in which we can acquire sufficient resources to replace Trojan's output." The resources to replace Trojan would be either new PGE resources (similar to resources included in Example 3 to illustrate lower and upper expected values), or resources expected to be available throughout the Northwest Power Planning Council (NPPC) region.

Immediate Shutdown: Immediately close the Trojan plant, but keep it in standby status for two years.

If PGE immediately closed the Trojan plant (by early 1993, given the period of time that the integrated resource plan was prepared) they would have to acquire 230 MWa of energy and 745 MW capacity in 1993. Replacement resources from 1993 to 1995 would be power and capacity purchased on the secondary market (similar to the resource purchases described in Example 1 to illustrate operational uncertainty). Replacement resources after 1995 would be the same as for 1996 Phaseout.

The desirability of the options depends on (1) the costs of operating Trojan and (2) the costs of resources that would replace Trojan.

To depict the uncertainty about the costs of operating Trojan, PGE developed precise probability distributions for the following variables with the simulation method described in Example 2:

* Trojan's annual O&M expenses and their escalation over time,

* Trojan's annual capital needs and their escalation over time, and

* Trojan's capacity factor.

The result of analysis with these precise probability distributions is an expected value and a probability distribution for the cost (in present value) of operating Trojan.

To depict the uncertainty about the costs of resources that would replace Trojan, a model of regional replacement power costs was used. (The NPPC did this modeling on behalf of PGE.) However, the model was only applied to replacement resources from 1996 on. The result is an expected value and a probability distribution for the cost of replacement resources from 1996 to 2010.

For now, let's confine the discussion to comparing continued operation to 1996 phaseout. The expected costs PGE came up with are $3.2 billion for continued operation and $2.86 billion for 1996 phaseout. Thus, 1996 phaseout has lower expected cost. However, have these expected costs accurately captured what is known and not known about each variable? The following statement by PGE answers this question.

> The expected values for Trojan inputs are not overly pessimistic. Certainly evidence existed that supported more pessimistic operating conditions. The expected values even reflect improvement over historical costs and capacity factors, which is consistent with some of the recent trends in the industry of stabilizing levels of nuclear regulatory activity. On the other hand, it is difficult to find hard evidence that would support more optimistic inputs. . .than the expected values. [p.4A.22]

Such discussions about the optimism or pessimism of expected values are an artifact of the precise probability approach. How can one "represent. . .[***the***] estimated frequency of occurrence" [p.4A.4] for the values of a variable such as the escalation in O&M costs from 1996 to 2010? Constructing a precise probability for such an uncertain variable must entail expressing a pessimism or optimism about its value. (Granted, with imprecise probability judgment is still brought to bear. However, one can represent all reasonable judgments, [i.e., all judgments based on the available evidence], in one imprecise probability function.)

As discussed above, sensitivity analysis is often performed to test how sensitive the expected values (and in turn, the decision choice) are to different input probabilities. Thereby, one can inject less or more pessimism or optimism and see how the options fare. However, as discussed above, how does one justify using certain precise probabilities instead of others? Perhaps to avoid this problem, as well as the combinatorial problems of assigning different precise probability distributions to variables, PGE performed a scenario analysis.

This analysis takes different sets of assumptions about conditions and costs (scenarios) and calculates the net present value (NPV) of continued operation and 1996 phaseout. Sensitivity analysis was performed by assuming different values for Trojan variables and recalculating costs. PGE took a "reference case" scenario that had values for Trojan variables very similar to the expected values calculated from the probabilistic analysis and tested how much the costs of the options changed if the values of these variables changed. The results of this analysis were expressed for each variable as exemplified below.

> *Fixed O&M Escalation Rate:* A 1% increase in the rate of escalation of fixed O&M would increase the NPV benefit of phase-out to $261 million. A 1% decrease would switch the result to favoring continued operation, but only by $31 million. A 1% increase or decrease in the escalation rate generally increases or decreases the NPV benefit of phaseout by between $100 and $200 million across cases [i.e., all scenarios considered].

PGE points out the drawbacks of such an approach.

> While interesting, these sensitivities provide little guidance because they do not convey any information about how likely a change from the reference case would be. In addition, each assumes holding all else equal. What they do allow is estimation of the effects of changing one or more of the inputs [provided it does not affect interrelations among them] if, for various reasons [which we would not know from this analysis], one believed the reference case assumptions for Trojan operations either too pessimistic or too optimistic. [p.4A.I8]

The use of imprecise probabilities avoids these problems because it allows one to represent a lack of knowledge about the probability of values obtaining, while still bringing to bear all the evidence one does have about the probability of a value obtaining. If this analysis were done with imprecise probabilities, the result would be an upper and lower expected value for the two options. The upper and lower expected values would capture the entire range of what is "expected" given the state of knowledge. Thus, one would not have to wonder whether these expectations are too optimistic or too pessimistic. With imprecise probabilities, sensitivity analysis is automatic. Furthermore, the upper and lower expected values would avoid having to assume all else is equal; that is, interrelations among variables are not assumed away.

Suppose PGE had analyzed the decision using upper and lower probabilities. It might have come up with lower and upper expected costs for continued operation of $3 billion to $4 billion, and for 1996 phaseout of $1.5 billion to $4 billion. Then 1996 phaseout would be preferred because it has a lower expected cost and the same upper expected cost. Of course, such analyses will not always result in such a straightforward result. However, as PGE states: "Choosing a future and identifying results that are unacceptable for reasons not in the quantitative analysis requires the exercise of judgment." We argue that such judgment could be exercised with a respect for uncertainty if the quantitative analysis represents all evidence practically available, as well as the limits of what is and can be known, using imprecise probability and the qualitative framework for understanding uncertainty.

Avoiding the second guessing of expected values is probably reason enough to justify the imprecise probability approach. However, it does not obviate the need for sensitivity analysis; it changes its focus. The sole focus of sensitivity analysis in the imprecise probability approach, and what we believe is its strongest justification, is to identify the value of collecting additional information before making a choice.

Returning again to the Trojan decision, suppose that imprecise probabilities were used and that the lower and upper expected values are as above, $3 billion to $4 billion for operation and $1.5 billion to $4 billion for phaseout. Thus, the utility decides not to order new steam generators. However, one is still left with the decision of whether or not to close Trojan before 1996. Further suppose that the expected value for immediate shutdown is $1.5 billion to $4.5 billion. Again, the utility would probably want to choose phaseout because it has a lower maximum expected cost ($4 as opposed to $4.5 billion) and the same minimum expected cost. However, the choice is different than the choice between phaseout and operation because the option of operating past 1996 was closed after new steam generators were not ordered in time to get them installed by 1996.[9] In comparing phaseout to continued operation, the utility could still decide to shut down the plant any time before 1996. The present decision, (i.e., to acquire resources to replace Trojan by 1996) is the same for either of these two options. Thus, the utility has the opportunity to gather more information about the variables affecting the costs of operating Trojan compared to shutting it down.

For this example, suppose that the utility had constructed an imprecise probability distribution over the variable "average costs of purchased power (PP) to replace Trojan." (As noted in Example 1, a value of $5/MWh was assumed for the cost premium sellers would extract, which was included in the PP of $41/MWh.) Suppose that the expected value ranges from $30 to $45, and analysts note the large amount of operational uncertainty arising from having little experience with having to purchase large blocks of nonfirm power for an extended amount of time. Sensitivity analysis could be done to see how lowering the uncertainty would affect the decision. (Recall that operational uncertainty can be reduced with more time and money.)

Suppose that the operational uncertainty could be reduced by "shopping" for power. Further suppose that an analysis was done that showed this new information could pin down the expected PP value to $30 to $40 if the upper part of the distribution were proven wrong. (Such analyses are the subject of the next chapter.) (Note that we would not consider the effect of the lower part of the distribution being wrong because this would not improve the desirability of shutdown, and as things

[9]In fact, the option was not entirely closed. The 1996 date was picked to simplify the problem. PGE states "Rather than attempt to model the uncertainties that exist with operating the plant after 1996 without replacing the steam generators, PGE decided to focus the economic comparison on continued operation and phase-out in 1996. This decision also sets a planned termination date and avoids the costs, confusion and uncertainty that would arise from attempting to pursue both continued operation and replacement resources for some amount of time." [p.4A.3, footnote 2]

stand, phaseout is already preferred.) Further suppose that the cost of shopping for power is relatively low, and that the information gained would be useful for other utility decisions. Sensitivity analysis results show that the $30 to $40 expected PP range reduces the upper expected cost of 1996 phaseout by about $100 million. Thus, this analysis by itself does not change the preferred option. At most, the upper expected cost to shutdown would be $4.4 billion, and it could be higher. However, if reducible operational uncertainty and application uncertainty is examined in other variables, the combined effect might warrant further study, provided the reductions in uncertainty were cost effective. In the next chapter, we present an example where the reduction in application and irreducible uncertainty over time changes a variable's expected value range.

For now, it should be realized that one can gain a lot by using imprecise probabilities to calculate upper and lower expected values. Once one understands the uncertainty, the uncertainty can be fully represented with two "gauges," obviating the need for second-guessing sensitivity analyses, and finally, one can determine what, if anything, can be done about the uncertainty and whether it is worthwhile to do it.

3.8 Summary

The added expressiveness of imprecise probability is necessary in the world of IRP. This is because many if not most of the key variables in IRP are plagued by high levels of irreducible and operational uncertainty, many are plagued by application uncertainty, and many if not most of these key variables are not amenable to statistical investigation. One cannot honestly represent lack of knowledge about key variables using precise probability. The extra gauge offered by imprecise probability is completely necessary.

IRP analysts should explore the use of imprecise probability for the reasons stated above. It is our belief that the experts will find it easier to express themselves using two gauges rather than one. And of course, imprecise probability collapses to precise probability when knowledge is high. There are straightforward methods to calculate lower and upper expected values. In addition, by using imprecise probability, one avoids the potential morass of sensitivity analysis using precise probabilities.

In our model of the IRP process, the IRP analysts need to be aware of all of these points. The expert needs to be able to express uncertainty using lower and upper probabilities. We believe that experts will find it easier to provide ranges, especially in cases where they also report reasons for high levels of irreducible and operational uncertainty.

The PUC staff and utility executives need to understand that IRP analysts will be presenting them with bounds on expected values of key decision variables. The IRP analysts need to be able to recommend to the decision makers what activities could be undertaken to reduce uncertainty in the key variables to narrow the bounds and hopefully make their decisions easier. Thus, the analysts need to conduct the second type of sensitivity analysis to investigate how much operational and application uncertainty can potentially be overcome, how much it might cost, and how such activities would improve the decision making process.

4. MANAGING UNCERTAINTY

In developing integrated resource plans and investigating difficult and complex decisions, such as whether or not to shut down a power plant, it is advisable to use every piece of information that could conceivably help improve resource acquisition decisions. This could be a daunting task for the IRP analyst who would be responsible for putting the pieces together in some fashion and who may have in the past been forced to discard difficult to quantify information. This section briefly reviews several methods that can be used to help analysts *design* or *construct* solutions to problems that entail uncertainty and numerous, seemingly disparate pieces of information.

4.1 Consensus Methods

What should the IRP analyst do when confronted with more than one expert opinion about an estimate for a key decision variable and/or outputs from more than one model that contain estimates of the key variable? The analyst could choose one expert opinion or one model output and throw away the rest. This would be a waste of information if indeed the other experts and models possess legitimacy. The analyst could try to persuade the experts and modelers to come to some agreement amongst themselves on an estimate. Many may view this as a preferable approach but it is unlikely to happen in many real-life contexts.

A third option for the analyst would be to use consensus methods. These methods provide a rigorous means to synthesize judgments about the same set of events rendered by different people or different models. To use consensus methods, the judgments must be complete statements about the set of events and must be rendered by different knowledge sources (i.e., by different experts, by computer models that are different in some significant manner, by different theoretical models that purport to explain the same aspects of reality). The mathematics of consensus are well-understood and have been extended to encompass imprecise probabilities (Wagner 1989).

For example, let's assume that the IRP analyst is interested in estimating the remaining life for one of the utility's coal plants. The utility employs three people who can be considered experts in the maintenance and operation of coal plants and each has intimate knowledge of the plant in question. It is also known that the experts do not agree on how much longer the plant can last. To implement a consensus design, the IRP analyst would elicit from the experts lower and upper probabilities for a set of 4 to 5 lifetimes for the plant. A consensus formula can then be applied to these data to calculate consensus lower and upper probabilities for lifetimes for the plant, which can then be used in expected value calculations, etc.

4.2 Combination Methods

A common problem facing IRP analysts is how to synthesize disparate pieces of fragmentary evidence about some past, present or future state of the world. For example, suppose the analyst is faced with predicting a decision to be made by one of the utility's large industrial customers about whether to buy power from a competitor utility, build its own cogeneration unit, continue to buy power from the analyst's utility, or shut down the plant altogether. Much like a detective or attorney, the analyst should go about collecting pieces of information about the firm that might shed light on what the decision will be. Combination rules can be used to synthesize these pieces of evidence.

To illustrate, let's assume that the analyst is able to collect the following information: (1) the chairman of the firm has publicly stated that one of its four North American plants will need to be shut down this year; (2) the competitor utility has little excess capacity; (3) the site of the industrial plant is conducive for cogeneration; and (4) the utility has never interrupted power to the firm over the past 20 years.

To apply a combination rule, the analyst first assigns lower and upper probabilities to each of the four decision outcomes for each piece of evidence. Then a combination rule, such as Dempster's Rule of Combination (Shafer 1976) or another proposed by Tonn (1994), could be used to construct aggregate lower and upper probability assessments over the decision set. The analyst could just assign one set of lower and upper probabilities based on all of the evidence, but this would seem to be a cognitively daunting task. Doing separate assessments and using a combination rule is easier cognitively and is also more flexible if new information concerning any piece of evidence is obtained.

4.3 Conditioning Methods

Conditionalization (or conditioning) is a term which covers an array of methods that incorporate conditional probabilities (e.g., the probability of A, given B is true). Conditionalization is useful for (1) updating uncertainty measures when new information is received, (2) sharpening uncertainty measures with ancillary evidence, and (3) synthesizing evidence represented as marginal and conditional probabilities.

Let's revisit the case where the IRP analyst is constructing an imprecise probability to express what is known about a large industrial's electricity purchase decision in the future. The first row in Figure 4.1 summarizes, in an illustrative fashion, the results of applying a combination rule to the four pieces of evidence mentioned in 4.3. Visual inspection of the lower and upper probability gauges indicates that the present evidence does not support the buy from competitor decision but somewhat evenly supports the other three potential decisions.

Now let's assume a new piece of information arrives. The chairman of the large industrial announces that the plant in the utility's service area will not be shut down. In other words, one of the potential decisions has been eliminated with certainty from the set of potential decisions. *Exclusion conditioning methods* would be used to reallocate the lower and upper probabilities among the remaining decisions. Application of such a method is illustrated in the second row of Figure 4.1, which indicates tighter lower and upper probability bounds for the remaining decisions. Conditioning methods of this type are available for both precise and imprecise probability (Sundberg and Wagner 1992).

Next, let's assume that the IRP analyst comes across a report that addresses electricity purchase decisions of large industrials. Over 50 cases were assessed and statistical results were presented that indicate the proportion of industrials that chose buying from a competitor utility, cogeneration, and buying from its original utility. The IRP analysts believes that this study is of high quality (e.g., is characterized by low operational uncertainty) and is of value to this context (e.g., is characterized by low application uncertainty). Using a method known as *diagnostic conditionalization* (Wagner and Tonn 1990), the statistics can be used to sharpen the current imprecise probabilities that are based on a real example of a decision context. Diagnostic conditionalization produces precise probabilities from imprecise probabilities, as is illustrated in the third row of Figure 4.1.

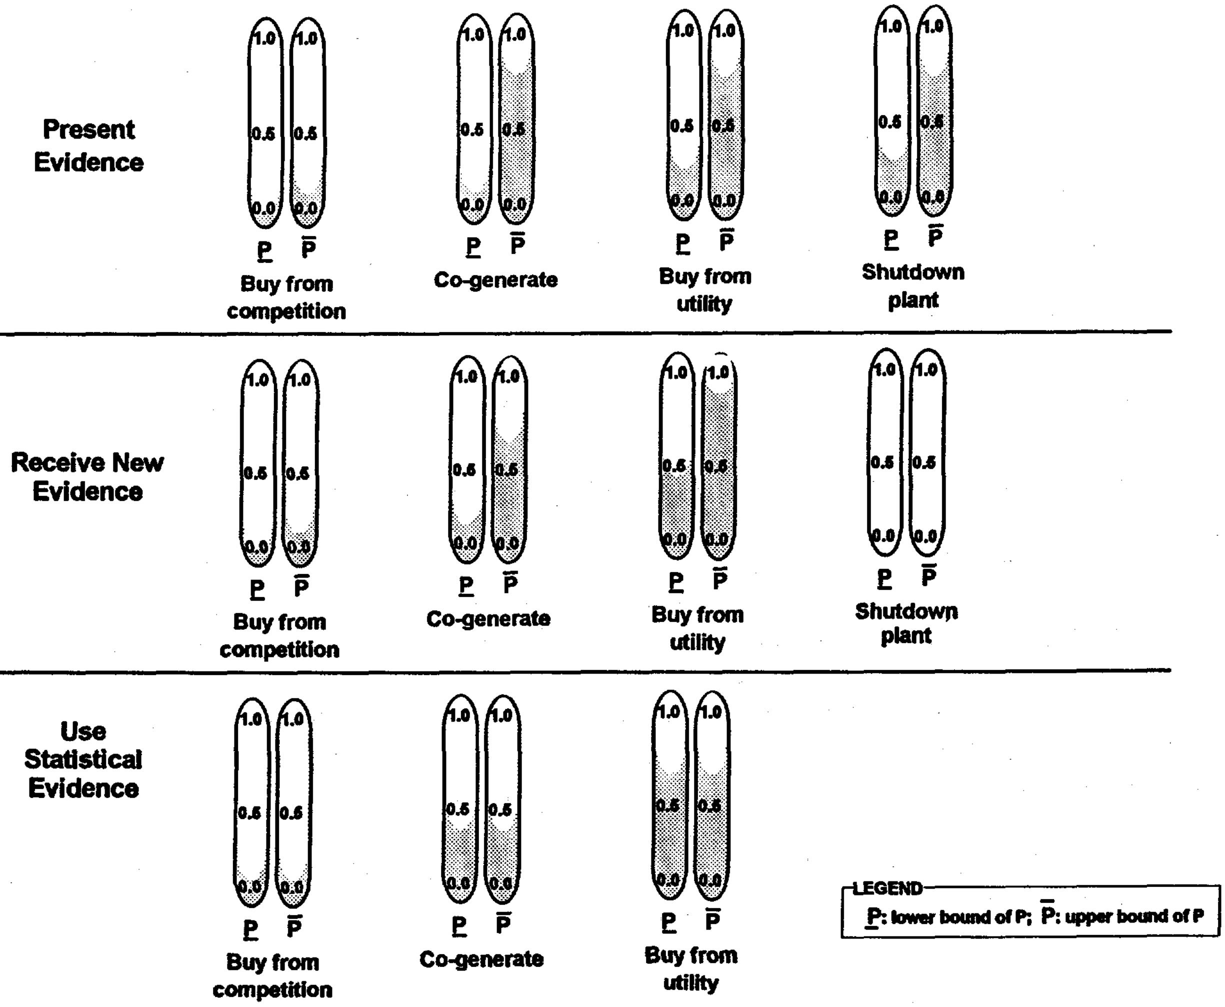

Fig. 4.1. Two types of conditioning.

The third, and last, conditionalization method discussed herein is known as Bayes Formula. It is a method for synthesizing probabilistic information concerning several events. One formulation of Bayes Formula is the following:

$$P(A \mid B) = \frac{P(A) * P(B \mid A)}{P(A) * P(B \mid A) + P(\bar{A}) * P(B \mid \bar{A})} \quad (4)$$

where P(A|B) is the probability of A occurring given that B has occurred (which is known as the posterior probability of A), P(B|A) is the probability of B occurring given that A has occurred, and P(A) and P(B) are prior probabilities of A and B occurring.

Bayes Formula, along with other formulations known as Bayes Rule and Bayes Theorem, is the foundation for a school of thought known as Bayesianism. Adherents of this school of thought believe that Equation 4 is an excellent manner to frame problems that involve uncertainty. Bayesians are subjectivists in that they believe that probabilities are introspective or personal. On the other hand, it is probably a fair observation that Bayesians have not yet supported the use of imprecise probability. As a consequence, Equation 4 has yet to be generalized for imprecise probabilities.

To illustrate an application of Bayes Formula, let's assume the IRP analyst is interested in calculating the probability of a scenario where a regional economic recession is prompted by an oil price shock. In other words, the goal is to calculate the conditional probability of recession given oil price shock. To use Bayes Formula, the analyst would need to collect data or ask experts (i.e., economists knowledgeable in this area) to estimate the following probabilities: probability of recession in any given year, probability of oil price shock given a recession, probability of there not being a recession, and the probability of an oil price shock given no recession.

4.5 Notes on Constructing Designs for Problem Solving Under Uncertainty

The beauty of the three types of methods presented above is that they provide the IRP analyst with a powerful toolkit to construct designs for problem solving under uncertainty. The methods can be used one at a time or in various combinations. The output from the application of a consensus rule can be used as a piece of evidence in the application of a combination rule. The output of a combination rule can be conditioned with new evidence or past statistics, as noted above. A conditionalized precise probability can be input into Bayes Rule. The challenge for the IRP analyst is to arrange the pieces of the puzzle in the most logical fashion.[10]

The following example presents just such a design related to the nuclear power plant shutdown decision.

---

[10]It should be noted that the literature on the subject of designs appears to pit Bayesianism against the Dempster-Shafer approach to reasoning under uncertainty (e.g., Shafer and Tversky 1985, Hobbes 1994). We do not see the approaches as competitors but as two different approaches for the IRP analysts' toolkit. The approach most appropriate for the problem at hand ought to be chosen, and as indicated, both can be synthesized into one design.

4.6 Design for a Power Plant Shutdown Decision

One of the key variables affecting the desirability of operating the Trojan nuclear plant is the plant's performance, expressed as its capacity factor (CF). This example demonstrates a design that a utility could use to:

(1) develop an imprecise probability function over Trojan's forecasted 1993-95 CF using combination methods,

(2) synthesize this estimate with other expert judgments about the probability function over Trojan's 1993-95 CF using consensus methods, and

(3) update the probability function based on new information using conditioning methods.

(1) Suppose that a utility analyst collects the following evidence about Trojan's forecasted capacity factor from 1993-95.

E1: Trojan's low variable cost generally ensures that it is dispatched whenever it is available, however there were several periods in the early 1980s when the plant's output was displaced with low-cost hydropower.

E2: Trojan historical data shows an average CF of 55% over the last 16 years.

E3: Degradation of Trojan's steam generator tubes are the cause for them needing replacement in 1996. A premature failure of the tubes would cause an unplanned outage. Even if the tubes do not fail, degradation beyond NRC specifications would cause a forced outage for repair (by plugging or sleeving the tubes).

E4: Due to steam generator tube degradation, the NRC granted permission to operate Trojan under one-year interim repair criteria. Regardless of the change over time in the condition of the tubes, a change in NRC's requirements could cause a forced outage. More generally, however, recent trends of stabilizing levels of regulatory activity in the nuclear industry suggest improvement over historical CFs.

The analyst decides that the range of possible CF values is 0% to 88%. He then specifies a LCDF and a UCDF given each piece of evidence. More specifically, for a sample of values CF_0, $CF_1, \ldots, CF_x$, in the possible range, the analyst would specify a lower probability that the true value is at least CF_x given E1, $\underline{P}(CF_x \backslash E1)$, and an upper probability that the true value is at least CF_x, $\overline{P}(CF_x \backslash E1)$. This would be done for the other pieces of evidence, E2 through E4. Then a combination rule is applied to generate an aggregate LCDF and UCDF over the range of possible values (see Figure 4.2).

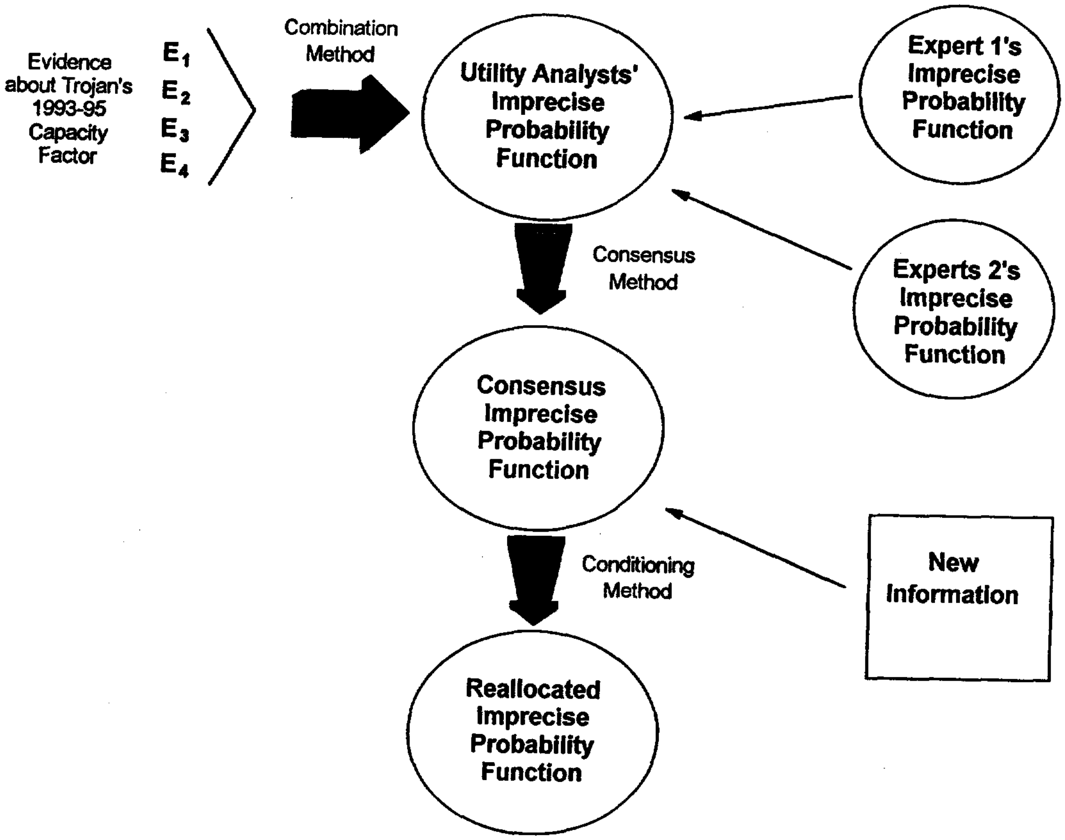

Fig. 4.2. Resolving uncertainty about Trojan's capacity factor.

(2) In addition to those of its own analysts, PGE obtained estimates of Trojan's CF from a consulting firm and BPA. Because the weighing of evidence requires judgment, it would be helpful to consider the judgments of these other experts. An analyst could elicit these other experts' judgments in the form of lower and upper CDFs assigned to a range of CFs these other experts deemed possible. Then the utility analyst could use a consensus formula to combine all three sets of lower and upper CDFs into consensus lower and upper CDFs. These consensus imprecise CDFs could be used in the analysis of the Trojan decision to calculate upper and lower expected values for Trojan costs.

(3) There may also be a need to generate new imprecise CDFs based on new information using conditioning methods. One case where PGE might have used conditioning methods is to revise its estimate of Trojan's CF based on the events of November 9, 1991. At that time, Trojan immediately shut down after a small leak occurred in a tube of one of its four steam generators. To reevaluate the Trojan shutdown decision, PGE could have conditioned the consensus imprecise CDFs based on this occurrence. This could be done by reallocating the upper and lower CDFs over a narrower range of possible CFs. More specifically, analysts might have narrowed the range of possible CFs to 0% to 60% based on the fact that reaching a CF of over 60% is no longer possible due to the down time while the steam generators would be repaired. An exclusion conditioning formula could be used to reallocate the upper and lower CDFs upon the new, narrower, range of possible CFs.

Another case where conditioning methods might be used is to evaluate the benefit of acquiring more information. In Example 4, we demonstrated how new information might narrow the range of expected values for replacement power purchase costs. Suppose there exists potential information that would narrow the range of *what values are possible for a variable*, but does not otherwise shed light on which of the remaining values will obtain (i.e., does not provide information that makes any of the remaining possibilities more or less probable relative to each other). In cases like this, exclusion conditioning could be used to generate reallocated lower and upper CDFs over the narrowed range of possible values. This would be done just as it was for new information that has already come to light, but the reallocated upper and lower probabilities would represent the potential information. Then, upper and lower expected values could be recalculated with the reallocated upper and lower CDFs to determine the benefit that acquiring the additional information would have.

4.7 Summary

This section is of primary interest to IRP analysts, who need to compile, organize, and synthesize as much information as possible in preparing integrated resource plans. The analysts need to be aware of the tools that are available, how they are applied, and how they can be used to create designs. As illustrated in the examples, using imprecise probabilities and the methods introduced above, analysts have no need to discard any piece of potentially valuable information from their analyses.

Experts should also be interested in these methods. For one reason, within their area of expertise, they may be faced with questions about how to synthesize pieces of information. Also, they need to appreciate how their contributions may be used as inputs to a sophisticated design. PUC commissioners and staff and utility decision makers need to have a general understanding of the methods, especially the consensus methods because how expert opinions are used is often a controversial topic. IRP analysts will need to prepare informative visual aids that explain how the methods are used without complicated mathematics and symbolism. For example, visual aids such as Figure 4.2 can be created to "tell a story" of how the pieces of information were synthesized that can be readily followed by a nontechnical audience.

5. CONCLUSIONS

Uncertainty is a challenging concept. As this report illustrates, one can approach the topic on many levels, both qualitatively and quantitatively. IRP has a long way to go to substantially handle uncertainty. Integrated resource plans need to have improved discussions of reasons for uncertainty in estimates, following the suggestions of 2.0. We recommend that integrated resource plans represent uncertainty quantitatively and that imprecise probability be used. Much work is needed to illustrate how the methods presented in 4.0 can be synthesized in sophisticated designs to support IRP.

Many other activities in the area of uncertainty and IRP can also be pursued. For the IRP practitioner, software needs to be developed around the methods presented in 4.0. Software could also be developed to assist IRP analysts elicit estimates and imprecise probabilities from experts. The methods themselves require additional refinement, especially the combination methods and a generalization of Bayes Formula would also be valuable. Methods for assisting all the IRP players communicate with each other about uncertainty also need to be enhanced. Visual methods should receive special attention in this regard.

Handling uncertainty is important in the electric power industry beyond IRP. The industry is undergoing significant restructuring, which some argue, may prove to be the demise of IRP. Maybe, maybe not. The important point is that all the entities in the electricity business will still need to do strategic planning, and PUCs and other government agencies will still need to protect the public interest. These players will still rely on informed, quantitative analyses, which will still be plagued with uncertainty. All the concepts of uncertainty discussed herein apply to these contexts, too.

6. ACKNOWLEDGEMENTS

We thank the many people who contributed ideas and perspectives for this report. We thank Eric Hirst and Carl Wagner for their reviews and comments on this draft report. This report also benefitted from comments on a previous, much more technical report on IRP and uncertainty. Comments from the following people helped shape this report: Tony Artman, John Lowell, Karl Wieland, Doug Logan, Nancy Esteb, Joe Eto, Richard Sonnenblick, James Valverde, Jim Hammitt, Stan Hadley, Lynn Maxwell, Steve Elliott, and Eric Hirst. Artman, Logan, Eto, Sonnenblick, Valverde, Hadley and Hirst also provided valuable comments on this report.

7. REFERENCES

Bernoulli J. , *Ars Conjectandi*, 1713.

Buchanan, B. and Shortliffe, E., *Rule-Based Expert Systems*, Addison-Wesley: Reading, Mass, 1984.

DeFinetti, B., "Foresight: Its Logical Laws, Its Subjective Sources," *Studies in Subjective Probability*, Eds. H. Kyburg and H. Smokler, John Wiley: New York, 93-118, 1964.

Fishburn, P., *Nonlinear Preference and Utility Theory*, Johns Hopkins Press: Baltimore, Maryland, 1988.

Functowitz, S. and Ravitz, J., *Uncertainty and Quality in Science for Policy*, Kluwer: Dordrecht, 1990.

Hacking, I., *The Emergence of Probability*, Cambridge University Press: Cambridge, England, 1975.

Hirst, E., *A Good Integrated Resource Plan: Guidelines for Electric Utilities and Regulators*, Oak Ridge National Laboratory, ORNL/CON-354, 1992.

Hirst, E. and Schweitzer, M., "Electric-Utility Resource Planning and Decision-Making: The Importance of Uncertainty," *Risk Analysis 10*,(1), 137-146, 1990.

Hobbes, B., "Bayesian Methods for Analyzing Risks from Climate Change," *Workshop on Risk Estimation Techniques for Extreme Weather Events*, International Geosphere-Biosphere Programme, Feb. 1994.

Jaynes, E., "Information Theory and Statistical Mechanics," *Physical Review 106*, 620-630, and *108*, 171-190, 1957.

Ramsey, F., "Truth and Probability," *The Foundations of Mathematics*, Ed. R. Braithwaite, Routledge and Kegan Paul: London, 156-198, 1931.

Savage, L., "The Foundations of Statistics Reconsidered," in *Studies in Subjective Probability*, Eds. H. Kyburg and H. Smokler, John Wiley: New York, 173-188, 1964.

Schmeidler, D., "Integral Representation Without Additivity," *Proc. American Mathematical Society 97*, 255-261, 1986.

Shafer, G., "Non-additive Probabilities in the Work of Bernoulli and Lambert," *Archive for History of Exact Sciences* 19, 309-370, 1978.

Shafer, G., *A Mathematical Theory of Evidence*, Princeton University Press: Princeton, New Jersey, 1976.

Shafer, G. and Tversky, A., "Languages and Designs for Probability Judgments," Cognitive Science 9, 309-339, 1985.

Shannon, C., "A Mathematical Theory of Communication," *Bell System Technical Journal 27*, 379-423 and 623-656, 1948.

Sundberg, C. and Wagner, C., "Generalized Finite Differences and Bayesian Conditioning of Choquet Capacities," *Advances in Applied Mathematics 13*, 262-272, 1992.

Tonn, B., "A Combination Rule for Belief Functions," draft paper, Oak Ridge National Laboratory, Oak Ridge, TN, 1994.

Tonn, B. and Schaffhauser, A., "Towards a More General Theoretical and Mathematical Model of Probability for Policy Analysis," *Environment and Planning A 24*, 1337-1353, 1992.

Tonn, B., and Wagner, C., *Handling Uncertainty in Quantitative Estimates in Integrated Resource Planning,* ORNL/CON-388, Oak Ridge National Laboratory, Oak Ridge, Tennessee, 1995.

Wagner, C., "Imprecise Probability and Integrated Resource Planning," Research Report, Dept. of Mathematics, University of Tennessee, 1994.

Wagner, C., "Consensus for Belief Functions and Related Uncertainty Measures," Theory and Decision 26, 295-304, 1989.

Wagner, C. and Tonn B., "Diagnostic Conditionalization," *Mathematical Social Sciences, 19*, 159-165, 1990.

Walley, P., *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall: London, 1991.

Wallsten, T., and Budescu, D., "Encoding Subjective Probabilities: A Psychological and Psychometric Review," *Management Science 29*, (2), 151-173, 1983.

Zadeh, L., "Fuzzy Sets as a Basis for a Theory of Possibility," *Fuzzy Sets Systems 1*, 3-28, 1978.

APPENDIX A: COMMENTS AND RESPONSES

Several of the reviewers posed important but technical questions about some of the ideas presented in this report. We felt that we could not address every technical issue raised without making the report too difficult to read for the intended audience. The purpose of this appendix is to address some of these questions. People interested in a mathematical presentation of the methods presented in 3.0 and 4.0 are urged to see Tonn and Wagner (1995).

(1) Why does it matter that precise probability represents complete ignorance the same way as strong statistical knowledge?

In 3.0, we argue that one the most important reasons for using imprecise probabilities is that precise probabilities can be viewed as representing in the same way both complete ignorance and strong statistical knowledge. A small number of reviewers questioned whether this point is important enough to prompt people to use imprecise probabilities over precise probabilities. One reviewer simply stated that we did not *prove* that this situation is a problem with precise probability.

We do not suppose to possess the ability to prove, in a formal sense, our argument. We can only offer an extended explanation for our views, which to us are strongly and intuitively reasonable. To begin our extended explanation, we need to start with the qualitative characterization of uncertainty presented in 2.0. Our position is that in many decision making situations, especially those dealing with IRP, problems will be plagued by high levels of irreducible, operational, and application uncertainty. Within one integrated resource plan, there may well be cases where one literally has no knowledge about key variables, a situation of complete ignorance, and there may well be instances where one has exceptional databases from which one can derive strong statistically based knowledge about key variables.

Thus, information, in a technical sense, about variables may range from zero to some maximum level attainable (if there is irreducible uncertainty) and possibly perfect information for a very limited set of variables. We work under the assumption that probability methods provide one method for measuring the amount of information one has about a variable. An absolute requirement for a good measuring device is that its measurement readings must be unambiguous. That is, each measurement reading must have a one to one mapping to what the device has been designed to measure. Thus, a thermometer provides a one to one mapping between the level of mercury in the thermometer and the actual temperature of the air (We are ignoring the reliability of the measurements at this point but address this problem in comment 4 below). A thermometer would be rather useless if, say a measurement of 60 °F on the thermometer could relate to both -30 °F and to 60 °F actual air temperatures.

Precise probability is deficient as a "device" to measure information one has at one's disposal about a variable because there is ambiguity in its measurements. As clearly shown (see Fig. 3.1), and not disputed by classical probabilists, there is ambiguity in uniform probabilities. Do they represent strong statistical knowledge or no knowledge whatsoever? Because the Principle of

Insufficient Reason instructs classical probabilists to assign the uniform function to a probability distribution when one has no knowledge, mathematically one cannot tell the difference when one has strong statistical evidence to support the uniform function (e.g., from throwing a fair die thousands of times).

We cannot imagine a more serious problem with a mathematical method purported to be valuable to decision makers, who are urged to apply decision making heuristics that blindly use probabilities as inputs regardless of what amount of information actually underlie their calculation. We think that decisions—alternatives for action and for more information collection—ought to be different when one has no knowledge versus cases where one has strong statistical knowledge. Using precise probabilities, decisions may be the same in both cases because inputs to decision methods would be the same in both cases.

If the probabilities are the same in either case, then why should anyone endeavor to gain statistical knowledge? Why not simply make decisions without getting information. Note that gaining even the strongest statistical knowledge does not make an indeterminate outcome determinate. There is still a range (or a set) of outcomes to consider. Gaining statistical knowledge does, however, supply some information about the outcomes. Clearly there is a benefit to gaining this information on which to base probabilities. If this information is valuable to have, then lacking this information must be undesirable. In any case, where there is significant ignorance about a probability, a precise probability will not represent this ignorance.

This point is the most important and most basic justification for using imprecise probability. The reason why imprecise probability is needed is because there is almost always irreducible, operational, and application uncertainty about the values of key variables in policy contexts and precise probability is inadequate to represent any significant amount of uncertainty. If one agrees that having more information is superior to having less information, then one should also agree that a situation where probabilities are based on more information should be distinguished from a situation where probabilities are based on less information.

We realize that people can wonder why, after precise probability has well served humanity all these years, we argue that it has a serious flaw. Our short history of uncertainty presented in 3.0 tries to explain how this could be. Simply put, the original ideas put forth by eminent thinkers about mathematical formulations of probability had much in common with imprecise probability and less in common with precise probability. Indeed, these thinkers were interested in evidential reasoning such as presented in 4.0. Subsequent thinkers were more interested in statistics related to insurance annuities and mathematical deduction, and finally in scientific data analysis. Classical probability evolved out of these concerns, not from concerns firmly focused on evidential reasoning and decision making, as IRP practitioners, artificial intelligence researchers, and policy and decision analysts understand them in the late 20th century.

Any student of the history of science and ideas will tell you that ideas gain inertia and change takes time. We argue that classical probability has gained an enormous amount of inertia but that now is the time to reassess, and re-evaluate its usefulness for applied decision making contexts. Imprecise probability is not the last word on the subject of representing uncertainty for applied decision making, but it has better qualities than precise probability (and even dissolves into precise probability given strong knowledge) and inertia is not a valid reason to uncritically support precise probability.

(2) Some argue that sensitivity analysis is not ad hoc, and the choices of which probabilities to vary in a sensitivity analysis need not be arbitrary.

Several reviewers have had experience with sensitivity analysis. They do not see the process as being ad hoc nor do they view the choices of different probabilities as being arbitrary. The typical goal for sensitivity analysis is to explore thresholds of changes in variable values that would necessitate a shift in decisions. Thus, the process is to first develop a base case, consisting of the most likely variable values and then explore under what conditions decisions would change from the base case decision. Under this description of sensitivity analysis, uncertainty about precise probabilities is only one reason, maybe even a minor reason, to explore different variable values.

Given this description, sensitivity analysis using precise probability still doesn't seem like a fruitful endeavor. What does one do if one finds that plausible changes in variable values does shift decisions from one to another? One reaction would be to collect more information to reduce the range of plausible variable values. But as mentioned above, precise probability is not very good at representing the amount of information brought to bear on a problem. Since precise probabilities admit no imprecision or ignorance, the value of reducing ignorance cannot be fully addressed. The value of information in the precise probability paradigm is limited to considering how new information would change the values of a probability distribution. It cannot consider the effect of reducing ignorance. Thus, while the motivation underlying this aspect of sensitivity analysis is strong, precise probability does not offer a strong response. At least with imprecise probabilities, more work could reduce the range between upper and lower probabilities, given that some level of operational and application uncertainty exists.

Another reaction would be to work on the decision set. This report has shied away from discussing decision making because it is an enormous topic. In addition, we believe that many decision heuristics can be straightforwardly used with imprecise probabilities. However, we also believe that using imprecise probabilities will go hand in hand with formative changes in the way we think about decision making. It is quite predictable, in our view, that many using imprecise probabilities will find, when all is said and done, that they do not know as much about a situation that they thought they might nor will they ever know as much about a situation as they would like. We make no apologies for presenting reality in a harsher focus if this is what is needed to improve decision making. Sets of decision alternatives structured for more certain situations may not contain any "good" alternatives for situations with large magnitudes of uncertainty. Changing the decision sets to be more contingent, more reversible, more incremental will be ways to find decision alternatives that will fare better in uncertain futures.

Lastly, as discussed in 3.0, conducting sensitivity analysis if one is uncertain about precise probabilities is not necessary when one uses imprecise probabilities. The upper and lower probabilities are firm bounds that represent one's knowledge at hand. That's it. If one feels the need to change the upper and lower probabilities because of uncertainty about these bounds, then the bounds have not been properly specified. There is no need for this kind of sensitivity analysis within the theory of imprecise probability.

(3) Some argue that using upper and lower probabilities will cause a combinatorial explosion of expected values for complex models where numerous variable interact with each other.

In other words, several reviewers do not understand how upper and lower expected values for variables can be used in a straightforward manner to calculate target expected values (e.g., upper and lower expected power plant maintenance costs). For example, variables A, B, and C could be used to determine the value for variable Z and all four could have upper and lower expected values. Wouldn't one manually need to use different combinations of A, B, and C to calculate many expected values for Z?

The answer is no. In Chapter 5.0 of Tonn and Wagner (1995), Carl Wagner has laid out a straightforward approach to use operations research techniques to find the minimum expected value for Z and the maximum expected value for Z. There is no combinatorial explosion. This report contains an example of Wagner's method that has been applied to a real problem using real data. Wagner's method does not depend upon Monte Carlo methods, just nonlinear optimization algorithms. Wagner's method can be implemented using available computer codes. In fact, all the methods presented in this report can be computerized, and many have been. A coherent, commercially available package for analysts has yet to be developed.

(4) Are imprecise probabilities more difficult to use than precise probabilities, from the expert's perspective?

This point wasn't made directly by any of the reviewers but is hinted at in several places. We think there is some confusion about uncertainty elicitation. Much has been written about uncertainty elicitation, mainly by psychologists. The primary findings are that people are not necessarily good at providing precise probabilities and that elicitation aids like probability wheels often help. Our argument is that people should be better at providing imprecise probabilities, because they have more room to represent their uncertainty about the context. Imprecise probability wheels and other aids can also be used to assist people in the elicitation of imprecise probabilities. Not much research has been done by psychologists to support or reject our argument, although our discussions with cognitive psychologists leads us to believe that there should be no problem in using methods developed to elicit precise probabilities to elicit imprecise probabilities.

If properly designed, an elicitation protocol will allow subjects to provide reliably either precise probabilities, if they have a high level of knowledge, or imprecise probabilities if they do not. Subjects can provide a single CDF if they have a high level of knowledge or a LCDF and a UCDF if they do not. There is a continuum at work here, not an either or situation, where one must use either precise or imprecise probabilities. Whatever is provided by the expert is used in the analysis as is. Combination rules are not used to merge LCDFs and UCDFs to get CDFs. One uses what one gets from the experts, period, without artificial alternations.

(5) Is Bayesianism distinct from imprecise probability?

There are many people who refer to themselves as Bayesians. It is our understanding that they believe that all probabilities are subjective in nature and believe that Bayes Rule and its variants should be used as *the* structure to assess problems that entail uncertainty. As a group, they have been hostile to a theoretical approach to managing uncertainty, which is referred to as Dempster-Shafer Theory (DST). Very briefly, DST is based upon the elicitation and combination of belief functions, which are special versions of imprecise probability functions. Instead of eliciting probabilities, DST calls for the elicitation of basic probability mass (there are mathematical transformations to go from basic probabilities to upper and lower probabilities and back again, Shafer 1976). Bayesians object to DST because they object to the elicitation of basic probability mass over probabilities and the use of combination rules over Bayes Rules and its variants. It appears to us that Bayesians also object to imprecise probability, because it is seen as a direct extension of DST.

We believe that Bayesians do not fully understand our approach. Our formulation presented in this report is **not** a direct extension of DST. First, we advocate the elicitation of upper and lower probabilities, not basic probability mass. We do not understand how people would even go about thinking about basic probability mass. There is no indication in the psychological literature that people think in these terms of basic probability mass and there is not psychological research on the topic. Second, given that Bayesians believe that all probabilities are subjective, which we also believe, we do not see how being a Bayesian per se leads one to object the use of upper and lower probabilities. In fact, imprecise probability allows a broader range of subjectivity to enter into the analysis.

Third, it should be noted again that Bayes Rule and its variants can be generalized to handle upper and lower probabilities, although we have not conducted this research. Thus, Bayes Rule and its variants would still be available.

Fourth, we do not advocate that the decision between probabilistic methods is between Bayes Rule on one hand or combination rules on the other. As we have tried to show in 4.0, each has its place and proper application context. One cannot be substituted for the other. Their proper application contexts are not comparable. We believe that Bayesians have tried to fit many problems into a Bayes Rule formulation that could have better been structured using combination and consensus rules. Conversely, combination rules are not appropriate in contexts that call for Bayes rule. The literature has done a terrible job in even acknowledging that the methods are not competitors but are complementary tools that can be used together to construct solutions to problems of uncertainty.

If Bayesians agree that a generalized Bayes Rule can be used with other methods (e.g., combination and consensus) for problem solving, along with the use of upper and lower probabilities, and basically the same elicitation methods that they are used to, then there is really no difference between Bayesianism and the theory of imprecise probability and the methods presented herein.

(6) What is meant by information?

The concept of information has a very technical definition within classical probability. The amount of information in a probability function is defined to be its entropy. The famous formula for entropy is

$$H = - \sum P_i * \log(P_i).$$

Provided perfect information, the equation yields a value of 0.0, meaning an absence of entropy. Maximum entropy is when the equation is provided a uniform probability distribution. Thus, ignorance is expressed as maximum entropy, and as argued above, strong statistical knowledge may also yield maximum entropy. It should be noted that this formula was developed for use in the communications industry to track the flow of bits of information through telephone lines, not for direct use applied decision making contexts.

A better information measure would be one where the equation yields 0.0 when one has no knowledge, a different value for strong statistical knowledge, and a maximum value, normalized to 1.0, when one has perfect knowledge. With respect to imprecise probability, the equation would yield 0.0 when the lower probability is 0.0 and the upper probability is 1.0 for each potential event, i.e., the case of no knowledge. The equation would yield a normalized 1.0 when the lower probability and the upper probability equal 1.0 for one event. The uniform probability distribution, where the lower and upper probabilities are equal for all the potential events, would have a value in between, preferably toward 1.0 because the uniform probability distribution actually does contain a relatively large amount of information.

Thus, it is important to point out that in our view, imprecise probability offers a more expanded approach not only to conceive of the concept of information, but from a technical viewpoint, precise probability, in conjunction with the above formula, restricts information representation. Work on an information measure for imprecise probabilities has been slow. The equation above is not capable of meeting these criteria for imprecise probabilities. We have developed an initial equation for belief functions. This is an important area for future research.