Regulation under Uncertainty

OR

Uncertainty Analysis and the Use of Expert Judgment

RFF Workshop on Expert Judgment:
Promises and Pitfalls
2006 March 13

M. Granger Morgan
Head, Department of
Engineering and Public Policy
Carnegie Mellon University
tel: 412-268-2672
e-mail: granger.morgan@andrew.cmu.edu

This morning I will talk about:

- Sources of uncertainty and the characterization of uncertainty.
- Two basic types of uncertainty.
    - Uncertainty about coefficient values.
    - Uncertainty about model functional form.
- Performing uncertainty analysis and making decisions in the face of uncertainty.
- Some summary guidance on reporting, characterizing and analyzing uncertainty.

Probability

Probability is the basic language of uncertainty.

I will adopt a personalistic view of probability (sometimes also called a subjectivist or Bayesian view).

In this view, probability is a statement of the degree of belief that a person has that a specified event will occur given all the relevant information currently known by that person.

$P(X|i)$ where:

X is the uncertain event

i is the person's state of information.

The clairvoyant test

Even if we take a personalist view of probability, the event or quantity of interest must be well specified for a probability, or a probability distribution, to be meaningful.

"The retail price of gasoline in 2008" does not pass this test. A clairvoyant would need to know things such as:

- Where will the gas be purchased?
- At what time of year?
- What octane?

Does a subjectivist view mean your probability can be arbitrary?

NO, because if they are legitimate probabilities, they must

- conform with the axioms of probability
- be consistent with available empirical data.

Lots of people ask, why deal with probability? Why not just use subjective words such as "likely" and "unlikely" to describe uncertainties? *There are very good reasons not to do this*.

The risks of using qualitative uncertainty language

Qualitative uncertainty language is inadequate because:

- the same words can mean very different things to different people.
- the same words can mean very different things to the same person in different contexts.
- important differences in experts' judgments about mechanisms (functional relationships), and about how well key coefficients are known, can be easily masked in qualitative discussions.

Carnegie Mellon University

Mapping words to probabilities

This figure shows the range of probabilities that people are asked to assign probabilities to words, absent any specific context.

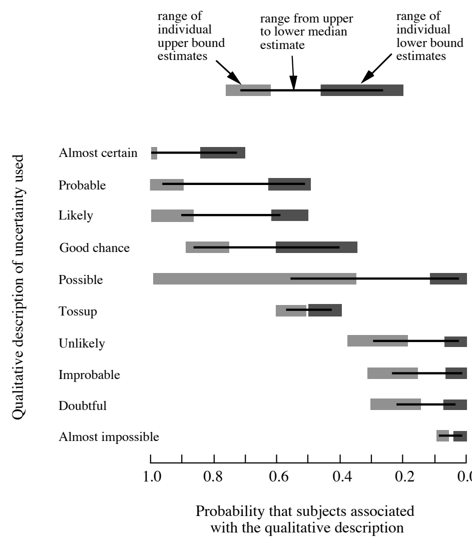

Figure adapted from Wallsten et al., 1986.

Department of Engineering and Public Policy

7

Carnegie Mellon University

Ex Com of EPA SAB

The minimum probability associated with the word "likely" spanned four orders of magnitude.

The maximum probability associated with the word "not likely" spanned more than five orders of magnitude.

There was an overlap of the probability associated with the word "likely" and that associated with the word "unlikely"!

Figure from Morgan, *HERA*, 1998.

Department of Engineering and Public Policy

8

The bottom line

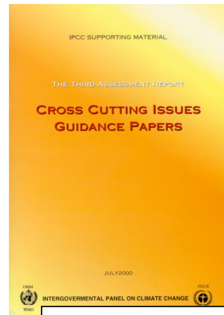

Without at least some quantification, qualitative descriptions of uncertainty convey little, if any, useful information.

The climate assessment community is gradually learning this lesson.

Steve Schneider and Richard Moss have worked hard to promote a better treatment of uncertainty in the work of the IPCC.

At my insistence, U.S. national assessment synthesis team gave quantitative definitions to five probability words:

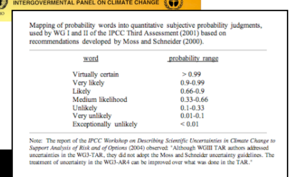

Mapping of probability words into quantitative subjective probability judgments, used by WG I and II of the IPCC Third Assessment (2001) based on recommendations developed by Moss and Schneider (2000).

word	probability range
Virtually certain	> 0.99
Very likely	0.9-0.99
Likely	0.66-0.9
Medium likelihood	0.33-0.66
Unlikely	0.1-0.33
Very unlikely	0.01-0.1
Exceptionally unlikely	< 0.01

Note: The report of the *IPCC Workshop on Describing Scientific Uncertainties in Climate Change to Support Analysis of Risk and of Options* (2004) observed: "Although WGIII TAR authors addressed uncertainties in the WG3-TAR, they did not adopt the Moss and Schneider uncertainty guidelines. The treatment of uncertainty in the WG3-AR4 can be improved over what was done in the TAR."

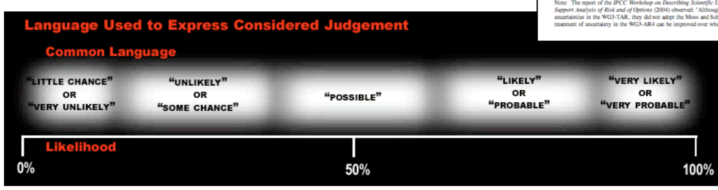

BUT, in other fields…

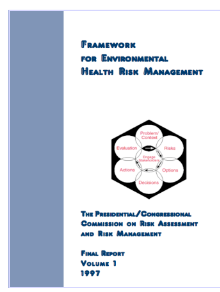

…such as biomedical and health effects, progress has been *much* slower.

A concrete example of this is provided by the recommendations of Presidential/Congressional Commission on Risk Assessment and Risk Management (1997) which recommended…"against routine use of formal quantitative analysis of uncertainty in risk estimation, particularly that related to evaluating toxicology."

While analysts were encouraged to provide "*qualitative* descriptions of risk-related uncertainty," the Commission concluded that "quantitative uncertainty analyses of risk estimates are seldom necessary and are not useful on a routine basis to support decision-making."

Slowly such views are giving way, but progress is slow.

Carnegie Mellon University

This afternoon I will talk about:

- Sources of uncertainty and the characterization of uncertainty.
- Two basic types of uncertainty.
  - Uncertainty about coefficient values.
  - Uncertainty about model functional form.
- Designing and performing expert elicitation.
- Performing uncertainty analysis
- Some summary guidance on reporting, characterizing and analyzing uncertainty.

Department of Engineering and Public Policy 11

Carnegie Mellon University

We must consider two quite different kinds of uncertainty

1. Situations in which we know the relevant variables and the functional relationships among them, but we no not know the values of key coefficients (e.g., the "climate sensitivity").
2. Situations in which we are not sure what all the relevant variables are, or the functional relationships among them (e.g., will rising energy prices induce more technical innovation?).

Both are challenging, but the first is much more easily addressed than the second. I'll talk more about the second when I talk about uncertainty analysis.

Department of Engineering and Public Policy 12

Uncertainty about quantities

Type of quantity	Examples	Treatment of uncertainty
Empirical parameter or chance variable	Thermal efficiency, oxidation rate, fuel price	Probabilistic, parametric, or switchover
Defined constant	Atomic weight, π, joules per kilowatt-hr	Certain by definition
Decision variable	Plant size (utility), emissions cap (EPA)	Parametric or switchover
Value parameter	Discount rate, "value of life," risk tolerance	Parametric or switchover
Index variable	Longitude and latitude, height, time period	Certain by definition
Model domain parameter	Geographic region, time horizon, time increment	Parametric or switchover
Outcome criterion	Net present value, utility	Determined by treatment of its inputs

From Morgan and Henrion, *Uncertainty*, Cambridge, 1990/98.

PDFs and CDFs

A number of examples I am about to show are in the form of probability density functions (PDFs) or cumulative distribution functions (CDFs).

Since some of you may not make regular use of PDFs and CDF's, let me take just a moment to remind you...

Probability density function or PDF

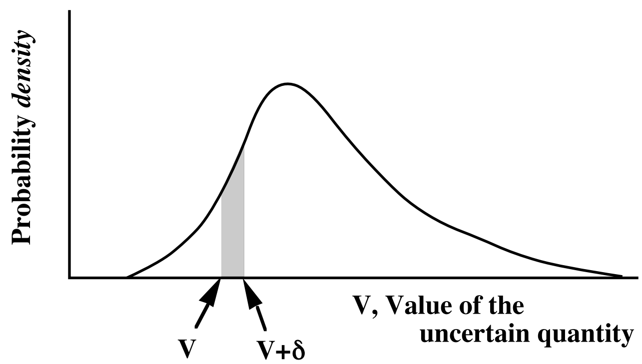

Cumulative distribution function or CDF

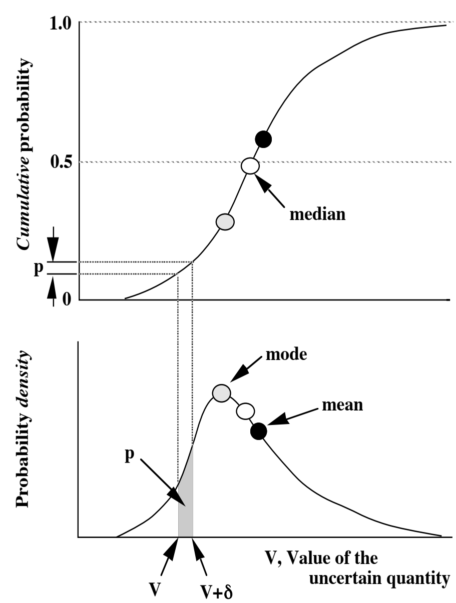

NOTE: In asymmetric distributions with long tails, the mean may be much much larger than the median.

If I have good data...

...in the form of many observations of a random process, then I can construct a probability distribution that describes that process. For example, suppose I have the 145 years of rainfall data for San Diego, and I am prepared to assume that over that period San Diego's climate has been "stationary" (that is the basic underlying processes that create the year-to-year variability have not changed)…

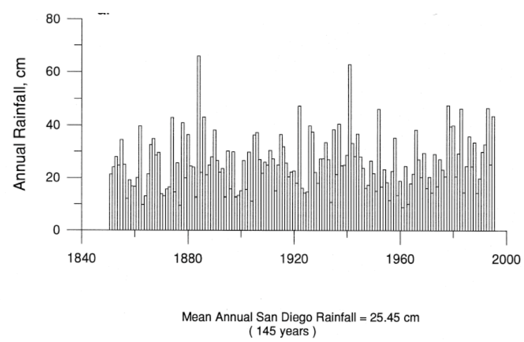

Source: Inman et al., Scripps, 1998.

Then if I want…

…a PDF for future San Diego annual rainfall, the simplest approach would be to construct a histogram from the data, as illustrated to the right.

If I want to make a prediction for some *specific* future year, I might go on to look for time patterns in the data. Even better, I might try to relate those time patterns to known slow patterns of variation in the regional climate, and modify my PDF accordingly.

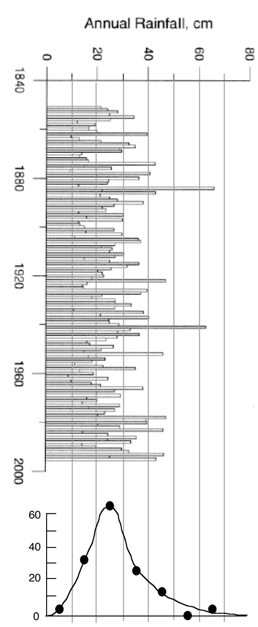

In that way…

…I could construct a PDF and CDF for future San Diego rain-fall that would look roughly like this.

However, suppose that what I really care about is the probability that very large rainfall events will occur.

Since there have only been two years in the past 145 years when rainfall has been above 60 cm/yr over, I'll need to augment my data with some model or physical theory.

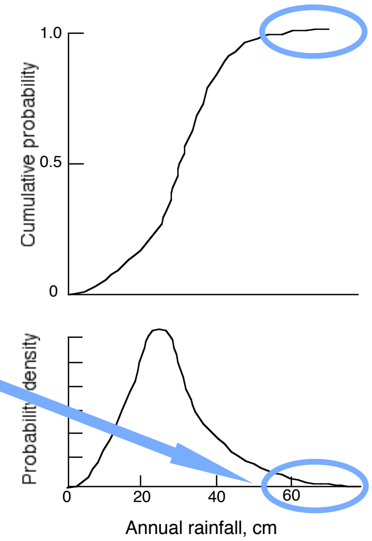

In summary…

…one should use available data, and well-established physical and statistical theory, to describe uncertainty whenever either or both are available.

However, often the available data and theory are not exactly relevant to the problem at hand, or they are not sufficiently complete to support the full objective construction of a probability distribution.

In such cases, I may have to rely on expert judgment. This brings us to the problem of how to "elicit" expert judgment.

Carnegie Mellon University

Expert elicitation takes time and care

Eliciting subjective probabilistic judgments requires careful preparation and execution.

Developing and testing an appropriate interview protocol typically takes several months. Each interview is likely to require several hours.

When addressing complex, scientifically subtle questions of the sorts involved with most problems in climate change, there are no satisfactory short cuts. Attempts to simplify and speed up the process almost always lead to shoddy results.

Department of Engineering and Public Policy

21

Carnegie Mellon University

Over confidence is a ubiquitous problem

21 different studies on questions with known answers:

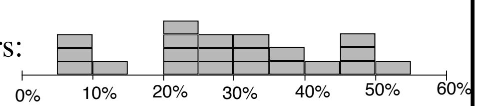

Percentage of estimates in which the true value lay outside of the respondent's assessed 98% confidence interval.

For details see Morgan and Henrion, *Uncertatinty*, Cambridge Univ. Press, 1990, pg 117.

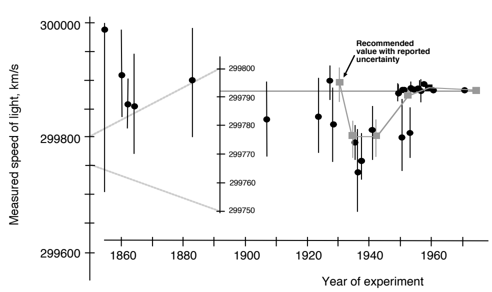

Estimates of the speed of light

For details see: Henrion and Fischhoff, "Assessing Uncertatinty in Physical Constants," *American Journal of Physics*, 54, pp791-798, 1986.

Department of Engineering and Public Policy

22

Cognitive heuristics

When ordinary people or experts make judgments about uncertain events, such as numbers of deaths from chance events, they use simple mental rules of thumb called "cognitive heuristics."

In many day-to-day circumstances, these serve us very well, but in some instances they can lead to bias - such as over confidence - in the judgments we make.

This can be a problem for experts too.

The three slides that follow illustrate three key heuristics: "availability," "anchoring and adjustment," and "representativeness."

Carnegie Mellon University

Cognitive bias

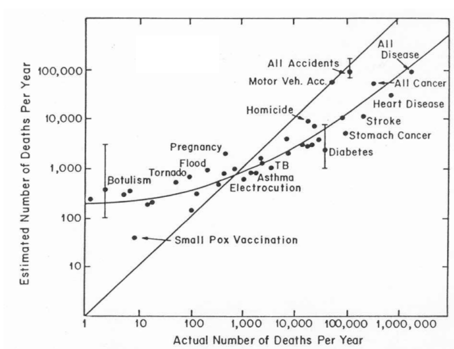

from Lichtenstein et al., 1978.

Availability: probability judgment is driven by ease with which people can think of previous occurrences of the event or can imagine such occurrences.

Cognitive bias…(Cont.)

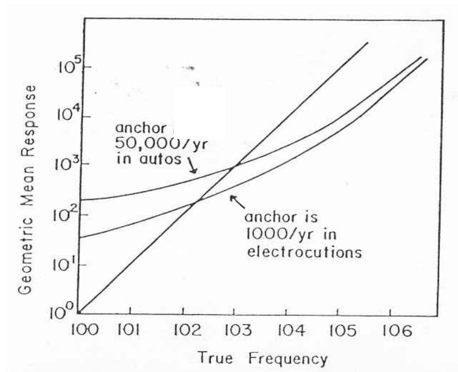

from Lichtenstein et al., 1978.

Anchoring and adjustment: probability judgment is frequently driven by the starting point which becomes an "anchor."

Cognitive bias…(Cont.)

I flip a fair coin 8 times. Which of the following two outcomes is more likely?

Outcome 1: T, T, T, T, H, H, H, H

Outcome 2: T, H, T, H, H, T, H, T

Of course, the two specific sequences are equally likely...but the second seems more likely because it looks more representative of the underlying random process.

Representativeness: people judge the likelihood that an object belongs to a particular class in terms of how much it resembles that class.

Expert elicitation …(Cont.)

Over the past three decades, my colleagues and I have developed and performed a number of substantively detailed expert elicitations. These have been designed to obtain experts' *considered* judgments. Examples include work on:

Health effects of air pollution from coal-fired power plants.

- M. Granger Morgan, Samuel C. Morris, Alan K. Meier and Debra L. Shenk, "A Probabilistic Methodology for Estimating Air Pollution Health Effects from Coal-Fired Power Plants," *Energy Systems and Policy, 2*, 287-310, 1978.
- M. Granger Morgan, Samuel C. Morris, William R. Rish and Alan K. Meier, "Sulfur Control in Coal-Fired Power Plants: A Probabilistic Approach to Policy Analysis," *Journal of the Air Pollution Control Association, 28*, 993-997, 1978.
- M. Granger Morgan, Samuel C. Morris, Max Henrion, Deborah A.L. Amaral and William R. Rish, "Technical Uncertainty in Quantitative Policy Analysis: A Sulfur Air Pollution Example," *Risk Analysis, 4,* 201-216, 1984 September.
- M. Granger Morgan, Samuel C. Morris, Max Henrion and Deborah A. L. Amaral, "Uncertainty in Environmental Risk Assessment: A case study involving sulfur transport and health effects," *Environmental Science and Technology, 19*, 662-667, 1985 August.

Expert elicitation…(Cont.)

Climate science, climate impacts and mitigation technology:

- M. Granger Morgan and David Keith, "Subjective Judgments by Climate Experts," *Environmental Science & Technology*, *29*(10), 468-476, October 1995.
- Elizabeth A. Casman, M. Granger Morgan and Hadi Dowlatabadi, "Mixed Levels of Uncertainty in Complex Policy Models," *Risk Analysis*, *19*(1), 33-42, 1999.
- M. Granger Morgan, Louis F. Pitelka and Elena Shevliakova, "Elicitation of Expert Judgments of Climate Change Impacts on Forest Ecosystems," *Climatic Change*, *49*, 279-307, 2001.
- Anand B. Rao, Edward S. Rubin and M. Granger Morgan, "Evaluation of Potential Cost Reductions from Improved CO_2 Capture Systems,"*Proceedings of the 2nd National Conference on Carbon Sequestration,* Alexandria, VA, May 5-8, 2003.
- M. Granger Morgan, Peter J. Adams and David W. Keith, "Elicitation Of Expert Judgments of Aerosol Forcing," *Climatic Change*, in press.
- Kirsten Zickfeld, Anders Levermann, M. Granger Morgan, Till Kuhlbrodt, Stefan Rahmstorf, and David W. Keith, "Present State and Future Fate of the Atlantic Meridional Overturning Circulation as Viewed by Experts," in review at *Climatic Change*.

Bounding uncertain health risks:

- M. Granger Morgan, "The Neglected Art of Bounding Analysis," *Environmental Science & Technology*, *35*, pp. 162A-164A, April 1, 2001.
- Minh Ha-Duong, Elizabeth A. Casman, and M. Granger Morgan, "Bounding Poorly Characterized Risks: A lung cancer example," *Risk Analysis*, *24*(5), 1071-1084, 2004.
- Elizabeth Casman and M. Granger Morgan, "Use of Expert Judgment to Bound Lung Cancer Risks," *Environmental Science & Technology*, *39*, 5911-5920, 2005.

Expert elicitation…(Cont.)

In all our elicitation studies, we've focused on creating a process that allows the experts to provide their carefully considered judgment, supported by all the resources they may care to use. Thus, we have:

- Prepared a background review of the relevant literatures.
- Carefully iterated the questions with selected experts and run pilot studies with younger (Post-doc) experts to distil and refine the questions.
- Conducted interviews in experts' offices with full resources at hand.
- Provide ample opportunity for subsequent review and revision of the judgments provided.

All of these efforts have involved the development of new question formats that fit the issues at hand.

SO_X health effects

…elicited subjective judgments about oxidation rates and wet and dry deposition rates for SO_2 and $SO_4^=$ from nine atmospheric science experts.

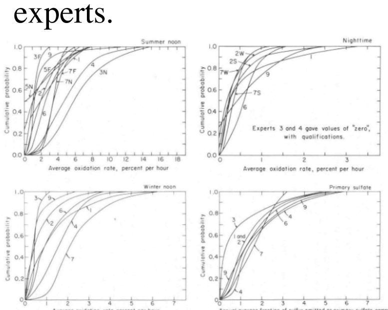

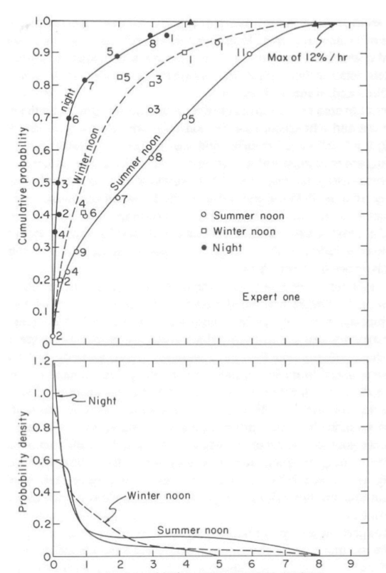

Source: G. Morgan, S. Morris, M. Henrion, D. Amaral, W. Rish, "Technical Uncertainty in Quantitative Policy Analysis: A sulfur air pollution example," *Risk Analysis*, *4*, 201-216, 1984.

Carnegie Mellon University

SO_X health effects…(Cont.)

Then for each expert we built a separate plume model, exposing people on the ground using known population densities.

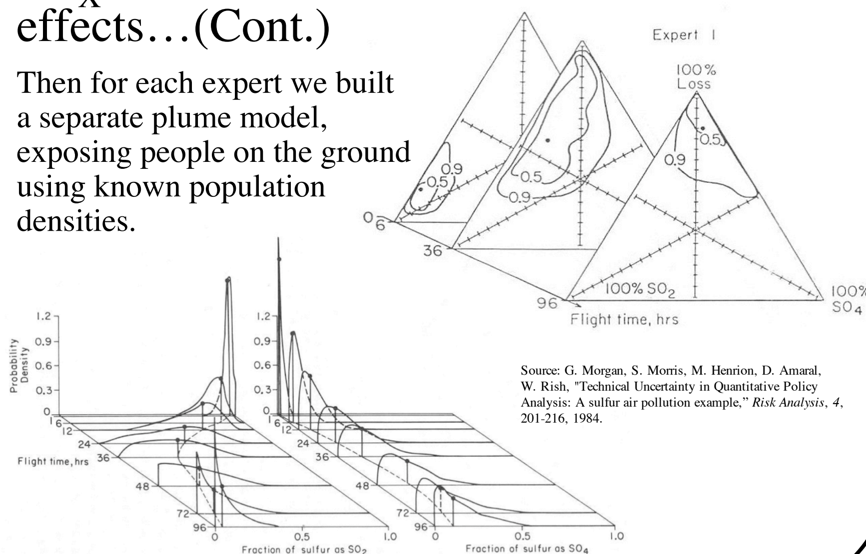

Source: G. Morgan, S. Morris, M. Henrion, D. Amaral, W. Rish, "Technical Uncertainty in Quantitative Policy Analysis: A sulfur air pollution example," *Risk Analysis*, *4*, 201-216, 1984.

Department of Engineering and Public Policy

31

Carnegie Mellon University

SO_X health effects…(Cont.)

We elicited health damage functions for sulfate aerosol from seven health experts (of whom only five were able to give us probabilistic estimates).

Finally, for each air expert, we did an analysis using the health damage function of each health expert…

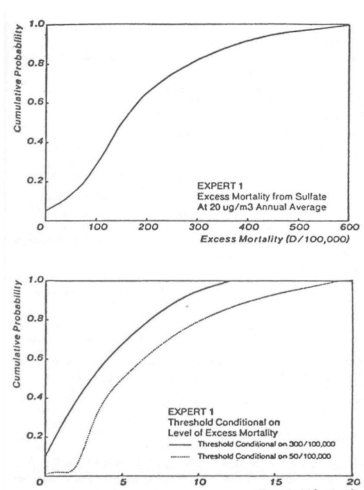

Source: D. Amaral, "Estimating Uncertainty in Policy Analysis: Health effects from inhaled sulfur oxides," Ph.D. thesis, Department of Engineering and Public Policy, Carnegie Mellon, 1983.

Department of Engineering and Public Policy

32

Carnegie Mellon University

SO_X health effects…(Cont.)

The results showed that while there was great discussion about uncertainty in the atmospheric science, the uncertainty about the health damage functions *completely* dominated the estimate of health impacts.

CDF of deaths/yr from a new (in 1984) super-critical $1GW_e$ FGD equipped coal-fired plant in Pittsburgh. Availability factor = 73%, net efficiency = 35%.

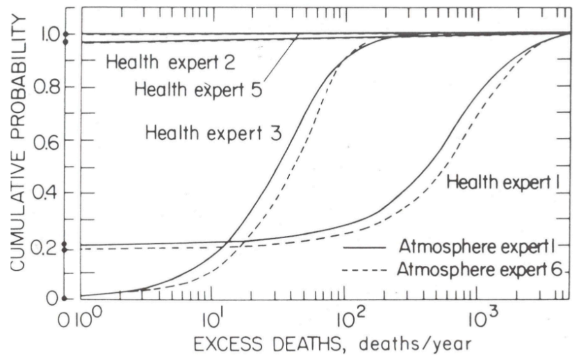

Source: G. Morgan, S. Morris, M. Henrion, D. Amaral, W. Rish, "Technical Uncertainty in Quantitative Policy Analysis: A sulfur air pollution example," *Risk Analysis*, *4*, 201-216, 1984.

Department of Engineering and Public Policy

33

Carnegie Mellon University

Multiple experts

When different experts hold different views it is often best not to combine the results, but rather to explore the implications of each expert's views so that decision makers have a clear understanding of whether and how much the differences matter in the context of the overall decision.

However, sophisticated methods have been developed that allow experts to work together to combine judgments so as to yield a single overall composite judgment.

The community of seismologists have made the greatest progress in this direction through a series of very detailed studies of seismic risks to built structures (Hanks, T.C., 1997; Budnitz, R.J. et al., 1995).

Department of Engineering and Public Policy

34

Uncertainty versus variability

Variability involves random change over time or space (e.g., "the mid-day temperature in Beijing in May is variable").

Recently, in the U.S., some people have been drawing a sharp distinction between variability and uncertainty. While the two are different, and sometimes require different treatments, the distinction can be overdrawn. In many contexts, variability is simply one of several sources of uncertainty (Morgan and Henrion, 1990).

One motivation people have for trying to sharpen the distinction is that variability can often be measured objectively, while other forms of uncertainty require subjective judgment.

This morning I will talk about:

- Sources of uncertainty and the characterization of uncertainty.
- Two basic types of uncertainty.
  - Uncertainty about coefficient values.
  - Uncertainty about model functional form.
- Performing uncertainty analysis and making decisions in the face of uncertainty.
- Some summary guidance on reporting, characterizing and analyzing uncertainty.

Uncertainty about model form

Often uncertainty about model form is as or more important than uncertainty about values of coefficients. Until recently there had been little practical progress in dealing with such uncertainty, but now there are several good examples:

- John Evans and his colleagues at the Harvard School of Public Health (e.g., Evans et al., 1994).
- Alan Cornell and others in the seismic risk (e.g., Budnitz et al., 1995).
- Hadi Dowlatabadi and colleagues at Carnegie Mellon in Integrated Assessment of Climate Change - ICAM (e.g., Morgan and Dowlatabadi, 1996).
- Also on climate, Lempert and colleagues at RAND (e.g., Lempert, Popper, Bankes, 2003).

John Evans and colleagues….

…have developed a method which lays out a "probability tree" to describe all the plausible ways in which a chemical agent might cause harm. Then experts are asked to assess probabilities on each branch.

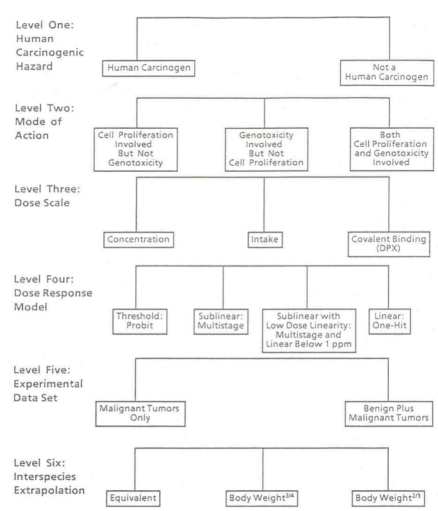

For details see: John S. Evans et al., "A distributional approach to characterizing low-dose cancer risk," *Risk Analysis, 14*, 25-34, 1994; and John S. Evans et al., "Use of probabilistic expert judgment in uncertainty analysis of carcinogenic potency," *Regulatory Toxicology and Pharmacology*, *20*, 15-36, 1994.

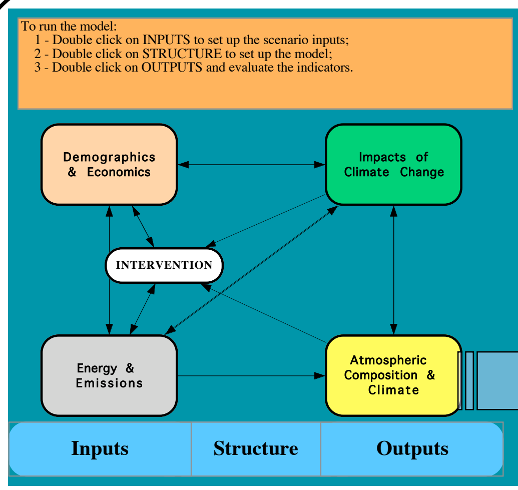

ICAM
Integrated Climate
Assessment Model

A very large hierarchically organized stochastic simulation model built in Analytica®.

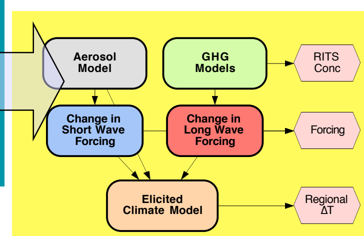

See for example:
Hadi Dowlatabadi and M. Granger Morgan, "A Model Framework for Integrated Studies of the Climate Problem," *Energy Policy*, *21*(3), 209-221, March 1993.
and
M. Granger Morgan and Hadi Dowlatabadi, "Learning from Integrated Assessment of Climate Change," *Climatic Change*, *34*, 337-368, 1996.

ICAM deals with...

…both of the types of uncertainty I've talked about:

1. It deals with uncertain coefficients by assigning PDFs to them and then performing stochastic simulation to propagate the uncertainty through the model.
2. It deals with uncertainty about model functional form (e.g., will rising energy prices induce more technical innovation?) by introducing multiple alternative models which can be chosen by throwing "switches."

ICAM

There is not enough time to present any details from our work with the ICAM integrated assessment model. Here are a few conclusions from that work:

- Different sets of plausible model assumptions give dramatically different results.
- No policy we have looked at is dominant over the wide range of plausible futures we've examined.
- The regional differences in outcomes are so vast that few if any policies would pass muster globally for similar decision rules.
- Different metrics of aggregate outcomes (e.g., \$s *versus* hours of labor) skew the results to reflect the OECD or developing regional issues respectively.

These findings lead us...

...to switch from trying to project and examine the future, to using the modeling framework as a test-bed to evaluate the relative robustness, across a wide range of plausible model futures, of alternative strategies that regional actors in the model might adopt.

We populated the model's regions with simple decision agents and asked, which behavioral strategies are robust in the face of uncertain futures, which get us in trouble.

Thus, for example, it turns out that tracking and responding to atmospheric concentration is more likely to lead regional policy makers in the model to stable strategies than tracking and responding to emissions.

Our conclusion

Prediction and policy optimization are pretty silly analytical objectives for much assessment and analysis related to the climate problem.

It makes much more sense to:

- Acknowledge that describing and bounding a range of futures may often be the best we can do.
- Recognize that climate is not the only thing that is changing, and address the problem in that context.
- Focus on developing adaptive strategies and evaluating their likely robustness in the face of a range of possible climate, social, economic and ecological futures.

Recent work by Robert Lempert and colleagues takes a very similar approach (e.g., Lempert, Popper, Bankes, 2003).

The next several slides…

…which I will not show because time is short, talk about two other important topics:

- A caution about the use of scenario analysis
- Some methods for dealing with extreme uncertaitinty.

I'll be happy to talk off-line or during a discussion session with anyone who is interested.

Jump to slide 50

Scenarios…

…can be a useful device to help think about the future, *but they can also be dangerous*. Remember the discussion of cognitive heuristics.

Here is a scenario from an experiment run by Slovic, Fischhoff, and Lichtenstein:

Tom is of high intelligence, although lacking in true creativity. He has a need for order and clarity, and for neat and tidy systems in which every detail finds its appropriate place. His writing is rather dull and mechanical, occasionally enlivened by somewhat corny puns and by flashes of imagination of the sci-fi type. He has a strong drive for competence. He seems to have little feel and little sympathy for other people and does not enjoy interacting with others.

In light of these data…

…what is the probability that:

Tom W. will select journal-ism as his college major? P = 0.21	Tom W. will select journalism as his college major but become unhappy with his choice? P = 0.39	Tom W. will select journalism as his college major but become unhappy with his choice and switch to engineering? P = 0.41

From: P. Slovic, B. Fischhoff, and S. Lichtenstein, "Cognitive Processes and Societal Risk Taking," *Cognition and Social Behavior*, 1976.

Scenarios…

…that describe just some single point in the space of outcomes are of limited use and logically cannot be assigned a probability.

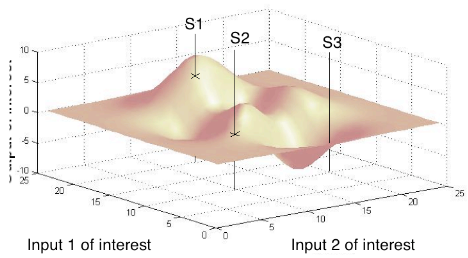

If scenarios are to be used, its better to span the space of interest and then assign a probability to each.

$$\sum_{i=1}^{n} Si = 1$$

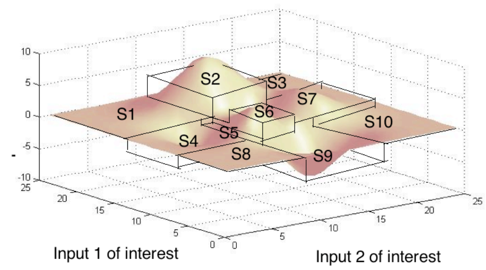

Cut the long causal chains

Typically, it is also better not to use detailed scenarios but rather to use simpler parametric methods.

Thus, for example, if future oil prices appear to be critical to a specific class of decisions, rather than develop long detailed stories about how those prices might be shaped by future developments in the U.S. and Canadian Arctic, the Middle East, and the Former Soviet Union, it is better to reflect on all the possibilities and then truncate the causal chain by positing a range of possible future oil prices and work from there.

In multiplicative models, uniform PDFs are often quite adequate to get good first-order estimates.

Limited domain of model validity

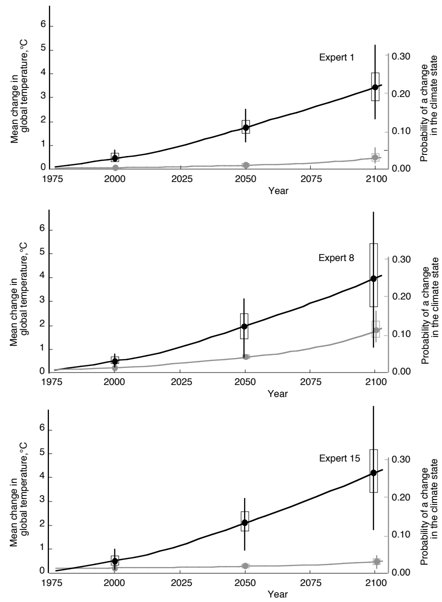

Examples of warming estimated via the ICAM model (dark curves) and probability that the associated climate forcing will induce a state change in the climate system (light curves) using the probabilistic judgments of three different climate experts.

Model switching

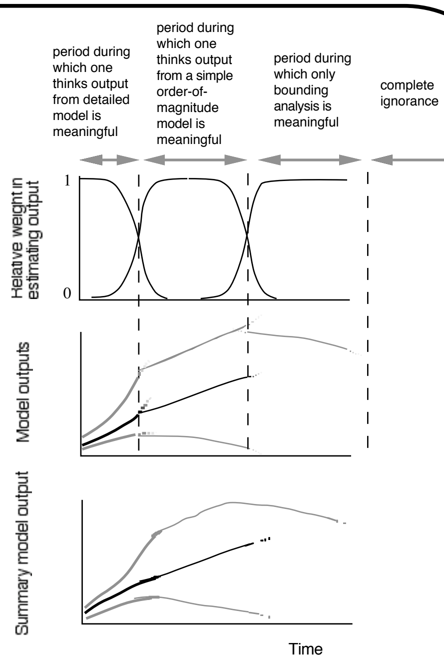

Schematic illustration of the strategy of switching to progressively simpler models as one moves into less well understood regions of the problem phase space, in this case, over time.

Illustration of model switching

Results of applying the model switch-over strategy to the ICAM demographic model (until about 2050) and an estimate of the upper-bound estimate of global population carrying capacity based on J. S. Cohen.

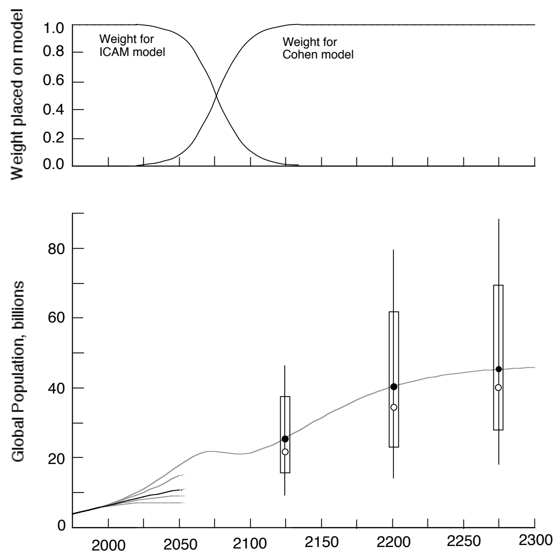

This morning I will talk about:

- Sources of uncertainty and the characterization of uncertainty.
- Two basic types of uncertainty.
  - Uncertainty about coefficient values.
  - Uncertainty about model functional form.
- Performing uncertainty analysis and making decisions in the face of uncertainty.
- Some summary guidance on reporting, characterizing and analyzing uncertainty.

Carnegie Mellon University

Propagation and analysis of uncertainty

Consider a basic model of the form $y = f(x_1, x_2)$

Simplest form of uncertainty analysis is sensitivity analysis using the slopes of y with respect to the inputs x_1 and x_2

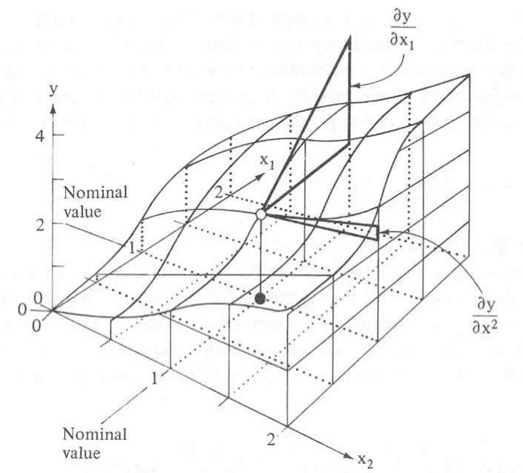

Source: Morgan and Henrion, 1990

Carnegie Mellon University

Nominal range sensitivity

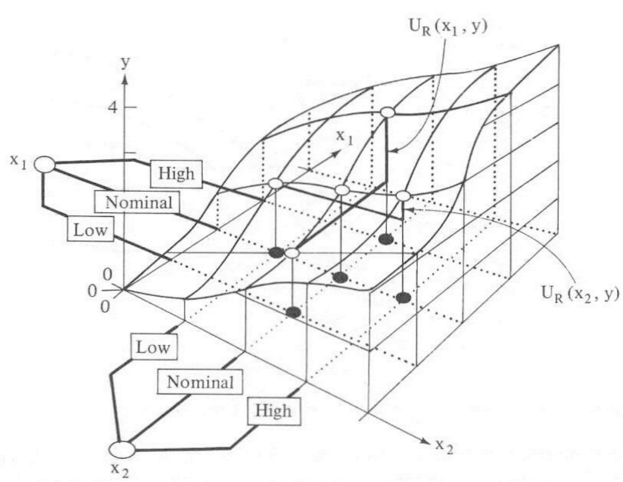

Source: Morgan and Henrion, 1990

Propagation of continuous distributions

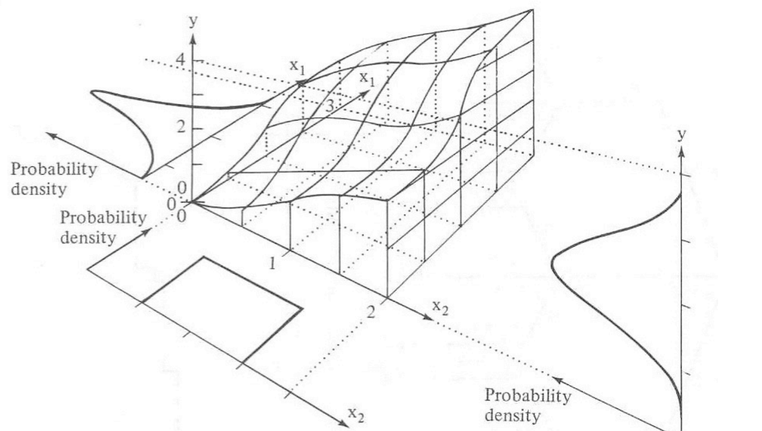

There are many software tools available such as Analytica®, @risk®, Crystal Ball®.

Source: Morgan and Henrion, 1990

Tools for analysis

Tools for continuous processes:

- Exposure models
- Dose response functions
- etc.

Tools for discrete events:

- Failure modes and effects analysis
- Fault tree models
- etc.

Use of some of these tools used to be very challenging and time consuming. Today such analysis is facilitated by many software tools (e.g., Analytica®, @risk®, Crystal Ball®, etc.).

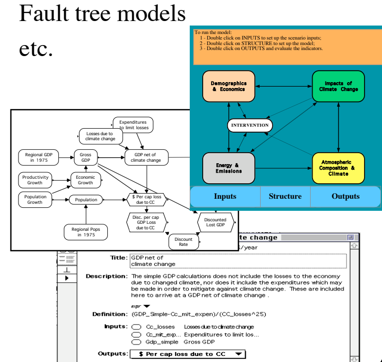

All that fancy stuff…

…is only useful when one has some basis to predict functional relationships and quantify uncertainties about coefficient values.

When one does not have that luxury, order-of-magnitude arguments and bounding analysis (based on things like conservation laws, etc.) may be the best one can do.

The conventional decision-analytic model assumes that research reduces uncertainty:

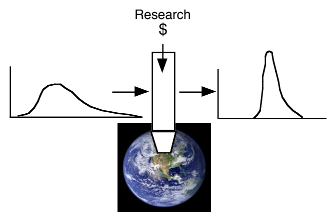

Research and uncertainty…(Cont.)

Sometimes this is what happens, but often, for variables or processes we really care about, the reality is that uncertainty either remains unchanged…

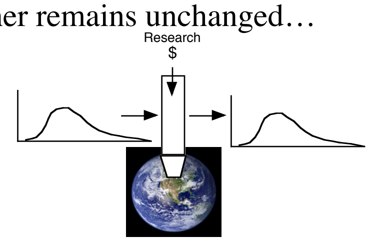

or even grows…

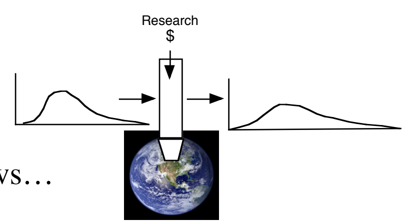

When the latter happens it is often because we find that the underlying model we had assumed is incomplete or wrong.

There are formal methods to support decision making…

…such as B-C analysis, decision analysis, analysis based on multi-attribute utility theory, etc.

Things analysts need to remember:

- Be explicit about decision rules (e.g., public health *versus* science; degree of precaution; etc.).
- Get the value judgments right and keep them explicit.
- Don't leave out important issues just because they don't easily fit.
- Remember that many methods become opaque very quickly, especially to non-experts.
- When uncertainty is high, it is often best to look for adaptive as opposed to optimal strategies.

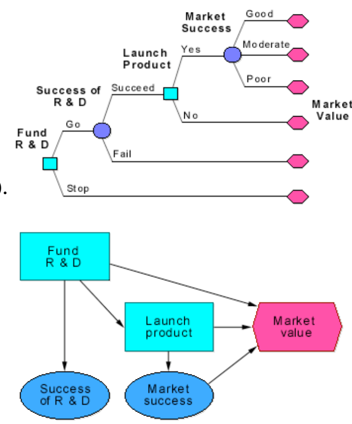

Source: Henrion, Lumina Systems

A classic strategy for solving problems:

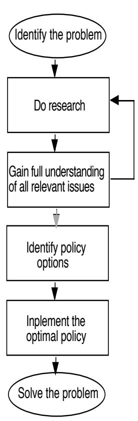

When uncertainty is high (and perhaps irreducible) an iterative adaptive strategy is better:

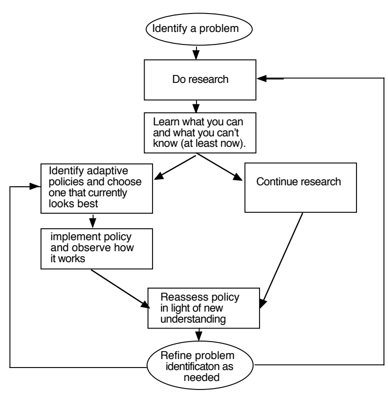

This morning I will talk about:

- Sources of uncertainty and the characterization of uncertainty.
- Two basic types of uncertainty.
  - Uncertainty about coefficient values.
  - Uncertainty about model functional form.
- Performing uncertainty analysis and making decisions in the face of uncertainty.
- Some summary guidance on reporting, characterizing and analyzing uncertainty.

Carnegie Mellon University

CCSP Guidance Document

Along with several colleagues I am currently completing a guidance document for the U.S. Climate Change Science Program.

The final slides reproduce our draft summary advice.

We'd welcome feedback and suggestions.

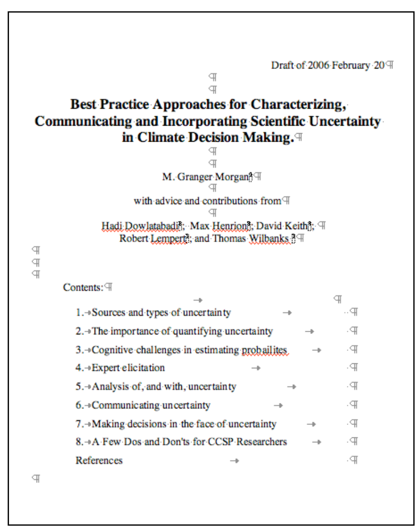

Draft of 2006 February 20

Best Practice Approaches for Characterizing, Communicating and Incorporating Scientific Uncertainty in Climate Decision Making.

M. Granger Morgan

with advice and contributions from

Hadi Dowlatabadi; Max Henrion; David Keith; Robert Lempert; and Thomas Wilbanks

Contents:

Department of Engineering and Public Policy

Doing a good job…

…of characterizing and dealing with uncertainty can never be reduced to a simple cookbook. One must always think critically and continually ask questions such as:

- Does what we are doing make sense?
- Are there other important factors which are as or more important than the factors we are considering?
- Are there key correlation structures in the problem which are being ignored?
- Are there normative assumptions and judgments about which we are not being explicit?

That said, the following are a few words of guidance…

Reporting uncertainty

When qualitative uncertainty words (such as likely and unlikely) are used, it is important to clarify the range of subjective probability values that are to be associated with those words. Unless there is some compelling reason to do otherwise, we recommend the use of the framework shown below:

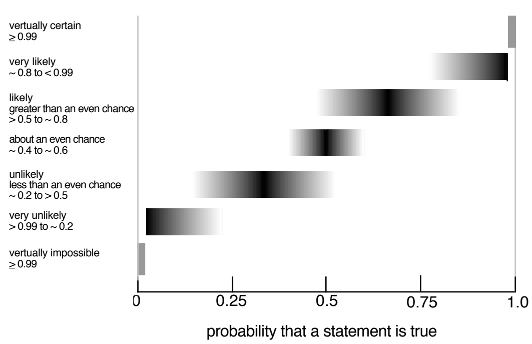

Reporting uncertainty…(Cont.)

Another strategy is to display the judgment explicitly as shown:

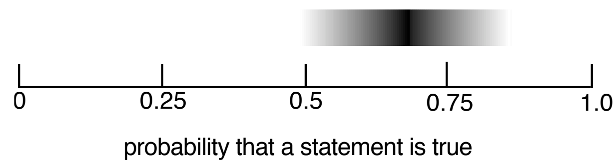

This approach provides somewhat greater precision and allows some limited indication of secondary uncertainty for those who feel uncomfortable making precise probability judgments.

Carnegie Mellon University

Reporting uncertainty…(Cont.)

In any document that reports uncertainties in conventional scientific format (e.g 3.5 ± 0.7), it is important to be explicit about what uncertainty is being included and what is not, and to confirm that the range is plus or minus one standard deviation. This reporting format is generally not appropriate for large uncertainties or where distributions have a lower or upper bound and hence are not symmetric.

Care should be taken in plotting and labeling the vertical axes when reporting PDFs. The units are probability *density* (i.e., probability per unit interval along the horizontal axis), *not* probability.

Since many people find it difficult to read and correctly interpret PDFs and CDFs, when space allows it is best practice to plot the CDF together with the PDF on the same x-axis.

Reporting uncertainty…(Cont.)

When many uncertain results must be reported, box plots (first popularized by Tukey, 1977) are often the best way to do this in a compact manner. There are several conventions. Our recommendation is shown below, but what is most important is to be clear about the notation.

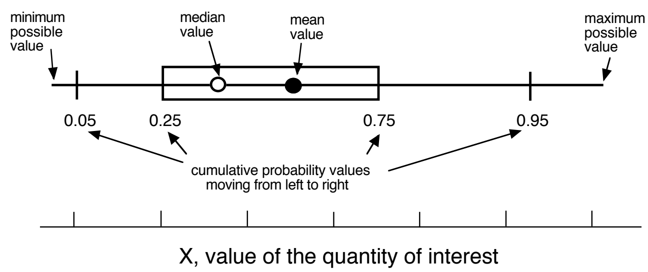

Tukey, John W., *Exploratory Data Analysis*, Addison-Wesley, 688pp., 1977.

Reporting uncertainty…(Cont.)

While there may be circumstances in which it is desirable or necessary to address and deal with second-order uncertainty (e.g., how sure an expert is about the shape of an elicited CDF), one should be very careful to determine that the added level of such complication will aide in, and will not unnecessarily complicate, subsequent use of the results.

Characterizing and analyzing uncertainty

Unless there are compelling reasons to do otherwise, conventional probability is the best tool for characterizing and analyzing uncertainty.

Characterizing and analyzing…(Cont.)

The elicitation of expert judgment, often in the form of subjective probability distributions, can be a useful way to combine the formal knowledge in a field as reflected in the literature with the informal knowledge and physical intuition of experts. Elicitation is not a substitute for doing the needed science, but it can be a very useful tool in support of research planning, private decision making, and the formulation of public policy.

HOWEVER the design and execution of a good expert elicitation takes time and requires a careful integration of knowledge of the relevant substantive domain with knowledge of behavioral decision science.

Characterizing and analyzing…(Cont.)

When eliciting probability distributions from multiple experts, if they disagree significantly, it is generally better to report the distributions separately then to combine them into an artificial "consensus" distribution.

There are a variety of software tools available to support probabilistic analysis using Monte Carlo and related techniques. As with any powerful analytical tool, their proper use requires careful thought and care.

Characterizing and analyzing…(Cont.)

In performing uncertainty analysis, it is important to think carefully about possible sources of correlation. One simple procedure for getting a sense of how important this may be is to run the analysis with key variables uncorrelated and then run it again with key variables perfectly correlated. Often, in answering questions about aggregate parameter values, experts assume correlation structures between the various components of the aggregate value being elicited. Sometimes it is important to elicit the component uncertainties separately from the aggregate uncertainty in order to reason out why specific correlation structures are being assumed.

Characterizing and analyzing…(Cont.)

Methods for describing and dealing with data pedigree (see for example Funtowicz and Ravetz, 1990) have not been developed to the point that they can be effectively incorporated in probabilisitic analysis. However, the quality of the data on which judgments are based is clearly important and should be addressed, especially when uncertain information of varying quality and reliability is combined in a single analysis.

While full probabilistic analysis can be useful, in many context simple parametric analysis, or back-to-front analysis (that works backwards from an end point of interest) may be as or more effective in identifying key unknowns and critical levels of knowledge needed to make better decisions.

Funtowicz, S.O. and J.R. Ravetz, *Uncertainty and quality in science for policy*, Kluwer Academic Publishers, 229 pp,1990.

Characterizing and analyzing…(Cont.)

Scenarios analysis can be useful, but also carries risks. Specific detailed scenarios can become cognitively compelling, with the result that people may overlook many other pathways to the same end-points. It is often best to "cut the long causal chains" and focus on the possible range of a few key variables which can most affect outcomes of interest.

Scenarios which describe a single point (or line) in a multi-dimensional space, cannot be assigned probabilities. If, as is often the case, it will be useful to assign probabilities to scenarios, they should be defined in terms of intervals in the space of interest, not in terms of point values.

Characterizing and analyzing…(Cont.)

Analysis that yields predictions is very helpful when our knowledge is sufficient to make meaningful predictions. However, the past history of success in such efforts suggests great caution (see for example Ch.s 3 and 6 in Smil, 2005). When meaningful prediction is not possible, alternative strategies, such as searching for responses or policies that will be robust across a wide range of possible futures, deserve careful consideration.

Smil, Vaclav, *Energy at the Crossroads,* MIT Press, 448pp., 2005.

Characterizing and analyzing…(Cont.)

For some problems there comes a time when uncertainty is so high that conventional modes of probabilistic analysis (including decision analysis) may no longer make sense. While it is not easy to identify this point, investigators should continually ask themselves whether what they are doing makes sense and whether a much simpler approach, such as a bounding or order-of-magnitude analysis, might be superior (see for example Casman et al., 1999).

Casman, Elizabeth A., M. Granger Morgan and Hadi Dowlatabadi, "Mixed Levels of Uncertainty in Complex Policy Models," *Risk Analysis*, *19*(1), 33-42, 1999.

The use of subjective probability

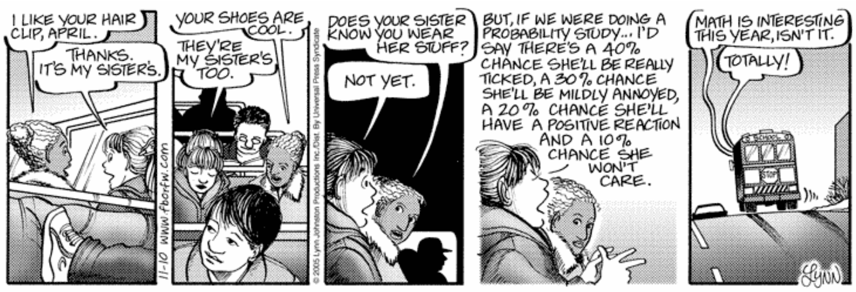

Lynn Johnson, Thur., Nov. 10, 2005.