Robin Gregory

IAN: Get back on time and all that stuff. Hello again. I'm Ian. As I was thinking in more discrete terms about this workshop and exactly who should be talking about what, Robin Gregory expanding on uncertainty and risk was one of my immediate choices. As it turns out, on top of being imminently qualified, he's also been very gracious from the beginning. So without further ado, I give you Robin Gregory.

ROBIN: Thank you, Ian. I don't know if I've ever been called gracious before. So thank you. From what I hear, it's been a really good conference, so I'll try to not lower the bar by too much.

I'm not a hydrologist; I'm not a water manager. My background's in natural resource economics and psychology and ecology—Buzz Holling, Carl Walters and Tony Scott at UBC along with Dan Kahneman and Paul Slovic. Most of my work now focuses on decision making and decision aiding techniques for tough problems where uncertainty is a big player, both factual uncertainty and uncertainties on the values side, where my role is to help people think through how they feel about something and to reveal what matters to them. This requires addressing tradeoffs and values uncertainty as well as biological uncertainty. It's generally kind of a mix. I also do a lot of teaching, research, and some consulting. In the past few years I've worked on projects for the US Fish and Wildlife Service, EPA, NOAA, Fisheries and Oceans Canada, BC Hydro, and other agencies involved in resource management decision making. And along with Lance Gunderson – whom I think is here somewhere? -- I have been involved for several years in some work for the Lake Champlain Basin program. I find that often the most challenging and interesting way to do research is to get my hands dirty in applied projects, because that's really where you know if stuff works or not.

I'll try to give a short introduction here, in about half an hour, to some ideas dealing with resource management under uncertainty. I'll make sure I leave a bit of room at the end for questions, and I am also sticking around afterwards for at least the first part of the afternoon. So if there's something I say that's of interest and you want to talk more, ask questions, tell me why I'm wrong, or add comments I'm very open to that. It's hard to have one-on-one conversations in a big room like this.

Many of the case studies I'll discuss were done with Nate Dieckmann and Marcus Mayorga, both with Decision Research in Eugene, or with Michael Harstone and Graham Long, with Compass Resource Management in Vancouver, Canada. I live in Vancouver but go back and forth a lot between Oregon and British Columbia.

As a start to thinking about communicating effectively about uncertainty -- as a decision analyst type and somebody who does a lot of work on decision making -- I always say let's start by understanding the decision context: What's the choice that's being addressed, who are the decision makers, who are the participating stakeholders, what is the history of the issues, how is the problem structured, what really are the key objectives? Is there some agreement among the major players, the major stakeholders, on what matters? How can the fundamental objectives be expressed and measured? We've just heard in the previous session the importance of having measurable indicators -- how can you compare the consequences of actions, what's the timeframe? And then always there are constraints, and I think these are really important and can say a lot about your entry point to a problem. If there's a lack of trust, for example, then often there's a lot of upfront work one needs to do before even starting to hear the real story about what's going on. You certainly can't arrive on the scene and say, "Tell me what's going on, tell me the major sources of uncertainty." There's a lot of preamble before you can get to that. And I think it's also very important to think about *how* decisions are made, in addition to *what* decisions are being made.

These considerations all result in decisions being tough to address. And it's especially difficult for people to think about decisions where there is uncertainty and this uncertainty is present across multiple objectives. Almost every problem you guys look at and almost every problem I have worked on is characterized by having multiple dimensions, multiple objectives, with uncertainty about the outcomes of different management actions. Unfortunately, many people tend to anchor on one objective. In fact, some stakeholders think that's their job -- to anchor on one objective, so that they are seen as the economics person or the environment person or the tribal interest person. These people end up anchoring on a specific alternative, one that does a good job of addressing their particular concern. And so, a lot of the work that I think needs to be done upfront on problems as a way to begin to get to some of the key sources of uncertainty is basically opening up and un-anchoring those people so they see the problem more broadly, and a lot of discussion, a lot of deliberation needs to go on to get there.

I think it's also important to be humble about what we can do. In part this is because of external factors that are really significant, so even on major problems where we might want to do a whole lot we're often very constrained in terms of the changes that are possible. The actual levers that can be controlled as water managers or foresters or fisheries experts are often pretty limited because of the external factors that are present, which may sometimes be political, sometimes climate change, or other concerns.

Another point, second to last in this slide, is that it's very important to integrate the more intuitive and more thoughtful modes of decision making. These are often called S-1 and S-2, based on work that a number of psychologists have been doing. Daniel Kahneman, in his excellent book *Thinking, Fast and Slow*, provides a number of great descriptions of automatic and more thoughtful modes of thinking. We all react basically using two modes. One is more automatic, quicker, more emotional—that's the S-1. The other is more thoughtful and deliberate, cognitive. I think we often make mistakes in terms of trying to deal with uncertainty and getting people to address uncertainty by focusing solely on the more cognitive, rational side, because uncertainty also brings up a lot of feelings and emotions -- the more automatic response side of things. I'll talk a bit more about this distinction later on.

It's pretty obvious that decisions that involve risking and uncertainty are typically complex. They involve a lot more than science. There's the reasoned side, to be sure, but there's also emotion, there's power, there's politics, risk perceptions, culture, psychology, the role of the media -- a lot of different factors, all of which form important parts of the universe in which we all are trying to work. And in the interest of de-mystification I'll briefly mention a couple of books that have been really important in terms of helping me to think about things: *Smart Choices*, which introduces decision analysis approaches, *The Feeling of Risk* by Paul Slovic, who knows more than anyone about risk perceptions, and this *Thinking, Fast and Slow* by Danny Kahneman, which has turned out to be an incredible bestseller and has sold two million copies, which nobody would ever have thought. possible And then a book that I had a hand in, *Structured Decision Making*, that borrows heavily from these other books and outlines an approach to thinking about decisions characterized by multiple interests, multiple participants, conflicting information, and uncertainty. So back to uncertainty and how to communicate it and work with it. My first consulting jobs were back in 1974, working on salmon management in British Columbia. We always started with good intentions and came up with a sensible plan but then we would hit the implementation side, and that always runs us into people and what hopefully is not a brick wall of uncertainty about the key facts.

As this next slide shows, there are many different sources of uncertainty. One is natural variation, just the way the world works – I don't need to go over that with you. You all understand that. Measurement error is another source of uncertainty: it's hard to measure many of the things that we really care about. Often, there's also uncertainty in our models. And many times our judgments are biased, anchored on our most recent experiences or biased by our experiences in the past (although of course the future is also informed by our experiences in the past). So it's a nice tension there between biases, which may prevent us from seeing something important, and experiences, which may allow us to see something important. And often, there is linguistic uncertainty or vagueness, a lack of specificity, so that you have two people talking with each other but they don't really fully understand what each other is saying – an effect is said to be small or large, unlikely or probable – although there's a lot of head nodding and agreement as if they did.

Some of these sources of uncertainty are within the control of managers, some not. And I think part of the problem is that we all see things differently. Different people will have different opinions about what's going on -- the famous blind men and the elephant story comes to mind. Different people also will have different opinions about what matters to them. So a bottom line is that when you bring uncertainty into a multi-stakeholder and multi-objective deliberative context, good scientific analysis is necessary -- and I imagine many people in the room have predominantly scientific training, in the natural sciences -- but it's not sufficient.

One of the building blocks of my research is that I think we don't pay enough attention to the communication of uncertainty with respect to how it informs environmental management decisions. I think we often believe we know more about future effects than we actually do, which has the effect that uncertainty is minimized as part of communcations. Part of it's overconfidence: we're supposed to be experts, and so we want to act as if we know a lot and as a result we are often overconfident in our own knowledge. And this is particularly so when we have models that churn out numbers. Those numbers can be very seductive. We can easily believe numbers that we know we really shouldn't. So there's often an emphasis, I think too much emphasis, on the results of computer studies and models rather than how well people actually understand those outputs and how well people actually understand and incorporate uncertainty in their choices. And that can lead to the result that less scientific stakeholders will end up feeling marginalized, which isn't something we want. It also can lead to a discouragement of dialogue. So, to communicate uncertainty effectively, I think one of the many bottom lines is that experts need to provide clear assessments of what they do and don't know. But in addition, I think scientists need to open up to the fact that uncertainty is difficult for laypeople, and it's often difficult for experts as well. Whenever we ignore or understate uncertainty, I think our end users will lack important information and often they too will be overconfident. As a result, the options or alternatives they select may be inferior. If – on the other hand -- we overstate uncertainty, the again end users may ignore or overweight some alternatives and, again, make poor choices.

There is a need to find a balance between completeness and understandability. You don't need to present the minutia of uncertainty on absolutely every variable or every parameter that you're looking at. But you do need to present uncertainty effectively on the key drivers, the key things. That goes back from the biology side back to that first slide, where we were talking about objectives or what matters. If there are important biological drivers that people don't recognize, then work with these stakeholders and help them to see that those things are important. You need to find that balance between the completeness of the picture of uncertainty you're presenting and how understandable it is -- not overly complicated, but it has to be sufficiently detailed so that you can make defensible decisions and basically present uncertainty at an appropriate level of detail.

The next slide demonstrates that there are a lot of different expressions of uncertainty. Sometimes we get binary responses, a yes or no – will some action reduce emissions, will something increase temperatures? Sometimes we're working with verbal expressions of uncertainty or evaluative labels, which characterize uncertainty as being low, medium, or high or characterize a probability as slight, a remote chance, or highly probable. Using verbal expressions introduce a lot of room for miscommunication and so we often turn to numerical probabilities or frequencies and odds, or perhaps percentages. Another common way to express uncertainty is by using ranges, either two-point ranges—you just have a high & low estimate—or a three-point range, which shows the high/low plus a best estimate. Many of you I think, may also work with box whisker diagrams or box plots that show a low and a high estimate and have that nice little box from the 25 to 75 percent range on the probability distribution and then also show a mean or median. Occasionally, depending on the problem and the context for discussions, we can work with full probability distributions. And then there's a variety of other approaches, belief functions and so on -- lots and lots of different ways to talk about uncertainty. Obviously, it's a very important choice -- which approach or which combination of approaches is suitable given the audience and given the nature of the problem.

One other thing that I want to emphasize here (and we can come back to it later) is that it is not obvious how to best link participants with uncertainty expressions. I've worked with technical audiences who only wanted to work with verbal probabilities and evaluative labels, and I've worked with tribal audiences, tribal elders who wanted to work with probability distributions and belief functions. So the warning is to beware of making uninformed assumptions about the training or preferences of the audience. I think it's up to us as people involved in these decision-making processes, whether we take half an hour or an hour or half a day, to understand what people know and then sometimes to conduct some additional training in how to express uncertainty. It's amazing how in almost every group I've worked with the participants are at least as bright as I am and at least as quick to catch on to ideas. Yet I think that there is a tendency to assume perhaps too much of technical audiences, scientific audiences, and perhaps going the other way to assume too little of some public audiences, and I think the onus is on us to not make that mistake.

A bunch of the research I've been doing lately involves wondering whether people will differ in how they interpret ranges based on, for example, their worldviews & their motivations, or based on their numeric abilities -- their ability to make sense of numbers. I also have been wondering – along with Nate, Ellen, and others -- will some of these factors translate into systematic differences between members of the public and experts, so that when we present uncertainty in a way that makes sense to us and to our peers, will it also be interpreted in that same way by members of the public.

Another really important question is whether the presentation of uncertainty influences how people perceive the underlying probability distribution for outcomes: Is it a fairly flat distribution, is it a nice normal distribution, is it skewed to one side? This can make a huge difference in terms of the actions we choose, and particularly if we're working in a context like endangered species it might make a big difference in terms of where some of the key thresholds are and whether we might end up with some catastrophic outcome, which not only would be bad for practicing biologists but might also be against the law.

Some of our recent experiments also involve the use of more precise methods for communicating uncertainty, such as box plots. In this example you have a 90 percent confidence interval showing the extreme outcomes, a plot in the middle showing the 25 to 75 percent range, and a point somewhere showing the mean or median. This relates to the previous discussion about how expert predictions of uncertainty often display overconfidence, and here the point is that in many cases the outliers really matter -- the uncertainty may look as though it is well characterized when perhaps it isn’t. So just by showing the complete distribution, as here, emphasizes that there may be a lot of action and a lot of surprises in either of those tails. I think we often tend to interpret tails, perhaps overly so, in the light of past events. And I think we have to keep in mind about thresholds, going back to the endangered species context. If we've been trying everything with an endangered species and nothing has been working, maybe we want to take a shot at doing something new because we're really hoping to get it correct, as shown by the upside tail on the right of the graph. Or, maybe we can't try the same thing any longer because we're too worried that we're going to hit the downside tail on the left and wipe out this species.

What we do in this circumstance depends on our understanding of the uncertain factual information but also on our value judgments. Certainly if you have a really good understanding of the uncertainty, it can help you to justify your decisions and to make a defensible argument to peers or to decision makers. Coming back to the point that experts tend to be overconfident and often don’t give sufficient attention to the tails of the distribution, when working with experts, I am a big believer in undertaking formal expert judgment elicitations. This is in contrast to what I think a lot of us tend to do, which is simply kicking around ideas with peers, perhaps putting assumptions down on paper but maybe not. In such cases working more formally with expert judgment techniques, which are well documented in the literature—I’m sure many of you are pretty familiar with this work —can help expose fundamental assumptions regarding how a problem is thought about. And I have yet to engage in a formal expert judgment elicitation process and not have the experts who I'm working with be surprised at some of the outcomes. Often, I'll be working with people who have spent 30 or 40 years as peers. They may have gone to the same grad schools. They probably worked together on many of the same studies. If you get down to a level of detail where you are really trying to expose assumptions, I have found that always, always, these individuals have something to learn from each other because to some degree they will view problems differently. And remember that expert elicitations are not designed to show who's smart and who's dumb. They are designed to enhance the learning that we can all provide, to share perspectives on a problem and to thereby become better informed. Too often, simply because we're experts, we tend to think we aren't going to be as prone to anchoring on past events or as easily swayed by the most recent information or other kinds of things that are talked about in the literature as judgmental heuristics or judgmental biases. And that's simply not true. There's nothing in our training as experts that prohibits us from also being affected by these very human judgmental biases and tendencies. I also find that experts often are skeptical of expert judgment elicitations because of a fear that it's going to encourage or foster disagreements among a team of people who need to work together. And again, I think it's just the opposite: the use of formal expert judgment techniques helps establish a better shared understanding of the problem and that really helps to form and emphasizes the bond between people.

These next slides illustrate two expert judgment problems. With both slides you have the probabilities on the Y axis and results from four different experts, over two different problems. The one on the left—same experts, different problem—I would say is quite a well-structured problem. We've done a good job of specifying what is it we need to look at and there's a pretty good level of understanding among these four individuals. If we do something, if we plant conifers, will 1,000 salmon come back to this river -- whatever the problem is. So on the left-hand side, you can see there's quite a bit of agreement and these guys are saying, well, there's fairly close agreement among the four of us in terms of the probability associated with this event happening.

The slide on the right tells quite a different story. Here you see one expert saying, “Who the heck knows what's going to happen?” and his or her high and low assessments cover basically the whole probability range. The event could occur with a probability of 0.1 or it could show a probability of 1.0. So this person is saying I don't know, we don't know, nobody knows. Expert 2 is saying I've got it nailed: There's a very high probability that this event is actually going to take place (e.g., the fish will come back, the output will exceed xxx tons) and I'm very certain of that. Expert three is saying “I'm also pretty sure what will take place” but notice there's no overlap at all in the predicted distributions of Experts 2 and 3. In this particular case, Expert 2 and Expert 3 thought before we did this elicitation that they viewed the problem the same way. Only after we took this complex problem and decomposed it were these two experts able to see that they fundamentally disagreed on the interpretation of this issue, so much so that there is no overlap at all in their distributions. And Expert 4 is saying he also is pretty unsure, but it's probably not going to work out great -- the probability is lower than 0.5 that the desired outcome is going to take place.

For a problem that looks like the one on the left-hand side, I'd say it is pretty well structured – the uncertainties are low enough and (at least for this part of the puzzle) we probably know enough to go forward with a proposed action. However, when the result of an expert judgment elicitation looks like it does on the right, I'd say: Stop. You may have defined the problem wrong. Or there may be more dialogue that needs to take place among your experts. Either way, you are not ready to proceed. Okay. It's time for a cartoon: it's hard to think about consequences before they happen, which is the situation we nearly always find ourselves to be in. This is a lunchtime talk; I probably should have more cartoons.

The main research goal of the recent NSF awards -- which I think led to some of the papers that Ian and Tom saw and liked – is to explore the ability of ranges, evaluative labels, box plots, and other expressions of uncertainty to increase comprehension and informed incorporation of uncertainty. We're asking questions about what's happening in terms of peoples' understanding of the underlying uncertainty distributions --will there be some systematic differences between expert and public? And will the weight of a concern vary with changes in the uncertainty format? So I'll really quickly go through the results of a couple of recent studies. There are papers that summarize these results. If anybody wants, I can send things to you or at least tell you where these papers can be found. For this first study we had 367 public participants, randomly drawn from a web panel, and compared their responses to those of 67 experts drawn from U.S. Fish and Wildlife Service. The experts were older and more educated, no surprises here. We wanted to study how people draw meaning from different uncertainty expressions. We chose only three: numerical ranges, evaluative labels, and a combined condition. We provided a scenario dealing with vegetation management and presented information to respondents, both experts and public, in terms of a consequence table that showed the implications for key objectives of three different environmental management alternatives. To conduct the experiment we varied only the middle objective, comparing different ways to present the uncertainty associated with how the alternatives were predicted to affect the resident moose population.

What we found was that there were significant differences in how the public and expert participants interpreted this information. You can see results for the lay group on the left, the expert group on the right. In this case, the question was: In which option do the scientists have the least amount of confidence in the estimated saved moose population? So all we were asking was getting people in the study to read the information in the consequence table correctly -- this is just an accuracy test. The expert group nailed all three presentations. But you can see the lay group had a lot of trouble for this question: for the evaluative labels and combined uncertainty presentations they did pretty well, but it’s clear that the range information was really hard for them to interpret correctly.

A second comprehension question asked: for which alternative are increases in the moose population most likely? As you see here, the experts again did pretty well, although not as well with the labels or combined formats. But the lay group in this second test question did best when uncertainty was expressed using ranges whereas for the previous question they did worse with the range format. So this very simple manipulation has led to some very different interpretations and could result in quite different choices between experts and members of the public based on the different uncertainty presentations and whether people were most sensitive to uncertainty presentations using verbal labels or numeric ranges. That's troubling because we had thought that we could present things in a clear enough way so that this systematic split wouldn't happen -- and as often occurs with experiments, things didn't work out the way we thought they would.

The moral of this first study is that different audiences may well focus on different elements of uncertainty formats, and because of that, they may reach different conclusions about what management plans should be implemented. So they may go for different choices not because of a values difference, not because they want something different, but because they're reading the information you give them about uncertainty in different ways.

We later ran another study where we were wondering about to what extent our choices are value consistent. Here again we manipulated the uncertainty format, looking at a combined condition (numerical range plus evaluative labels) and contrasting this to numerical ranges only. We again used a consequence table format and found, as in study one, that the presentation using evaluative labels for describing uncertainty was favored by the lay group. But we also found that labels produced more value inconsistent choices. We used the results of an earlier study, where we asked people what mattered most to them, and found that manipulations didn't make much difference using uncertainty ranges. But when using labels to describe uncertainty we found that we could move people around much more easily and provided value-inconsistent choices. In other words, it's a dangerous tool because the evaluative labels, using words to characterize uncertainty, has such power over how lay people construct their choices.

Here's an application of some of these concepts, using the case study example of recovery planning for endangered Upper Columbia River white sturgeon (based on the results of a study I led with Graham Long and several BC Hydro colleagues). Here we were working with a technical expert group that disagreed a lot about what could be done to help these Upper Columbia River sturgeon, which are heading to extinction due to a failure to recruit successfully. There probably are some people in the room who worked on this problem between U.S. Fish and Wildlife Service and Department of Fisheries and Oceans -- let me know afterwards. Our work involved a technical group of about 25 people. We developed influence diagrams here to help characterize their hypothesis pathways, with a focus -- the big issue around recruitment failure for Upper Columbia white sturgeon -- having to do with several large hydroelectric dams that were built and then how they have changed water flows, turbidity, predator-prey relationships, and other factors. It’s an important problem to both Canadian and U.S. managers, and it's also mixed up with industrial effluent from Teck Mining Corporation. A lot of very good thinking has gone into the problem for over a decade but if you look at a graph of what's happening with recruitment for Upper Columbia River white sturgeon it's just going down, down, down, a very sad story. So my colleague Graham Long and I were brought in to say, “Well, is there anything in what we are doing with thinking about uncertainty that might help organize your various hypotheses and alternatives?”

We looked at some of the different causes, all the things that might be affecting recruitments since the dams went in. And then using an SDM approach (structured decision making techniques) we organized the group into thinking about different hypotheses by getting a “science court” underway where participants would argue for or against different hypotheses. And we tried to reach a common understanding within the group about which of these hypotheses best addressed the uncertainties associated with reasons for recruitment failure.

This is shown by the influence diagram on the slide. Each of about 200 different actions could be mapped onto the diagram in terms of the different steps along the way, ending up with the recruitment failure on the right but starting off with the primary causes, dam, etc. on the left. There were three key questions that we asked of the experts. First, what percent of ongoing recruitment failure is attributed to this hypothesis? We asked them to distribute points to show what percentage of recruitment failure they thought belonged to this hypothesis and then, second, How certain are you in your assessment -- so this directly gets to the confidence and knowledge level of the experts, based on whether they might be wrong by plus or minus various percentages. And then the third question: How likely is it that further research could confirm that this hypothesis accounts for at least 20 percent of ongoing recruitment failure? Through interactions with the group we already had reduced the 200 hypotheses to about 40 hypotheses, and this additional three-question process allowed us to start to prioritize the remaining hypotheses so that the group could get on with a small set of near-term and longer-term actions in the face of this really extreme biological uncertainty and in the face of a species which was going towards extinction fast.

We ended up with this depiction of the problem, which I think is a helpful way to look at it. The left axis, the Y axis, has to do with the certainty associated with the hypotheses, and the bottom (X) axis has to do with the importance rating of the hypothesis. It’s worth noting here that neither Graham nor I knew much about Upper Columbia white sturgeon, all we did was help guide the discussions among experts and help them to organize the hypotheses and address the associated uncertainty. So this figure shows that any hypothesis that's in the upper-right corner means that the experts thought that it was quite important and that it had a pretty high degree of certainty in their opinion. So to me, those are the prime suspects but these are really a priority for action, not so much for research but for action. Whereas down at the bottom, the "likely guilty" bunch of hypotheses are the things that are really important but there's not much certainty associated with them. So if you have research dollars, I'd pour them into the bottom right. If you have action dollars, I'd pour them into the upper right. In contrast, the hypotheses grouped on the lower left, which are low in importance and low in uncertainty, well if they're low in importance, why pour a lot of money into them given that you have scarce funds?

Think of this as just kind of a fancy bookshelf. Somebody spent $30,000 or $40,000 on this project in order to come up with a good bookshelf so that we could bin the different projects. But I believe the process really helped to distinguish the various actions and helped move the overall project forward in terms of what actually should be done rather than spending year after year—and I say this with great respect—but year after year after year, meeting after meeting with everybody just saying it's a complex world. It is a complex world, and that's why I think some kind of binning or organizing to address uncertainty directly is helpful.

With this in mind, we also sent all of you an internet survey about uncertainty before this meeting. It focused on three major topics. All I have at this point are preliminary analyses: there are perhaps 200 people in the conference and these analyses were done after we had about 75 responses. I thank you very, very much, all those of you who completed the survey. And all those of you who didn't do the survey I know had very important and good reasons why you didn't do it.

The survey focused it on a comparison of different expressions for expressing uncertainty: formats included no uncertainty, a three-point range, and a two-point range plus evaluative label. We looked at the effects of different ways of expressing uncertainty and the importance given to objectives, and then we also asked you some descriptive questions – how important is uncertainty is in your work with different audiences, etc.

Here is a quick summary of some of what you told us. The left-hand side of this slide looks at different ways of characterizing the uncertainty – no uncertainty, low uncertainty, and high uncertainty – and then we show (as you may remember) the example with three different management options, where flows would either be 10 cubic meters per second or 20cms or 40cms. And it's interesting if you look at the no uncertainty condition, the middle option under no uncertainty, and you can see that it really wasn't very popular. When there was no uncertainty, two-thirds of you selected the 40 cms option, whereas under high uncertainty conditions that result completely flipped around. There you have 60 percent of the people (under high uncertainty) selecting the mid-range, and about a quarter of you were saying okay, let's go for the 40 cms option. So uncertainty was a big driver of your choices. It could be that if there are ways of addressing the high uncertainty, such as what we just talked about for white sturgeon, which would encourage you to change your preferred management option. The bottom line, or one of them is that as the level of uncertainty changes, your preferred management option also changes. It makes perfect sense. But, what happens there, the difference between no uncertainty and low uncertainty, and focusing on how much the choice of options changes, the question is: How do we responsibly characterize uncertainty as being one or the other while keeping in mind the very dramatic effect this has on what we end up doing as managers?

The second scenario shows the same kind of situation, although here it's a less dramatic change but the middle alternative again becomes more popular in the high uncertainty condition. Notice that option A is much less popular in the high uncertainty condition: it goes from about half, 50 percent, of you selecting that option to about a 25 percent selection rate.

That is a big change. And mix C is gradually building in terms of the percentage of you selecting it as the associated uncertainty increases.

So does how we characterize uncertainty have an effect on our choices, as managers or as public stakeholders? The answer is “Yes”. And the message is that if there are ways of working with uncertainty so we understand it better, ways that are likely to reduce the level of uncertainty associated with an action, then that's likely to change our management choice.

One more result from the survey you completed. We also asked you about the extent to which you talk about uncertainty when working with different audiences. Here are the results. With technical experts, about 40 percent of you said you talk about uncertainty more than 75 percent of the time, whereas with officials and other decision makers, only about 10 percent of you said you routinely discuss uncertainty with those groups more than 75 percent of the time. One of the messages I got from this is that we should be doing more training of the public and of officials and decision makers so that they can better understand uncertainty and then we would bring that into the picture more often with them than we do now. We also asked you about the formats you used for presenting uncertainty. Evaluative labels seem to be your favorite format for presenting uncertainty, followed by three-point intervals—showing extremes, high/low, plus a best estimate, and then the box plot, certainty equivalents, and so on. As an aside, I'm a fan of certainty equivalents and that may be something to talk about more since it seems that some of you are using them already. Interesting results here. If you combine these last next slides, I think it really suggests there's a lot of deep thinking to do about how uncertainty is being presented.

The last quick thing I'll say is a reminder that all these methods have advantages and disadvantages. I think evaluative labels can really help laypeople evaluate the meaning of uncertainty expressions and highlight important aspects of uncertainty, but people may end up relying on the format and not thinking as much as they should. If you provide evaluative labels, people may feel that you and other experts have done the thinking for them, and that is part of why evaluative labels can override people's own values so they end up making value-inconsistent choices. Numeric ranges have some real advantages and they are commonly understood, although the notion of a confidence interval is not. And I think we've got careful expert judgment elicitations that can help to do a good job on numeric ranges. But again they can really change how people view the distribution underlying the ranges.

In current experiments we're getting some really interesting results about how ranges encourage people to think about end points. In some cases providing a range encourages people to see a normal distribution; they interpret the range as a flat distribution rather than a more normal, bell-shaped distribution. I think that's quite worrying. Graphic displays can help but the danger is that with more complicated presentations, such as box plots, people may just tune out and ignore uncertainty entirely. So if you don't do the proper sort of upfront presentations, or training, something like a box plot may appear too complicated and then people will deal with the situation as if there were no uncertainty even though you're doing your best to represent the uncertainty because it is important. Thus there can be a big distinction between presenting uncertainty and actually getting people to think in probabilistic terms and use that information as part of the choices they make.

Using multiple uncertainty formats, that's the common advice—the logic being that if there are all these problems, then just throw out a bunch of different methods and people will make sense of it all —but we’re finding that this won't work because different people are going to interpret those different methods differently. Unfortunately, it appears that no single representation of uncertainty is going to work best in all contexts. That sounds wishy-washy, and I'm sorry, but that's true. Both numeric and nonnumeric evaluative structures are important and simple graphical representations can be great, but the different formats really need to be explained. I'm a big advocate for training and helping lay or expert participants to understand probabilities. And keep in mind that nobody needs more help with making decisions than decision makers, so I think we need to see that as part of our job also.

Another easy area where improvements can be made is that I think we need to be clear about the aspects of uncertainty that we want to make most salient. Not all objectives, not everything matters, so as managers we should go after those two or three things that really matter and bang them home in terms of making uncertainty important. Think about how communication might go wrong as well as how it might go right. Be humble: there is a lot more to learn. And this is one of my favorite uncertainty slides – another cartoon. Uncertainty continues to fascinate and interest us. The caption is that "these once-in-a-lifetime events are beginning to add up." As a person in his early sixties, I know the meaning of that slide. All right, time for questions: thank you very,much.

IAN: We are running a little bit late, but we can entertain one question. You'll have to high grade among yourselves the best question.

ROBIN: I'll stay right here afterwards in case there are others.

IAN: Anyone. All right, there we go.

SPEAKER 1: Okay. So Robin, I'm really curious about the elicitation and interpretation of uncertainty. But I'm wondering, how much does risk tolerance of the person you're getting the information, how do you think that plays into it?

ROBIN: It plays a big part, and I think it should. I think risk tolerance is important. I think with expert judgment elicitations, when you're working with the experts, you want to try to separate it out and understand what contributes to risk tolerance. Because often, there is a level of risk tolerance defined by legislation or by their agency, the group that someone is working with.

If there is really a lot of risk tolerance, often it's related to a contributing factor that can be pulled out in terms of the decomposition of the problem. So I think that understanding risk tolerance can transforms a single expert judgment elicitation into two or three because we're looking at those essential conditionalizing assumptions, the individual's risk tolerance as compared to the risk tolerance of the agency or the legislation under which they're working.