

Review of Holt, R.F. and G. Sutherland 2003. Central coast coarse filter ecosystem trends risk assessment: Base case. Draft, November 2003.

Review submitted by
Judith Anderson
6008 Leibly Avenue
Vancouver, BC V5E 3C8
email: judithanderson@shaw.ca or janderso@sfu.ca

I think the analysis presented in this draft is very valuable. The ecological information presented is impressive in its breadth, and I constantly found myself asking, "How could anyone even think of managing these forests without this sort of information? Yet, only a few years ago, this kind of spatial analysis was in its infancy. Now policy makers can visualize policy decisions as they play out on a generous spread of spatial and temporal information about specific ecosystems and locations.

Most of my suggestions for change have to do with the presentation of the information and the utility of the concept of risk as it's presented here.

Tightening up the discussion of risk

If we're to talk about "risk" at all in science, it's important to remind the reader about the nuts and bolts of "risk". Risk is defined as "exposure to the chance of injury or loss", and in quantitative analysis, it is calculated as the product of the (negative) *value of the loss* times the *probability* that that loss will occur (Morgan and Henrion 1990). Because not everyone views a particular outcome as bad, we have to know exactly who will experience the outcome as a loss, and the negative value assigned it must also be figured from that party's point of view (or pocketbook). In addition, the bad outcome in question must be clearly defined and identifiable in practice. The analyses of these components (value of the outcome from a particular stakeholder's point of view, and its probability of occurrence) should be kept rigorously separate. That's the only way they can be used meaningfully in a Bayesian context, whether it is a decision analysis or belief network.

The ETRA (Holt and Sutherland 2003b) defines the risk of interest as "that the coarse filter functions will be less likely to be maintained, and that species/processes/ecosystems will eventually be lost or degraded." That definition suggests the two components of risk, but it could be sharpened up. For example, we need to know how we would recognize the bad outcome when it occurs, who would perceive it as a loss and how we'd estimate that loss. In addition, the multiplicative relationship with probability should be made clear. This multiplicative relationship seems intuitive, but it has at least one intuitively disturbing property -- a large loss that occurs with small probability can produce the same calculated risk (expected loss) as a small loss that occurs with large probability. Examples of what this peculiarity might look like in this context would help analysts decide whether "risk" is really what they want to analyze. It may not be.

The risk components are, to my mind, never teased apart in the study. Instead, the ETRA focuses on the behaviour of an assumed indicator of the risk, the difference between the amount of old forest simulated under management and the amount of old forest simulated under a natural disturbance regime. Let's call that variable OFDiff. As an indicator of the bad outcome we're concerned about, OFDiff seems reasonable, and I would rather see the quantity being

analyzed referred to consistently as "OFDiff" or some other empirically descriptive name, rather than label it "risk" in text and tables. "Risk" may be the conclusion of the study, but you want readers to focus on the real variables and what is being done with them as long as possible before directing their attention to "risk". My comments on the psychology of risk explain why.

If decision makers are really to be able to interpret OFDiff's behaviour as a "risk", they need to know at least the *assumed* relationships between OFDiff and the probability of occurrence of one or more definite, specified, recognizable (bad) ecosystem outcomes. Empirical relationships between OFDiff and the risk of concern would be even better of course, but the data probably do not exist.

Some day, however, the relationships should be checked empirically, and when that is done, the probability of the bad ecosystem outcome will have to be represented by its *frequency* of occurrence. That is why it's essential to specify exactly what the bad outcome of interest is, and what its boundaries are -- frequencies are counts, and you can't count cases of the bad outcome unless you can identify and bound a case and say that it has definitely occurred.

Especially in the absence of those empirical relationships between OFDiff and the risk of concern, there are other disadvantages to renaming OFDiff "risk". OFDiff is already a highly derived, simulated variable -- why add another layer of assumptions to it? (See the discussion of tricky variables below). Also, if OFDiff is interpreted as risk, the dimensions of "risk" can't be the old dimensions of OFDiff (percent old forest). The dimensions of risk have to be related to the negative *value* (to someone) of the bad outcome -- yet another layer of assumptions if the data about value are not solid.

An additional problem with labeling the analysis results as "risk" is that "risk" apparently increases as a function of the absolute value of OFDiff, in both directions (ETRA Fig. 2). However, it seems unlikely that these opposite-oriented risks can represent the same bad outcome, much less the same causal mechanism. As percent old forest decreases from the "natural" expected amount, we're probably looking at progressive habitat loss for old forest-dependent species, for example. In contrast, as old forest increases beyond the "natural" expected amount, the damage to the ecosystem may come from loss of early seral stages or buildup of fuels. Lumping these very different processes together as "risk" seems unhelpful to me.

Here is my recommendation, which will come up throughout this review: Use the word "risk" as little as possible in the report. The analysis and graphs are most easily understood if they present variables as close as possible to the concrete ones from the real world and from the simulations. As a final step, the assumed relationships between OFDiff and probabilities of specific bad ecosystem outcomes could be presented, to transform the conclusions about OFDiff into reasonable speculations about "risk". But save that for last, after the analysis has been digested by the reader.

Verbal categories of risk

Verbal descriptors of risk carry a heavy psychological load (see below for a bit of discussion), but the same category boundaries (0-20%, 21-40%, and so on) seem quite useful to summarize OFDiff without any implicit assumptions about what they mean in terms of risk. In any case, the categories are arbitrary; in the absence of clear ecological thresholds, it's reasonable that they should divide the range of values evenly.

As for exploring other sets of boundaries, these divisions of the range of OFDiff are not really a "hypothesis" as suggested on p. 3 of the Executive Summary. They are a human creation;

they are neither true nor false. Also, alternative ways of categorizing OFDiff are labeled more or less or "cautious" in Figure 6. I'm not sure that's a useful label. The category boundaries are not really "cautious" themselves, nor should moving them in one direction or the other imply a cautious policy decision. Policy is the decision makers' bailiwick, and policy can be uncautious, even in the face of high risk -- "Damn the torpedoes" is a policy decision.

Tricky variables

One problem with highly-derived variables is that quantities that are apparently closely related often do not code in the same direction. So the graphs on pp. 15-18 have "percent old forest" on the y-axis, where high values mostly represent healthy ecosystems. In contrast, the risk classes for the summaries (Table 3) use the metric of percent deviation from expected natural, so high values represent ecosystems at risk. A similar directionality problem has led to confusion in the captions of Table 1 (see "Specific small problems" below).

Reversals of coding are not the only problem with these variables. Difference variables (e.g., OFDiff), and ratios (e.g. disturbance rates) are fraught with dangers, we all learned in Statistics 101. Because these types of variables are derived from two numbers, they bring with them doubled error and complicated error structures. Frequently these error structures produce subtle biases (two of these biases were discussed in the report).

Of course, the number crunching for this analysis requires complex variables. However, in the report for decision makers, it's probably better to present variables in a simple form wherever possible and try to keep the meanings of high and low values consistent. And at all stages, sharp vigilance is needed to avoid errors with difference variables and ratios, such as inadvertently reversing the direction of a derived variable or, worse, one of its components.

Specific small problems with the analysis and report.

Figure 1 seems to be repeated. The equivalent figure in the North Coast report (Holt & Sutherland 2003a) shows both AUs and BEC variants.

The figure of predicted percent old forest as a function of disturbance interval needs to be numbered.

Figure 2. The "(bad) outcome" from going over the expected natural distribution will be different in character and causal factors from the outcome if you fail to achieve that distribution. They are both labeled just "risk" -- that's not helpful.

Several of the figures (Figure 4 and the graphs on pp. 15-18) would be more understandable as X-Y time series graphs, not bar graphs. The low and high RONV benchmarks could be represented as horizontal lines going across the width of the graph. The resulting graph would be closer to the "risk model" presented in Figure 2. Moreover, that's also how this sort of simulation result is presented in some of the graphs in Williams and Buell (2003).

On the other hand, Figure 3, which reports frequencies, is fine as a bar graph.

The graphs into the future should indicate the increasing uncertainty of the prediction, at least in the downward direction (I assume the maximum expected old forest abundance is determined by the current age structure, to the extent that it is known).

Table 1. The text mostly talks about disturbance frequencies or rates; the labels for columns 1, 3, and 4 refer to disturbance intervals or return intervals, but the highest and lowest have been reversed (columns 3 & 4). Moreover, it seems that the two intervals from Banner are used to produce only one "lowest likely mean percent old forest" -- so was only Banner's highest frequency (lowest interval) actually used? Or does column 5 use both Banner numbers?

(Apropos, the percents in columns 2 and 5 should be whole numbers unless they all represent quantities less than 1%).

Table 2. It seems a small point, but in normal categorization, you can't have overlapping category boundaries. Therefore, we should see 0-20, 21-40, 41-60, etc. Also, the reader needs to be reminded what units these numbers represent (it seems to be OFDiff; if that's correct, it should be stated explicitly). As for sensitivity analysis on these category boundaries, I'm not sure it's helpful, since they are arbitrary to begin with and their meaning is purely subjective, and it's not clear how one would recognize a "better" or "worse" set of boundaries even if the distribution of ecosystems in the classes did change. Varying the boundaries to produce fuzzy categories, as reported in Holt and Sutherland 2003a p. 25, gets into the sticky area of fuzzy logic; it seems unnecessary to go there.

Table 3 has a couple of confusing headings. "Base risk values over time" would be better described (unless I'm mistaken) as "Percent difference between projected old forest coverage under base case management and projected old forest coverage under natural disturbance" (which I'm calling OFDiff). The next set of columns translate these numbers into risk categories anyway. "Confidence in outcome" is perhaps not the best heading in the context of a base case analysis that includes a Bayesian Belief Network. There is a lot more that should influence someone's confidence in the outcome than just the two known biases. How about something like this: "Known bias present (direction of bias on OFDiff)". You'd have to reverse the Ys and Ns of course, and use the phrase "biased the OFDiff outcome" rather than "affected the risk outcome" in the caption.

In all cases, it's helpful to distinguish empirical, current variables from expected/predicted/simulated versions of the same variables. See the discussion of accounting rules below.

Psychology of risk

Heuristics and biases. Because of human psychology, discussions of "risk" have a hard time remaining scientific. There is a huge literature in traditional psychology describing how people fail to reason logically about risks and uncertainty, instead relying on emotions, heuristics, and biases when formulating judgments about uncertain situations (e.g., Slovic 2000, Tversky and Kahneman 1974). To avoid these well-documented problems, it's best to keep the scientific discussion focused as long as possible on the separated components of the risk at hand (the bad outcome, its causes, and its costs to someone specific, and the estimate of its probability of occurrence). You might even find it useful to discuss the science of good outcomes, as well. Wait until all that is worked through before introducing the combined bad outcome-probability entity as a "risk".

Evolved mental modules. Psychologists who study human behaviour from an evolutionary perspective assert that our minds consist of a variety of interconnected, evolved mental "modules" designed to efficiently process information relevant to particular adaptive problems that our ancestors faced. For example, people easily and automatically detect cheating by others in a situation of social exchange (bartering, for example), but they can find it quite difficult to detect whether someone is violating an arbitrary, abstract rule that has no social context (Cosmides and Tooby 1997).

Because there is no such thing as a "general risk", evolutionary psychology suggests that natural selection has never produced a mental adaptation to process "risks" in the abstract. Evolutionary psychologists, in fact, explain people's "illogical" response to risk by

hypothesizing that different mental modules are activated by different kinds of uncertain threats. For example, the chance of a life-threatening illness in your own child (close kin) and the chance of an airplane accident involving your adult co-workers (unrelated group members) will be processed by different modules, even though they are both seriously bad events. Even if some utility analysis estimated the losses from those two events to be the same, and even if their probabilities were identical, it is very likely that the mental modules processing them would produce different outputs -- emotions, plans for action, etc. Therefore, the psychology of particular uncertain bad outcomes does not interact very well with generalized, objective analysis of "risks", despite efforts to make the analysis understandable for people.

Unfortunately, verbal categories of risk are among the seemingly sensible, well-intended efforts that are thwarted by human psychology. Patt and Schrag (2003) provide an interesting example of how verbal categories for probabilities can represent quite different numerical probability ranges, depending on the characteristics of the bad outcome. In particular, a costly, community-destroying bad outcome (such as a hurricane) will be considered "high risk" at a lower probability of occurrence than a more benign alarm such as an early snow storm.

Importance of frequencies. So far, I don't see much discussion of numerical probabilities in the ETRA, and that's appropriate, given that it's probably impossible to estimate the probabilities on the relevant hypotheses to even one significant digit. The verbal categories respect that problem nicely. However, I would guess that at some point, people will want to know what range of numerical probabilities the categories represent. In addition, people might want to know more detail about the Bayesian Belief Network. If numerical discussion of probabilities is needed, translating them into frequencies seems to make them more intuitively understandable (Cosmides and Tooby 1996, Gigerenzer 2000, Anderson 1998, Anderson 2001).

The psychological importance of stories and scenarios. Stories and case studies are perhaps the most effective means for conveying information to people and mobilizing knowledge when there is a problem to solve. That is not just because stories are fun; it is because expertise consists of (1) having experienced many cases of a problem and (2) being reminded of relevant cases or stories from one's experience when faced with a new problem (Riesbeck and Schank 1989). It seems to me that the great strength of this study is the depth and breadth of the scenarios it can generate, as we see in Williams and Buell (2003) and Holt and Sutherland (2003a).

The Bayesian Belief Network

From this draft, I don't have a clear picture of where the Bayesian Belief Network fits in and how its interface with SELES is set up. Holt and Sutherland (2003a) describe pros and cons of using a Bayesian Belief Network analysis, but its function would be much clearer if we had a more informative diagram than Figure 2 of Holt and Sutherland (2003a). Exactly what hypotheses form the Bayesian Belief Network, where do the priors and updating information (if any) come from, where do we see the posterior probability distributions on the hypotheses? In particular, what is the Bayesian Belief Network giving us that a parametric sensitivity analysis or Monte Carlo simulation could not give?

For example, there doesn't seem to be any indication of the probability distribution of the results graphed in Section 5.2 of Holt and Sutherland (2003a), yet it seems they are outputs from the Bayesian Belief Network. When the results are presented as point estimates, don't we lose a lot of the advantage of propagating uncertainty through the Bayesian Belief Network analysis, especially as the predictions extend far into the future?

Nor is there any indication of the probability distribution on the base risk or risk categories reported in Table 3 of Holt and Sutherland (2003b). The risk categories reported in this table don't seem to consistently reflect the values in Table 2; it's not clear to me how the Bayesian Belief Network has massaged the base risk values to produce a distribution and pick out a most probable risk category for the table. Why, for example, is CedarHigh CWHvm2 assigned VH risk at all times by the Bayesian Belief Network when its base risk values vary from 52 to 61?

Perhaps the analysis of one ecosystem could be presented from beginning to end as a demonstration of the process.

Concerns about the Bayesian methodology

Empirical support for Bayesian methods. If the Bayesian Belief Network really is returning probabilities on various hypotheses in this analysis, we need to remember that Bayesian procedures have little empirical backing in applied ecology. Once we're outside the realms of casinos, epidemiology, and weather reporting, what assurance do we have that ecological events given a Bayesian posterior probability of, say, 0.3, actually do occur 30 times out of every 100 opportunities? I don't think this should discourage people from using Bayesian methods, but it should always be emphasized that they are based on theory, and have not been empirically validated.

Bayesian analysis is not the only game in town when it comes to analyzing uncertainty. Monte Carlo simulation requires a large number of iterations based on probability distributions for model parameters; I would guess that is not possible given the complexity of these simulations. Parametric sensitivity analysis is a structured sensitivity analysis, which involves systematically choosing high and low values for important parameters in addition to the central estimate; it might be more feasible in this situation and provide just as useful information as Bayesian or Monte Carlo analysis. Scenario analysis, discussed below, seems particularly useful for this context.

The bimodal distribution in risk. The majority of ecosystems analyzed in the ETRA fall into the extreme "risk" categories, or, as I translate to myself, at extreme values of OFDiff: either the old forest distribution resembles what you'd expect under a natural disturbance regime, or it is very different. The common-sense explanation offered for that distribution, that these extremes reflect different management histories, seems useful and informative.

My concern is that this strongly bimodal distribution may make a Bayesian analysis on these ecosystems problematic. The reason is that Bayesian posterior probabilities calculated on extreme priors can be highly unreliable if the prior probability is not known with certainty, which is usually the case. Let's suppose we're considering only two hypotheses, H, and not-H. The table below shows the posterior probabilities on a hypothesis H, calculated from Bayes' Theorem as a function of the prior probability on H and the diagnosticity of the data at hand. (Data are "diagnostic" when the probability $\Pr(D|H)$ of observing them is close to one or zero, if H is true.)

Prior probability on H, $\Pr(H)$	Probability of data when H is true, $\Pr(D H)$				
	0.05	0.10	0.50	0.90	0.95
0.02	0.00	0.00	0.02	0.16	0.28
0.05	0.00	0.01	0.05	0.32	0.50
0.10	0.01	0.01	0.10	0.50	0.68

0.20	0.01	0.03	0.20	0.69	0.83
0.30	0.02	0.05	0.30	0.79	0.89
0.40	0.03	0.07	0.40	0.86	0.93
0.50	0.05	0.10	0.50	0.90	0.95
0.60	0.07	0.14	0.60	0.93	0.97
0.70	0.11	0.21	0.70	0.95	0.98
0.80	0.17	0.31	0.80	0.97	0.99
0.90	0.32	0.50	0.90	0.99	0.99
0.95	0.50	0.68	0.95	0.99	1.00
0.98	0.72	0.84	0.98	1.00	1.00

First, let's see what happens if the data are *not* diagnostic. The middle column with $\Pr(D|H) = 0.5$ (non-diagnostic data) shows that the posterior probabilities are identical to the prior values. They are now labeled "posterior" probabilities, but clearly not much has been learned, which is what you would expect from the uninformative data.

Now, let's look at the effects of highly diagnostic data in the first column, where $\Pr(D|H) = 0.05$. (That is borderline where a classical statistics test would reject the hypothesis H.) Where the prior and the data point in the same direction (prior = 0.02 to 0.4), the posterior probability is low and is relatively insensitive to the exact value of the prior. In the range where the prior is neutral or mildly supportive of H (prior = 0.5 to 0.6), the posterior is still low and rather insensitive to the exact value of the prior. However, when the prior probability is in the range of 0.7 to 0.98, contradicting the data, the posterior probability is *very sensitive* to the prior; the posterior could be anywhere from 0.11 to 0.72. If we translate this result to an evenly divided 5-point verbal scale, we would conclude that, when the prior on H is considered "high to very high" *and* the data are improbable given H, then the posterior could be anywhere from "very low" to "high", sliding from one verbal category to the next with only small changes in assumptions about the prior.

Column 2, $\Pr(D|H) = 0.10$, shows a similar pattern with less strongly diagnostic data. Again, the posterior probability changes dramatically as a function of the prior when the prior is above 0.7.

Columns 4 and 5 show the corresponding problem when the extreme values are at the other end of the distribution. As in columns 1 and 2, diagnostic data (here, supporting H), combined with a prior pointing in the opposite direction (here, prior between 0.02 and 0.4) produce highly unstable posteriors.

Schnute and Hillborn (1993) corroborate this problem for cases where multiple hypotheses are being considered. When data diverge strongly from a prior probability distribution, Bayesian analysis results in an intermediate distribution of posterior probabilities that can be difficult to reconcile with either input distribution and difficult to interpret or use.

In plain language: if you think you understand the system (you hold an extreme prior) and you get new data that surprise you (data are diagnostic in the opposite direction), you can't say anything about the probability(ies) of the hypothesis(es) you're studying. You need to think about the science and gather more data. This seems obvious, but could easily be overlooked if the data are being manipulated by software that just delivers you a nice tidy 3-significant-digit posterior probability (or firm verbal probability level) on H. Bayesian analysis won't make a silk purse out of preliminary science or inadequate data.

Now, in the ETRA context, I'm not sure exactly what information the Bayesian Belief Network is using, what hypotheses it updates probabilities for, or what numerical probability ranges each verbal risk category represents. Therefore it is possible that this concern about applying Bayesian analysis to bimodal distributions is misplaced. But I believe it should be kept in mind as a general precaution about Bayesian analysis, even if it will not cause problems in this particular application.

Sources of uncertainty not discussed

The ETRA is billed as a base case study. Presumably we'll have a report of alternative cases in the future, so I can see why there is not a lot of discussion about external sources of uncertainty (as opposed to uncertainties about data used directly in the analysis). Nonetheless, I think it would be helpful to forestall the concerns of readers who immediately think of big uncertainties they would like to see explored, by briefly mentioning important scenarios that will be looked at or have been covered by Williams and Buell 2003.

What about obvious sources of uncertainty, such as climate change, further trouble with export markets, changes in the exchange rate with the US dollar, the possible devolution in BC of public management and forest ownership to private management and/or ownership? These don't seem to be discussed explicitly (though they may be implicit in some of the Williams and Buell 2003 scenarios). All the issues just mentioned contribute to uncertainty, but they are not really stochastic, so a Bayesian probabilistic analysis won't help as much as political and economic analysis.

Climate change in particular is a social problem for which a Bayesian approach has limited utility. The political will needed to prevent or adapt to climate change is clearly in the realm of social science. One might imagine that a Bayesian Belief Network would be a good model for beliefs about the causes of climate change and the effectiveness of remedies, but in fact the process of changing scientific beliefs also seems not to be well-described by Bayesian updating. Instead, it is better understood as an process of cultural evolution (Blackmore 1999) depending critically on the social arrangements of scientists working within the field (Hull 1988, 2001).

With respect to ecosystem health, one could argue that climate change would raise the risk category of all these forest ecosystems by at least one or two levels, even if there is no more sophisticated way to include its effects in the analysis.

Presenting 200-year simulations to skeptics.

Let me start by saying that I think the focus on old forests is very wise, because of all the variables that could represent forest ecosystem health, old forest is most likely to be unaffected by unforeseen events -- the trees that will be old forest in 200 years are mostly already in place and measurable. However, other variables could be quite vulnerable to the concerns I'm about to express, and as a general principle, I feel very unconfident in predictions of anything 200 years from now; even 50 years is a stretch.

First, if we ran the LP model starting in 1800 with the then-current forest practices and ecosystem distributions, how well would it have predicted the current 2003 distribution? Would we do much better starting in 1880, 1900, or 1920? From what starting date would such a simulation have predicted current outcomes with any accuracy? Without that information, I suggest that there should be large and increasing error bars around any predicted trajectories for the ecosystem indicator.

Second, this sort of long-term simulation is basically a form of accounting. You keep track of variables over time and sum them over years, report on their final status, or report them as you go along. In accounting, there are rules about how far removed from reality a predicted asset can be and still be added together with real, current assets. Consider the following definitions (from <http://www.investorwords.com>).

Asset: Any item of economic value owned by an individual or corporation, especially that which could be converted to cash. Examples are cash, securities, accounts receivable, inventory, office equipment, a house, a car, and other property.

Accounts receivable. Money which is owed to a company by a customer for products and services provided on credit. This is treated as a current asset on a balance sheet. A specific sale is generally only treated as an account receivable after the customer is sent an invoice.

In other words, assets are mostly real, current values. Your permission to dip into the future is strictly limited -- accounts receivable can be assets, but only if they are pretty close to being realized, after an agreement has been reached with a customer and an invoice has been sent. Developments that may bring in cash in the more distant future may have their place in a company's accounts, but you can't call them "assets". You have to call them something else and report them in a different column.

I think long-term simulations need a similar sort of accounting rule. In the case of economic forecasting, such as Fig. 3.5 in Williams and Buell (2003), it should be made clear that the y-axis for most of the graph represents a different entity than the y-axis for the current decade, which is surely the limit of any forest company's accounts receivable. The current numbers might be called "current revenue" and the more distant numbers might be called "modeled revenue". Or perhaps we might use the terms "real revenue" and "imaginary revenue" (in parallel with real and imaginary numbers in math -- both are useful, but they can't be added together or share the same axis).

In the case of forest simulations, patches forecast on the basis of currently existing trees resemble accounts receivable. The "invoice has been sent": We know what to expect from them, barring disturbances that are not included in the model. In contrast, the characteristics of a patch forecast from regeneration after a cut 30 years from now should be reported in a different column: the "invoice" for that patch has not yet been sent. Again, labels should reflect this accounting structure -- "real" vs. "modeled" or "imaginary" forest, or some such. You can report both columns, of course, but if we accept the rules of accounting as a guide, they shouldn't be combined or reported on the same y-axis; they are different entities.

Other issues in communication of uncertainty

While it may be hard to adequately convey all the uncertainty in these long-term forecasts of forest condition, projecting out to 200 years does serve a powerful social purpose -- helping stakeholders to feel a commitment to the forest ecosystem in the long run. So I guess I would like to see these projections presented in a way that brings to life the social context of that future, linked back to the people who are having to make that tough, counterintuitive commitment now. Possible techniques:

- on graphs, include the past 200 years (with appropriate error bars) as well as 200 years into the future (with appropriate error bars).
- accompany the graphical information with a narrative about the past and the future. The future story could have several variations, and these could be linked to different trajectories drawn on the graph, within a general graphed area of possible future states. The future part of these graphs would look something like the multiple scenario graphs presented in the EGSA (except, as I suggest above, you'd need to redraw and relabel the y-axis beyond some point in the near future).
- Natural selection has given humans mental adaptations that help them to care for the long-term wellbeing of their children and their grandchildren. Commitment to the long term may therefore be facilitated if we remind people that their direct descendants will depend on the forests of the future. Our ancestors probably lived to see and help nurture their grandchildren but not great-grandchildren, so it is no surprise that, beyond grandchildren, concern and caring tend to get a bit hazy. To keep the links to future generations psychologically alive, one might therefore refer to the descendants four generations from now as "your grandchildren's grandchildren".
- First Nations members of the policy group could undoubtedly provide many suggestions for linking the future to the past and present.

Should this risk assessment morph into a scenario analysis?

Many of my concerns about the ETRA have to do with the risk-related concepts and methods presented in this draft. They need more development and better explanation. However, even if those improvements are made, I'm not convinced that focusing on risk will produce better policy.

On the other hand, I think the greatest strength of this study is the potential for exploring a variety of long-term, spatially-explicit scenarios. As long as we keep in mind the concerns about 200-year predictions and treat the scenarios as stories, scenarios are probably the best way to convey the information to decision makers and give them a visceral sense of the processes involved. Scenario analysis is an alternative or adjunct to risk assessment -- there is no definitive boundary between risk analysis and scenario analysis. See for example <http://www.usf.uni-kassel.de/usf/scenarios/documents/paperguidelines.pdf>. Risk analysis tends to focus on the quantitative aspects of the problem -- the losses associated with bad outcomes, and their probabilities of occurrence. Scenario analysis, focusing on narratives of different ways the future might unfold, feeds comfortably into human psychology and, most important, into the style of cognition that makes people experts. I think this study could profitably be presented as a scenario analysis, and the risk aspects could be de-emphasized.

References

- *Anderson, J. L. 1998. Embracing uncertainty: The interface of Bayesian statistics and cognitive psychology. *Conservation Ecology* 2(1):2. <http://www.consecol.org/vol2/iss1/art2>.
- *Anderson, J. L. 2001. Stone-age minds at work on 21st century conservation science: how cognitive psychology can inform conservation biology. *Conservation Biology in Practice* 2:18-25.
- *Blackmore, S. 1999. The memes' eye view. Pages 25-42 in R. Aunger, editor. *Darwinizing culture: the status of memetics as a science*. Oxford University Press, Oxford, United Kingdom.

- *Cosmides, L. & Tooby, J. (1996). Are humans good intuitive statisticians after all?: Rethinking some conclusions of the literature on judgment under uncertainty. *Cognition*, 58, 1-73.
- *Cosmides, L. & Tooby, J. (1997). The multimodular nature of human intelligence. In A. Schiebellek & J. W. Schopf (Eds.), *Origin and evolution of intelligence*. Center for the Study of the Evolution and Origin of Life, UCLA. (pp. 71-101).
- *Gigerenzer, G. 2000. Adaptive thinking: rationality in the real world. Oxford University Press, Oxford, United Kingdom.
- *Holt, R.F. and G. Sutherland 2003a. Environmental risk assessment: Base case. Coarse filter biodiversity. Final report. North Coast LRMP.
- *Holt, R.F. and G. Sutherland 2003b. Central coast coarse filter ecosystem trends risk assessment: Base case. Draft.
- *Hull, D. L. 1988. Science as a process: An evolutionary account of the social and conceptual development of science. University of Chicago Press, Chicago.
- *Hull, D. L. 2001. Science and selection: Essays on biological evolution and the philosophy of science. Cambridge University Press, Cambridge, United Kingdom.
- *Morgan, M. G. and M. Henrion. 1990. Uncertainty: a guide to dealing with uncertainty in quantitative risk and policy analysis. Cambridge University Press, Cambridge, United Kingdom.
- *Patt, A. G. and D. P. Schrag. 2003. Using specific language to describe risk and probability. *Climatic Change* 61: 17-30.
- *Riesbeck, C. K., and R. C. Schank. 1989. Inside case-based reasoning. Erlbaum, Hillsdale, New Jersey.
- *Schnute, J. and R. Hilborn. 1993. Analysis of contradictory data sources in fish stock assessment. *Can. J. Fish. Aquat. Sci.* 50:1916-1923.
- *Slovic, P. 2000. Perceived risk, trust, and democracy. Pages 499-513 in T. Connolly, H. R. Arkes, and K. R. Hammond, editors. *Judgment and decision making: an interdisciplinary reader*. Cambridge University Press, Cambridge, United Kingdom.
- *Tversky, A., and D. Kahneman. 1974. Judgment under uncertainty: heuristics and biases. *Science* 185:1124-1131.
- *Williams, D. and M. Buell. 2003. Economic gain spatial analysis -- Timber. CIT Central Coast Region. Draft.