

The Beijer Institute of Ecological Economics

DISCUSSION PAPER

Beijer Discussion Paper Series No. 173

Coping With Uncertainty: A Call for a New Science-Policy Forum

Ann Kinzig, David Starrett, Kenneth Arrow, Sara Aniyar, Bert Bolin, Partha Dasgupta, Paul Ehrlich, Carl Folke, Michael Hanemann, Geoff Heal, Michael Hoel, Bengt-Owe Jansson, AnnMari Jansson, Nils Kautsky, Simon Levin, Jane Lubchenco, Karl-Göran Mäler, Stephen W. Pacala, Stephen H. Schneider, Domenico Siniscalco, and Brian Walker. 2003.

Coping With Uncertainty: A Call for a New Science-Policy Forum

Ann Kinzig, David Starrett, and others¹

Introduction

Throughout their history, humans have devised a number of approaches and strategies to reducing (perceived) uncertainty or the psychological impacts of uncertainty, and to function and thrive in a partially unknown world. A certain world begets certain 'best strategies' for survival; an uncertain world necessitates a larger portfolio of options and approaches, and means for choosing among them.

We cannot hope to cover the whole of human history here. Instead, we focus on the ways in which uncertainty and approaches to reducing uncertainty have played out in recent history in two arenas—science and policy. We focus on these because it is the interplay between these that has such relevance for global environmental futures, and because the characterizations of uncertainty, and strategies for coping with it, differ so dramatically between the two.

The scientific process is built on the goal of advancing knowledge—penetrating and reducing the reaches of what is not known. Each advance is built on knowledge acquired earlier. The cost of 'incorrect knowledge' is therefore quite high, affecting not only that building block in the foundation but those that follow. Science has therefore evolved procedures whose primary goals are to protect against being wrong.

The policy process is built on the goal of addressing societal ills or challenges. Timeliness is frequently of the essence—action may have to precede knowledge. The errors that need be avoided are those associated with undue political or social costs, including harm to constituents, the economy, national security, or the environment.

Because of these fundamentally different goals, science and policy have evolved significantly different 'evidentiary standards'. Scientists apply relatively high evidentiary standards—insisting, for instance, that there be only a small probability (sometimes 5% but frequently lower than 1%) that an incorrect conclusion has been drawn. This means that many conclusions concerning the workings of the world that have a reasonable probability of being correct—say 75%—may not pass muster.

The policy process employs both stricter and looser evidentiary standards, largely based on the (perceived) costs of being wrong. If we are certain a patient is going to die shortly, for instance, there is little hazard in prescribing a drug whose efficacy is largely

¹ The ideas in this paper were developed at a Beijer Institute workshop at Askö marine laboratory Aug. 31 - Sep. 2 2002. Participants in the meeting were Kenneth Arrow, Sara Aniyar, Bert Bolin, Partha Dasgupta, Paul Ehrlich, Carl Folke, Michael Hanemann, Geoff Heal, Michael Hoel, Bengt-Owe Jansson, AnnMari Jansson, Nils Kautsky, Simon Levin, Jane Lubchenco, Karl-Göran Mäler, Stephen W. Pacala, Stephen H. Schneider, Domenico Siniscalco, and Brian Walker

unknown, but offers some hope of life extension. Here the ‘unknown’ interacts with our perceptions of risk in such a way as to allow us to employ relatively loose evidentiary standards. A drastically different evidentiary standard should apply for instance, in the case of cosmetics, where society insists on a vanishingly small chance that one might die from the application of mascara. The manufacturer must demonstrate the highest confidence that the product works as intended and is very unlikely to have undesirable side effects.

Difficulties arise due to these different standards in four ways. The first is merely failing to recognize that different standards in policy and science exist; this can result in a failure to communicate properly. Policymakers may believe that scientists are using a stricter or, more likely, looser evidentiary standard than is actually being employed. Thus, an outcome that a scientist believes is highly certain may be (mis)interpreted as only probable or likely in policy circles.

The second difficulty arises because scientists may fail to provide information that could be considered useful or even crucial in the policy arena, in large part because they may have failed to study the appropriate phenomena. The success of scientists depends on their ability to advance their field under the relatively strict evidentiary standards that apply. This selects for fairly reductionist approaches to studying phenomena whose drivers can be tightly controlled and manipulated. Until recently, this meant largely avoiding study of precisely the kinds of environmental systems that society is influencing most profoundly—global climate systems, entire ecosystems or landscapes. Fortunately, our capacity to study these complex, and often highly uncertain, systems has advanced tremendously, but the strong scientific tradition of applying strict evidentiary standards can still delay the introduction of information that could be considered crucial by policymakers.

The third problem arises because scientists may find it difficult to quantify the uncertainty associated with highly uncertain—and to some extent unknowable—futures. This has led, in some circles, to an avoidance of characterizing uncertainty at all. Witness, for example, the recent emissions scenarios of the IPCC, each of which was presented without an assessment of likelihood, and each of which was thus treated as equally likely by the public and policymakers. No rational assessment of probabilities of the various scenarios would conclude equal likelihood; rather such a process would conclude (potentially drastically) different likelihoods for each scenario, albeit with very low confidence (Schneider, 2002). We will discuss the need for likelihood assessment and relevant methodology in the next section.

The fourth problem has to do with the ways in which the neutral language of evidentiary standards can be used to mask what is ultimately a debate over values. This tactic is not limited to the (supposedly subjective) policymaker, but is employed by the (supposedly objective) scientist as well. Most if not all of the environmental problems faced by society are ‘complex all the way down’. They are rife with uncertainty and there can be numerous plausible solutions, leading to numerous possible futures. These futures will differ in terms of environmental protection, social justice, economic growth, and political freedoms, among other things. Scientists can help illuminate those futures and trade-offs, but they are no more ‘expert’ than any other citizen when choosing among the trade-offs. There is nothing objective about valuing environmental protection over political freedom, or economic growth over social justice, and no set of evidentiary standards can ultimately allow society to

navigate among these complex considerations. Unfortunately, they are often employed as if they can. Scientists advocate certain policies, frequently claiming a scientific objectivity when they do, but their positions ultimately derive from their own values. Policymakers invoke scientific uncertainty as the justification for inaction on certain issues, when ultimately their positions have more to do with the weight of economic and environmental considerations. In what follows, we will attempt to sort out the roles of objective versus subjective information and values. [ref Functowitz and Ravetz]

In the next section, we articulate a set of principles for dealing with uncertainty, give a series of examples to illustrate these, and show what sorts of errors will arise when science and policy fail to coordinate correctly.

Principles and Examples

How should scientists present information to the public and to decision makers? And in particular, how should information be presented when there are substantial uncertainties concerning what is known (or even knowable)? In discussing these questions, we shall invoke two well established principles: **Expected Utility theory** for making decisions at a point of time in the face of uncertainty, and **Bayesian updating** for revising our measures of uncertainty in the presence of new information. According to expected utility theory, for any action we might take we must first identify each of the possible outcomes (consequences) it might engender and determine the relative likelihood of each consequence. We then assign utilities to consequences in a way that reflects relative desirability of outcomes and our attitudes toward risk (more on this later). We then take the action that maximizes expected utility. Probabilities are determined using the best available information. When new information becomes available we use it to update employing Bayes' rule. If the new information makes an outcome more likely to be observed than it would have seemed given ex ante information, the probability of that outcome is increased using the formulas of conditional probability, and vice versa (Morgan and Henrion, 1990).

Neither of these principles is particularly intuitive to the average citizen but there is a long literature defending them against plausible alternatives². We will not survey that literature here although some of the supporting arguments will be implicit in our treatment of the subject.

Expected Utility

To illustrate the principle of expected utility and show how it impacts the way in which scientists should conduct and report their work, let us see how it requires us to optimally trade off type I versus type II error. Consider a stylized example involving whether or not to do a mastectomy. Suppose there are two states of the world, (a) in which the outcome generates a higher utility $U(a)$ than doing nothing and (b) in which it generates lower utility $U(b)$. Normalizing so that the status quo gets zero utility, we have $U(a) > 0 >$

² For a relatively non-technical introduction to the theory of choice under uncertainty, see Raiffa (1968). For an introduction to Bayesian methods, see Lee (1997). More detailed and comprehensive treatments of these subjects can be found in Berger (1985), Pratt et al (1995) and Press (2003).

$U(b) = -C(b)$ where $C(b)$ represents the costs, both monetary and psychological, of performing an unnecessary or undesirable procedure. Assume further that we know the probabilities of these states: $P(a)$, $P(b) = 1 - P(a)$. If we do the procedure we incur the danger of type I error (accepting the procedure when it is unjustified) which would happen with probability $P(b)$. But if we forgo the procedure we experience type II error with probability $P(a)$.

According to the principle of expected utility we should conduct the mastectomy if and only if the **expected** utility of proceeding exceeds zero (the expected utility of doing nothing). Here the rule would be to proceed when $P(a)U(a) > P(b)C(b)$. Observe that this rule trades off the two types of error, taking into account the benefits of being right relative to the costs of being wrong. In this regard, it is extremely important that scientists and practitioners see the rewards and benefits in the same way as the society being served. For example, if doctors have an inordinate fear of malpractice claims while the medical costs of procedures are covered by insurance, they may be induced to conduct procedures unless $P(a)$ is extremely small [$P(b)$ is extremely large], in order to protect themselves. Thus, they may conduct even when $P(a)U(a) < P(b)C(b)$. In the extreme they would see only the benefits of desirable procedures and not the costs of unnecessary ones, and would conduct any time that $P(a) > 0$. Clearly society, which must incur these costs, loses on balance.

In other areas of science we may see a bias in the opposite direction. As indicated earlier, scientists tend to impose high evidentiary standards, and frequently will not report results or recommend action unless $P(a)$ is very large, e.g..95 [$P(b)$ is very small]. Policy makers, relying on this may then fail to act when $P(a) < .95$, yet $P(a)U(a) > P(b)C(b)$, and we incur the reverse bias. Although it is probably too early to tell now, inaction on the consequences of global warming may be attributable to failures of scientist to take strong enough positions for fear of possibly being wrong. Indeed, this example suggests that reliance on any arbitrary “confidence level” is inimical to good decision-making; rather scientists should report results with whatever levels of confidence can be justified.

Of course, scientists must still impose a severe level of standard before accepting a new theory into the corpus of scientific knowledge due to the substantial costs associated with building on a false scientific base. Thus, all practitioners should get used to the idea that different standards are appropriate depending on context.

Subjective Probability

Perhaps the most controversial aspect of our operative principles involves the use of subjective probability. The concept of probability is easiest to explain and defend in situations where repeated trials can be performed. For example, most would agree that a die is “fair” if after a large number of repeated rolls, each of the six numbers comes up approximately one-sixth of the time. We are then comfortable in saying that the likelihood or **probability** of any particular number on a given throw is one sixth. In this regard, the insurance industry uses data on frequencies of illness, death, accident, property damage, etc. within a given demographic category to determine the likelihood that any individual in that category will be affected, and offers protection against the bad outcome using these numbers to determine premiums. Given their ability to spread risk over many people, they can offer coverage at nearly “fair” odds (based on the frequencies) and risk-averse individuals find it

advantageous to purchase coverage. This sort of inference and use of probability from frequency is well established and not very controversial.

However, when repeated trials are infeasible (as they often are in real world situations), it is not clear how probabilities could or should be assigned and used. For example, consider the event that nuclear bomb-making material falls into the hands of terrorists. Since the opportunity has only been around for a relatively short time and this outcome has never been documented, there is no way to infer likelihood from the principle of repeated trials. Yet, expected utility theory would require that we assign a “best guess” probability to this event.

The logical underpinnings of subjective probability as applied to decision-making are similar to the logic of value comparisons in welfare economics. Economists have long argued that one must consider tradeoff rates between the value of marketed goods (as measured in money units) and such intangibles as (for example) a statistical life. This comparison is necessitated by the fact that policy makers must make decisions on “projects” that have some monetary benefits but incur some risk of loss of life. Whatever decision is made clearly will have some implications concerning the implicit tradeoff used, and consistency across decisions dictates that the same tradeoff rate be used whenever the conflict arises. This tradeoff rate then defines the “shadow price” of a statistical life. Similar logic dictates the use of subjective probabilities. For example suppose we are considering employing a breeder reactor whose dollar net benefits have been measured as \$10 billion. We believe, however that its use will increase the possibility that bomb-making material falls into the hands of terrorists. Should that occur, we anticipate substantial costs; assume for sake of argument that we can measure these and they come out to \$1 trillion. Still a decision cannot be made without some judgment concerning likelihood and aversion to risk. Concerning risk, assume either risk neutrality or that the numbers above are already in utility units (we will return to issues involving attitudes toward risk later). Then if we go ahead with the project we are implicitly saying that the likelihood is increased by less than one chance in 100, whereas if we forgo we are making the reverse judgment. Again, consistency across decisions requires that we use the same number in all similar comparisons and this number is what we call the subjective probability.

There are many difficulties in measuring subjective probabilities and again, these parallel similar difficulties in measuring shadow prices. In some instances we may be able to learn something about relevant citizen’s attitudes from observed behavior. For example, people’s willingness to build bomb shelters during the cold war tells us something about the perceived likelihood of a nuclear exchange. Or the wage premium on life-threatening jobs should tell us something about the value of a statistical life. Unfortunately such “hedonic” methods rarely if ever allow us to exactly identify a shadow price or subjective probability (Starrett, 1988, Chapter 13), and in any case can only reflect the quality of information possessed by citizens. The bottom line is that decision makers still need to use best judgment without a firm objective basis in specifying shadow prices and subjective probabilities. Scientists therefore have an obligation to provide the highest possible quality of information concerning relative likelihood of events and values of intangibles (see, for instance, Moss and Schneider, 2000).

Bayesian updating

The main way in which scientists can inform concerning probabilities (aside from repeated trials when the frequency method is available) is through the conducting of controlled experiments and tests. In the area of health, the questions might be on whether or not large doses of vitamin C helps prevent the common cold or whether a glass of red wine per day reduces the risk of heart disease. Evidence is brought to bear by comparing the experience of groups who take the treatment with a control group of those who do not. According to Bayes' law, what we learn from such evidence depends on three things: the test results, accumulated past evidence, and the inherent accuracy of such tests. Suppose that a new test does suggest a positive relationship between red wine and reduced heart disease. This evidence will be most compelling in changing our estimates of likelihood if there is little previous evidence to go on and we have reason to believe that such tests are accurate. Unfortunately, there is something of a conflict here, because if we have little previous experience with such controlled studies, we have no way to judge their inherent accuracy.

Consider now another situation where there is considerable accumulated evidence, but it is inconclusive in that some studies point one way and some the other. One might think in this circumstance that a new study would be quite informative, but this is not generally true. The fact that previous studies of this sort have been inconclusive suggests that the method is not very accurate and adding one new study showing a positive relationship should not change our prior likelihood of a relationship very much. We suggest that this situation is fairly typical of the medical literature on nutrition where conflicting studies abound. In such circumstance, medical researchers do a disservice when they make broad claims on the basis of one new study. For one thing, such a study should not alter our priors much; but more importantly, when such claims turn out to be unjustified (as they sometimes do), the public loses confidence in the general efficacy of medical research and tends to ignore warnings that are fully justified on the basis of accumulated evidence.

Functional uncertainty and surprises

Of the many types of uncertainty, among the least recognized, understood, and dealt with is something we will call **functional** uncertainty. We may be ignorant concerning the relevant state space or more generally with the functional relationships involved in the policy issue at hand. There are no general principles to invoke here so we will simply illustrate with two examples. First, think about the Georges Bank fishery. Marine economists and ecologists early worked with a 'logistics' model in which each positive fish stock would be associated with a positive sustainable level of fishing. Only later did evidence accumulate that this probably was the wrong model and that a correct version might entail a threshold whereby once the stock is allowed to fall to a critical level, the stock would collapse even if all fishing was curtailed. The fishery would be left barren, with quite dire consequences for the fishing industry. As a second example, consider the North Atlantic Thermohaline Circulation, wherein Gulf Stream warm water is channeled north in exchange for a cold water return from arctic regions. We are quite uncertain exactly how much disturbance—in the form of temperature or precipitation change—could trigger an instability, although we do know such instability exists, just not where the thresholds are (Mastrandrea and Schneider, 2001). Without this knowledge, we cannot predict confidently when human-made

forces might shut it off. We will return to this example in our discussion of irreversibility below.

In each of these cases, functional uncertainty introduces the possibility of a “surprise,” and such surprises have been a recurring feature of science throughout its history. How should scientists and policy makers deal with the possibility of surprises? Obviously, by their very nature surprises cannot be fully anticipated, yet they can have far-reaching consequences. One strategy for dealing with surprises would be **avoidance**: namely do whatever we can to minimize the occurrence of a surprise. But even if desirable, this strategy would have the fatal flaw of being non-operational. Since we can never know for sure what will trigger a surprise, we cannot know which strategies minimize the trigger. Ex post, of course, we may observe the trigger and identify the desired strategy. For example, knowing what we know now, we realize that fishing should have been restricted earlier on the Georges Bank (Holmes, 1994). However, ex ante it was not clear when or by how much this should be done. Indeed, it is even possible that allowing stocks to over-produce could promote a pathogen that provides another equally destructive surprise.

Further, we would argue that even if the avoidance strategy could be made operational, it would not be desirable. Suppose, for example we knew that due to natural ecological uncertainty, there was some small probability that a fishery would collapse in any given year, and moreover that this likelihood was an increasing function of the level of fishing. Then following the avoidance strategy would require that we never fish. But such a strategy would violate the principle of tradeoff that underlies all good decision-making: The likelihood of a “bad” surprise must be traded off against the lost benefit (economic and social) from fishing. Again, we cannot avoid making a subjective judgment concerning the likelihood of surprise and using it to compare the risks against lost benefit. More generally, extreme strategies such as “always assume the worst (or the best)” are never desirable because they ignore relevant tradeoffs. There is an emerging literature on alternative strategies to use in managing surprise, including strategies to improve anticipation of surprise and cope with inevitable surprise (see, e.g., Schneider and others, 1998; Gunderson and others, 1995; and Berkes and others, 2002).

This should not be taken, however, as a necessary argument for continuation of status quo practices. The point is to weigh relevant trade-offs. There are numerous examples of over-fishing of marine resources (Jackson and others, 2001) and there is much to learn from what has happened in Gorges Bank and other areas. In the Baltic Sea, for example, the cod spawning stock is outside of its biological limits (HELCOM, 2002) in spite of many years of scientific advances to reduce catches. Annual assessments of the declining cod spawning stock and full awareness of the fact that spawning success depends on the unpredictable inflows of salty, oxygenated water from the North Sea—a “standing surprise”—did not persuade the decisionmakers to adopt precautionary practices, though in this case they should have. Adequate information was there, but fierce lobbying by fisheries organizations and the difficulties of obtaining international consensus prevented beneficial outcomes.

Surprises and the strategy of scientific pursuit

How much should scientists “fool” with Mother Nature? Should we engage in genetic engineering? Should we try to interfere with the dynamics of hurricanes? Should we build new types of nuclear reactors? Clearly any of these strategies take us into unknown places where the possibility of surprises increases. But again there are tradeoffs. For one thing, such projects, by taking us outside known parameters provide information for resolving some functional uncertainty. Genetic experimentation, for example, can tell us things about the chemistry of life that we might not be able to discover otherwise. For another, some surprises are good—indeed many of our most important scientific discoveries (for example X-rays) have come as surprises observed because a scientist did something that had not been done before and consequently observed something unexpected. This conflict between maximizing the opportunity for learning and minimizing the likelihood of “bad” surprises is fundamental to the scientific endeavor, yet has no easy resolution.

Attitudes toward risk

One of the things the general public expects from science is the reduction or elimination of uncertainty faced. It is clear to us now that uncertainty cannot be eliminated entirely but there remains the question as to how much in the way of resources we should devote to reducing uncertainty and how the presence of residual uncertainty should be incorporated in our decision-making rules.

We can measure the aversion to residual uncertainty by asking how much a person who is facing some random outcome (for example chance of traffic accident or property damage) would be willing to pay to obtain an outcome with certainty that has the same expected value. In most circumstances people are willing to pay something (which we refer to as the **risk premium**) and that explains how the insurance companies are able to make money by spreading the risk. Risk aversion is incorporated in decision-making by using a utility function which is nonlinear (in fact concave) on numeraire (monetary) units. Concavity means that we “object more” to downside losses than we are “happy with” upside gains and the more concave our function, the stronger this effect, and the greater the degree of risk-aversion represented (see Box 1).

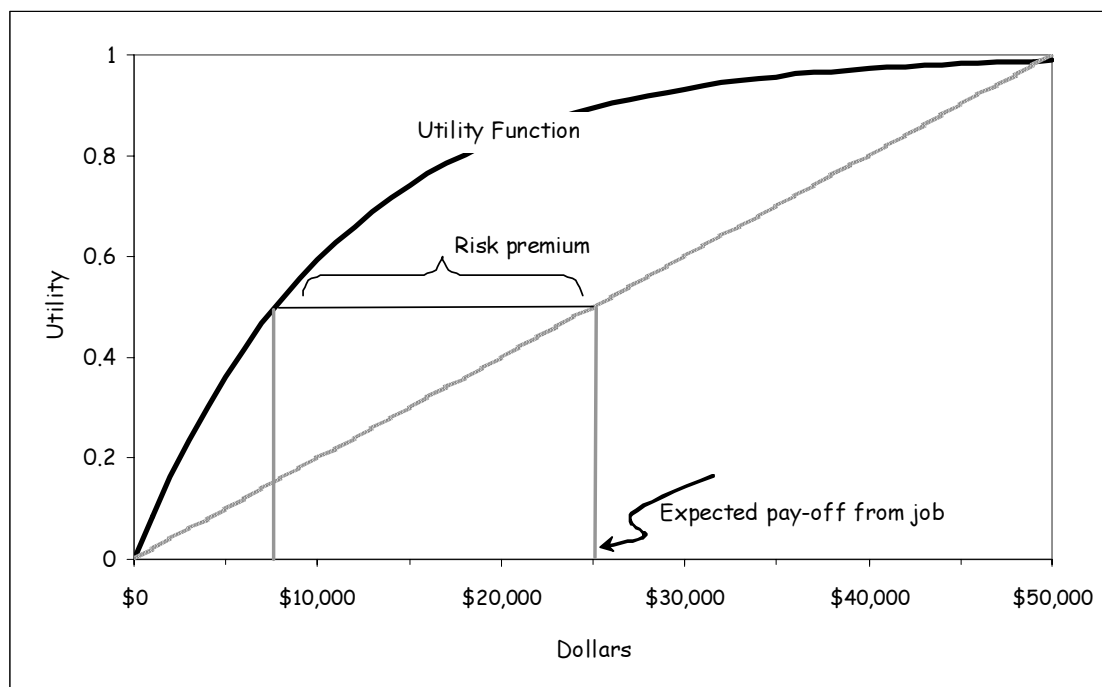
What degree of risk-aversion should be employed in public decision-making? It should be obvious that the answer depends on context. For example, in cases where risks are socially insurable we should provide the insurance whatever else is done and treat the residual problem in a risk-neutral way. However, few risks are completely insurable (whereas we can insure against the monetary costs of illness or accident, we cannot insure against the psychological costs) and some are not insurable at all. The latter include events that affect the social group as a whole such as ozone depletion or an asteroid impact, as well as the functional uncertainties discussed earlier.

When uninsurable risks are involved, we must add the degree of risk aversion to subjective probabilities and shadow values as things that need to be specified as a basis for good decision-making. The degree of risk aversion not only serves to inform current policy choices, but also gives guidance as to how much of our resources we should devote to obtaining risk reduction. In this regard, the size of risk premium (aggregated over all

members of society) measures the benefits of eliminating uncertainty, and therefore gives an upper bound on the value of resources that should be devoted to resolving the associated uncertainty. Again there is a subjective element here, since different agents will generally display different degrees of risk aversion.

Box 1: Determination of the risk premium

Consider a worker in a highly risky job. For illustrative purposes suppose half the time she will make \$50K and half the time nothing (so that her expected income is \$25K). We can determine the degree of aversion to risk by asking how much money the worker would accept for certain rather than have to face the job risk. A worker who will accept less than \$25K is said to be **risk-averse**, and the difference between \$25K and what she will accept is labeled the **risk premium**. In expected utility theory, risk aversion is captured by use of a utility function which is concave in dollars. If we normalize so that $U(50K)=1$ and $U(0)=0$, such a function looks as in the diagram below, with the expected utility of the risky job equal to 0.5. As drawn, the worker has said she would accept \$7.7K for sure [$U(7.7K)=0.5$], implying a risk premium of \$17.3K. Note that if the worker were risk-neutral (no risk premium) the implied utility function would be linear and further that the degree of risk aversion will be positively related to the degree of concavity in the utility function.



Irreversibilities and option values

One fairly universal feature of uncertainty is that the farther we look into the future, the more uncertain we are about functional relationships and outcomes. As one consequence, uncertainty is a far bigger concern for dealing with policies involving the long-term future than for those that are atemporal in nature. Emerging research suggests that the resilience of nearly all environmental systems is characterized by linked dynamics across spatial and temporal scales, including processes that play out on very long time scales (see,

for example, Gunderson and Holling, 2002). Thus, it is vital to assess long-term consequences and deal with the attendant uncertainties. A critical factor in determining how we deal with the long-run future involves **reversibility**. Certain kinds of actions/policies are reversible in that if near term consequences are not to our liking, we can take action that insulates the more distant future from these consequences. For example, in farming we may try a new fertilizer in one year. Unless this application has some long term leaching effect on the soil the action is reversible in the next year.

The most difficult situations in this regard occur when there is at least some possibility that one could never return to the status quo ante no matter what reversing policies are used. Returning to the Gulf Stream circulation example, some have suggested that global warming could decrease salinity sufficiently in mid-Atlantic waters through arctic ice melt to shut off the pump entirely (e.g., Rahmstorf 1999). Should this happen, the event would be effectively irreversible and have very unpleasant long-term consequences for Northern Europe, which could (paradoxically) enter a cooling phase while the rest of the world warmed.

How general is the presence of potential irreversibility in our decision-making world? Although the relevant science is not generally well understood, several likely examples have emerged (see, for instance, Carpenter and others, 2001; Scheffer and others, 2000). For instance, phosphorous loading in a lake can cause it to flip to a “dirty” state in a way that cannot be reversed by simply eliminating the loading. Also, it is not yet clear whether or not fisheries collapse is reversible by simply stopping the fishing. Obviously species lost to extinction can never be recovered, though more needs to be known about whether there are commensurate irreversible losses of essential ecosystem services. Clearly it is of great importance that we improve our understanding of functional uncertainties in complex systems so as to avoid irreversible effects before they occur.

When irreversibilities are present, an argument can be made for some form of a “precautionary principle.” In our example here, the potential of Gulf Stream ‘shut-down’ should make us especially cautious about overall global warming which might trigger the sufficiently decreased salinity scenario. Note also that when there are potential irreversibilities, the rate of discount takes on special importance, since the weight put on long-run irreversible costs is very sensitive to the rate of discount. Consequently, we need to be particularly careful to defend our choice of discount rate and acknowledge the value judgments inherent in that choice.

Unfortunately, there is no single strategy for dealing with irreversibilities. The precautionary principle suffers from some of the same difficulties as the “avoidance” principle discussed earlier. Since we can never be sure what action or inaction will trigger an irreversible shift, we cannot be sure of what to be precautionary about. However, when learning more about the natural environment is an option, there are extra reasons to postpone actions whose desirability will depend on the information learned. In this case, the option to act differently depending on what is learned has value that is lost if we act irreversibly too soon. The potential for global climate change presents one situation where option values are important. Global climate change is characterised by unpredictable local or even regional extreme events, but we are as yet unable to foresee where and when these may occur. Evolution will be gradual, i.e. over decades to a century and is unpredictable. This

give us time for constructive learning, but only if we do not take irreversible actions (or inactions) now. (For a detailed discussion in the context of climate change see Heal and Kristrom 2002. For a more basic review of environmental decision-making in the context of irreversibility, see Dasgupta and Heal 1979 chapter 13 section 4.)

Conclusions and Recommendations

We conclude that advances and changes must be made in the way science is conducted and uncertainty communicated. Scientists must become more effective and compelling communicators of both what is and isn't known. Politicians must bolster their ability to make decisions in the face of uncertainty and be clear about the role ideology and values play in interpreting uncertainty. They must also be willing, at times, to risk votes in order to follow societally beneficial and scientifically assessed pathways. We must develop institutions that maintain flexibility, continually reassess, and potentially change direction in the presence of possible surprises and irreversibilities. And we must advance public understanding of the nature of complex problems with inherent uncertainties, and acceptance of the fact that solutions and policies must be ever changing as systems evolve and knowledge advances. To foster these broad aims, we offer the following specific recommendations:

Recommendations

- Establish science-policy forums composed of leading figures in the science and policy communities. These forums should be constituted at the highest levels of government, within national academies of science, or within international science or policy bodies. These forums should provide input concerning the other recommendations outlined below and, further, provide oversight on the science-policy interface, offer constructive criticism on performance, and facilitate relevant communication links to the public at large.
- Overhaul undergraduate and graduate education to include courses in complex systems, decision analysis, Bayesian statistics, and public communication.
- Restructure scientific education at all levels to emphasize science as a process, and the Earth system as complex and, at times, highly uncertain in its dynamics.
- Establish short courses, workshops, and institutes through which current scientists and policymakers could be brought up to speed on methods and knowledge relevant to their roles in the decision making process.
- Educate natural-resource managers to view their systems as self-organizing, complex, and adaptive, with dynamics that play out on multiple scales of space and time, and across levels of organization. Train managers in the use of management 'experiments' that can be used to probe possible thresholds in the system, and distinguish between reducible and irreducible uncertainty.
- Develop organizations and processes for dealing with environmental and natural-resource challenges that maintain flexibility, foster reassessment, and change direction if needed.

Science is largely a publicly supported enterprise, and we must ensure that it is serving the public. We simply must have professional scientists fully engaged in the policy process. They must recognize, however, that their expertise is in elucidating probabilities and consequences. When it comes to more value-laden discussions, their voices should carry no more weight than any other informed citizen. Thus scientists need to be fully aware of their own values and subjective assessments when engaged in this process. Not all scientists need engage in such activities, but all should support those reputably engaged in the policy process, and create the conditions that allow and value that engagement. To do otherwise defaults on the debt owed the public.

Politicians already have a well-developed ability to make decisions in the face of uncertainty—they do so everyday. They have been less successful, however, in elucidating the ideological filters through which uncertainty is interpreted. Risk assessment generally cannot be divorced from ideology; the risks we are willing to tolerate, or decide we must avoid at all costs, are matters of values. Politicians should be clear about the extent to which their decisions rest on a scientific assessment of probabilities and outcomes, and to what extent they rest on a more subjective and ideologically driven interpretation of how the world works. Policymakers owe this level of honesty to their constituents, who as they advance their own understanding of the complex workings of the world, should come to demand it.

The general public thus must also become much more aware of the nature of environmental challenges and the uncertainties plaguing them. It has, to some extent, been the ignorance of the public that has allowed politicians and other policymakers to invoke scientific uncertainties as a justification for inaction. Uncertainty cuts both ways—there is no certainty or safety in inaction. Society must cultivate an educated citizenry focused on both their rights and responsibilities. The public has an absolute right to be at the table in discussions of environmental challenges; they also have a responsibility to educate themselves on the issues before taking that seat. Scientists must also recognize the right that stakeholders have to be at the table, and the necessity of a broad perspective when it comes to negotiating the values we will hold and the futures we will pursue.

References

- Berger, J. O. 1985. *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. Springer-Verlag, New York, NY.
- Berkes, F., J. Colding, and C. Folke, 2002. *Navigating Social-Ecological Systems*. Cambridge University Press.
- Carpenter, S.C, B.H. Walker, M. Anderies, and N. Abel. 2001. From metaphor to measurement: Resilience of what to what? *Ecosystems* **4**: 765-781.
- Dasgupta, P.S and G. Heal, 1979. *Economic Theory and Exhaustible Resources*, Cambridge University Press.
- Funtowicz, S.O., and J.R. Ravetz. 1991. *Uncertainty and Quality in Science for Policy, Theory and Decision Library Series A: Philosophy and Methodology of the Social Sciences 15*. Kluwer Academic Publishers, Dordrecht, The Netherlands.
- Gunderson, L. H., and C. S. Holling, eds. 2002. *Panarchy: Understanding Transformations in Human and Natural Systems*. Island Press, Washington, D.C., USA and London, England.
- Gunderson, L. H., C. S. Holling, and S. S. Light, eds. 1995. *Barriers and Bridges to the Renewal of Ecosystems and Institutions*. Columbia University Press, New York, New York, USA.
- Heal, G. and B. Kristrom, 2002 Uncertainty and climate change. *Environmental and Resource Economics* **22**: 3-39
- HELCOM 2002. *Environment of the Baltic Sea area 1994-1998*. Baltic Sea Environment Proceedings No. 82B, pp 215.
- Holmes, B. 1994. Biologists sort the lessons of fisheries collapse. *Science* **264**:1252-1253.
- Jackson, B.C., Kirby, M.X., Berger, W.H., Bjorndal, K.A., Botsford, L.W., Bourque, B.J., Bradbury, R. H., Cooke, R., Erlanson, J., Estes, J.A., Hughes, T.P., Kidwell, S., Lange, C.B., Lenihan, H.S., Pandolfi, J.M. Peterson, C.H. Steneck, R.S., Tegner, M.J. and R. Warner. 2001. Historical over-fishing and the recent collapse of coastal ecosystems. *Science* **293**:629-660.
- Lee, P. M. 1997. *Bayesian Statistics: An Introduction*. Arnold Press, London.
- Mastrandrea, M. and S.H. Schneider, 2001: Integrated assessment of abrupt climatic changes, *Climate Policy*, **1**: 433-449.
- Morgan, M.G., and M. Henrion, 1990. *Uncertainty: a Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press, New York, NY.
- Moss, R.H. and S.H. Schneider, 2000: "Uncertainties in the IPCC TAR: Recommendations to Lead Authors for More Consistent Assessment and Reporting," in *Guidance Papers on the Cross Cutting Issues of the Third Assessment Report of the IPCC*, Intergovernmental Panel on Climate Change, R. Pachauri, T. Taniguchi, and K. Tanaka (eds), Geneva, pp 31-51; available from the Global Industrial and Social Progress Research Institute: <http://www.gispri.or.jp>

- Pratt, J.W., H. Raiffa, and R. Schlaiffer, 1995: *Introduction to Statistical Decision Theory*. MIT Press, Cambridge, MA.
- Press, S.J., 2003: *Subjective and Objective Bayesian Statistics: Principles, Models and Applications*, 2nd. ed. Wiley-Interscience, Hoboken, NJ.
- Raiffa, H., 1968: *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*. Addison-Wesley, Reading, MA.
- Rahmstorf, S., 1999. Shifting sea in the greenhouse? *Nature* **399**: 523-524.
- Scheffer, M, S.C. Carpenter, J. Foley, C. Folke, and B.H. Walker, 2001. Catastrophic shifts in ecosystems. *Nature* **413**: 591-596
- Schneider, S.H., 2002a: Editorial Comment: Can we estimate the likelihood of climatic changes at 2100? *Climatic Change* **52**(4): 441-451.
- Schneider, S.H., B.L. Turner II, and H. Morehouse Garriga, 1998: Imaginable surprise in global change science. *Journal of Risk Research* **1**(2): 165-185.
- Starrett, D.A., 1988: *Foundations of Public Economics*. Cambridge University Press, New York, NY.