Risk & Sustainable Management Group

Murray Darling Program Working Paper: M05#3

Research supported by an Australian Research Council Federation Fellowship
http://www.arc.gov.au/grant_programs/discovery_federation.htm

The precautionary principle in environmental policy and the theory of choice under uncertainty

John Quiggin

Schools of Economics and Political Science
University of Queensland
Brisbane, 4072
rsmg@uq.edu.au

This version 15/4/05

The precautionary principle in environmental policy and the theory of choice under uncertainty

John Quiggin

Australian Research Council Federation Fellow

School of Economics and School of Political Science and International Studies

University of Queensland

EMAIL j.quiggin@uq.edu.au
PHONE + 61 7 3346 9646
FAX +61 7 3365 7299
http://www.uq.edu.au/economics/johnquiggin

I thank Nancy Wallace for helpful comments and criticism.

This research was supported by an Australian Research Council Federation Fellowship.

Abstract

The precautionary principle, presented as a guide to environmental policy decisions in the presence of uncertainty, has been the subject of vigorous debate. However, the has generally not been discussed in relation to formal theories of choice under uncertainty developed as generalizations of the expected utility model. In this paper, it is argued that a formal basis for the precautionary principle may be found in an incompleteness hypothesis regarding formal models of choice under uncertainty. The incompleteness hypothesis states that estimates derived from formal models of choice under uncertainty will generally be over-optimistic and that the errors will be greater, the less well-understood is the problem in question.

Keywords

precautionary principle, generalized expected utility theory

The precautionary principle in environmental policy and the theory of choice under uncertainty

Introduction

The concept of the 'precautionary principle' has been the subject of vigorous debate. As with other contested concepts in environmental theory and policy, most notably that of 'sustainability', the debate has proceeded in the absence of an agreed definition. As a starting point, it is useful to consider the definition implicit in this statement by Christine Todd Whitman, then governor of New Jersey and later Administrator of the United States Environmental Protection Agency, quoted in Appell (2001):

> Policymakers need to take a precautionary approach to environmental protection.... We must acknowledge that uncertainty is inherent in managing natural resources, recognize it is usually easier to prevent environmental damage than to repair it later, and shift the burden of proof away from those advocating protection toward those proposing an action that may be harmful.

As Whitman indicates, the precautionary principle is concerned with the formulation of choices under uncertainty. However, in the discussion of the precautionary principle, there has been only occasional reference to the literature on the theory of choice under uncertainty, a literature that spans economics, psychology and statistical decision theory. The absence of any formal framework for discussion has contributed to the confused nature of the debate, in which a multitude of definitions of the precautionary principle have been proposed and criticized.

Where formal models of choice under uncertainty have been used, it has often been observed that the use of theories that are more general than those in common use would imply the adoption of some form of precautionary principle. For example, Kinzig, Starrett et al (2002) propose the use of Bayesian decision theory in place of the classical inference model more commonly used in the assessment of statistical claims about health and other risks. Gollier, Jullien and Treich (2000) and Gollier and Treich (2003) consider the role of option value and irreversibility, as do Heal and Kristrom (2002). Bargiacchi (2003) considers the role of generalizations of expected utility theory, such as rank-dependent utility (Quiggin 1982), that allow for probability weighting.

The disparate nature of the issues considered by these authors suggests the need for a broader approach to the precautionary principle. In particular, it appears desirable to consider the issues in terms of characteristics of choice problems that can be described in general terms, rather than as parametric properties of particular models such as expected utility. This is the aim of the present paper.

The crucial idea is the 'incompleteness hypothesis' which states that, because formal choice models necessarily omit some aspects of decision problems from consideration, their use in poorly-understood problems will introduce a bias in favor of overly optimistic decisions. The incompleteness hypothesis implies support for the precautionary principle, considered as a procedural constraint on decision-making, rather than as a decision rule.

The paper is organised as follows. Section 1 contains some background information on the precautionary principle, though this does not amount to a survey of the voluminous literature on the topic. Section 2 presents the incompleteness hypothesis, which asserts that, because any particular model of decision under uncertainty inevitably omits some relevant factors, estimates derived from such a model will generally be over-optimistic. The errors will be greater, the less well-understood is the problem in question. This point is illustrated in relation to a sequence of models of decision under uncertainty, each more general than its predecessor. The core of the paper is Section 3, where a general form of the incompleteness hypothesis is shown to imply the desirability of using a version of the precautionary principle in decision-making. In Section 4, the implications of the analysis are discussed, with specific reference to climate change and policies to mitigate it. Finally, some concluding comments are presented.

1. Background

The literature on the precautionary principle is too large to permit the presentation of an adequate survey. VanderZwaag (1999) identifies fourteen different definitions of the principle. Despite the inevitable vagueness that results from discussing imprecisely defined concepts, some observations can be made.

The simplest applications of the precautionary principle have arisen in cases where a new activity is proposed, and concerns are raised that it may involve risks to the environment or to human health. Two of the most commonly-cited examples are the production and marketing of genetically modified foods and the exposure of the public to various kinds of electromagnetic fields, such as radio frequency fields.

In both examples, the low-risk course of action indicated by the precautionary principle is easily identified, namely to restrict exposure to the potential hazard in question. In other cases, the balance of risk is less clear. For example, application of the precautionary principle to constrain the growth of nuclear power might lead to an expansion of coal-fired electricity generation. Both options have potential adverse consequences that are poorly understood.

Debates over the precautionary principle have typically involved an interaction between scientific and legal standards of proof. Various standards of proof are considered in legal discussion, including proof beyond reasonable doubt, the balance of probabilities, and reasonable grounds for belief. Of these, only the balance of probabilities has a generally-accepted meaning in terms of formal decision theory.

In legal discussion, the standard of proof is often less important than the burden of proof. Advocates of a precautionary principle argue that those proposing an innovation should have the burden of proving it to be safe. By contrast, unless regulation is based, explicitly or implicitly, on a precautionary principle, the presumption in a liberal social order is that individuals (including corporations and other bodies endowed with some form of legal personality) should be free to pursue whatever activities they wish in the absence of evidence sufficient to show that such activities represent a danger to others. The burden of proof therefore falls on those seeking to show that constraints should be imposed.

Scientific notions of proof are also crucial. Definitions of the precautionary principle often refer to the need for decisions to be taken in the absence of conclusive prove. For example, Appell (2001) cites the following definition from the Wingspread conference, held in Racine, Wisconsin in 1998:

> When an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause and effect relationships are not fully established scientifically.

The introduction of notions of scientific proof raises further difficulties. As Kinzig et al. (2002) note, scientists are normally cautious about claiming that particular propositions have been 'proved' or 'established'. The same caution is exhibited by the principles of classical statistical inference, in which the null hypothesis of no correlation between variables is rejected only in the presence of evidence sufficient to yield a confidence level of 95 or 90 per cent.[1] Indeed, these

[1] The apparent conservatism of this approach is, to some extent, illusory. Formal or informal application of search procedures, popularly referred to as 'data mining', can induce rejection of the null hypothesis with a probability well above the stated significance level of 5 or 10 per cent, even when the data set is generated by random variables with no correlation.

procedures may be seen as a kind of precautionary principle, based on the presumption that it is worse to claim proof for a false statement (a Type 1 error in the terminology of classical inference) than to fail to claim proof for a true statement (a Type 2 error).

However, the cautious approach of classical inference is justifiable only when failure to reject the null hypothesis is, in some relevant sense, the 'safe' option. In a situation where it is necessary to choose some action, whether a given hypothesis is accepted or rejected, there is no general reason to suppose that it is safer to act on the basis of the assumption that the null hypothesis is true.

2. The incompleteness hypothesis and decision theory

The precautionary principle is best considered in relation to the standard prescription of normative theories of choice under uncertainty, namely, to choose the course of action that yields the highest expected (net) benefits. In this context, it is useful to begin by considering a claim that will be referred to as the incompleteness hypothesis.

In relation to any particular model of choice under uncertainty the incompleteness hypothesis asserts that, because the model fails to capture all relevant aspects of the problem, it will yield inaccurate estimates of the expected benefits of any given course of action. Further, the incompleteness hypothesis states that estimates will generally be over-optimistic and that the errors will be greater, the less well-understood is the problem in question. Implicitly or explicitly, the incompleteness hypothesis relies on the existence of an encompassing model, more complete than that under consideration, that would yield more accurate estimates than the model in use.

As will be discussed in more detail below, acceptance of the incompleteness hypothesis with respect to any given model of choice under uncertainty implies some form of precautionary principle in relation to decisions made using that model. It follows from the incompleteness hypothesis that the standard prescription of choosing the action that yields the highest expected benefits will lead to a bias in favor of choosing courses of action that are poorly understood. The need to correct this bias leads to the precautionary principle.

A number of examples of the incompleteness hypothesis will be considered. The examples take the form of a sequence of models, each more general than its predecessor, and each introducing new aspects of uncertainty. At each stage, the precautionary hypothesis holds true in general. From the viewpoint of the encompassing model, the predecessor model is a special case incorporating a bias towards over-optimism.

The best-projection approach and the expected value approach

The simplest, and still one of the most commonly used, method of evaluating a proposed course of action involving uncertain outcomes is to choose some particular projection of future uncertain events, including estimates of the values of unknown parameters, and to select the course of action that would yield the best outcome under that projection. This will be referred to as the best-projection approach. A closely related approach is that of the surprise-free projection (Kahn 1965).

The best-projection approach may be compared with more general approaches to benefit–cost analysis under uncertainty, in which a number of different possibilities are taken into account. The simplest such approach is sensitivity analysis, in which the consequences of varying individual parameters are assessed. A more general and systematic approach is expected-value analysis, in which a joint probability distribution over relevant parameters is used to calculate an expected value, expressed in monetary terms. Considered in the light of the expected value approach, two problems with the best-projection approach are relevant in consideration of the precautionary hypothesis.

First, it is common to use modal estimates for parameter values, and to combine them in a linear fashion. If probability distributions are skewed, or if variables are related in a nonlinear fashion, this will produce biased estimates. More commonly than not, skewness is usually associated with a long tail of unfavorable events and nonlinearity with undesirable interactions between variables.

One way of addressing the problem is to distinguish between 'pure' risk (variation about a central value) and 'downside' risk (the risk that a variable of interest may fall below the desired value). The problem of downside risk is discussed by Quiggin (2004).

The second problem is that of option value (Arrow and Fisher 1974; Henry 1974). In a situation of uncertainty, the best response will typically depend on the value of uncertain parameters. Hence, other things being equal, it is better to wait until uncertainty is resolved before making a decision, rather than implementing the decision that would yield the best expected outcome on the basis of available incomplete information. Thus, an action that does not foreclose future options should be preferred to one that is irreversible. The benefit from waiting is analogous to that of holding a financial option, and is referred to as 'option value'.

This point is developed in more detail by Gollier, Jullien and Treich (2000) and Heal and Kristrom (2002). In particular, consider a situation where a severe outcome might arise from a low-probability event or from the interaction of a number of adverse events. Use of the best-projection approach would normally lead to the exclusion of such events from consideration, and therefore to the adoption of an overly optimistic decision. By contrast, an expected value approach would take such events into consideration.

The problems with the best-projection approach have led to significantly overoptimistic estimates in many cases. The troubled history of *ex ante* project evaluation provides ample confirmation of the proposition that a best-projection approach will typically lead to overestimates of the returns from an investment project. A good recent summary of the evidence is given by Flyvberg, Bruzeliu, and Rothengatter (2003).

The expected value approach and the expected utility approach

As discussed in the previous section, the expected value approach involves consideration of the distribution of possible outcomes from any given course of action, and the selection of the action that yields the best mean outcome. Outcome values are typically expressed in monetary terms.

Although the expected value approach is widely used in the analysis of choices involving uncertainty, the framework favored by most economists for the analysis of such choices is that of expected utility theory. The expected value model is the special case of the expected utility model when utility is a linear function of wealth.

In the application of expected utility theory, it is normally assumed that the utility function is concave. This assumption is normally characterized as risk aversion, since it has the behavioral implication that any monetary outcome, received with certainty, is preferred to a risky prospect with the same expected value. As will be shown below, in general models of choice under uncertainty, risk averse behavior may arise from many different sources. In the expected utility model, risk aversion arises from the diminishing marginal utility of money (income, consumption or wealth).[2]

Considered in the light of expected utility theory, the expected value

[2] This way of expressing things relates to cardinal utility: a restatement in purely ordinalist terms might refer to the fact that the rate of substitution between income in states of nature A and B diminishes as income in state A increases.

approach is biased in favor of risky or uncertain options. As has been noted, risk aversion implies that any monetary outcome, received with certainty, is preferred to a risky prospect with the same expected value. More generally, of two distributions yielding the same expected value, the less risky[3] is to be preferred. By continuity, a certain outcome will be preferred to a risky prospect with a slightly higher expected value.

Hence, if a number of projects are being compared, use of the expected value approach will lead to a bias in favor of a more risky approach. This is true even in the presence of *ad hoc* corrections such as the use of a discount rate that is higher than the real bond rate. Such corrections penalize long-term projects but do not correct appropriately for a bias in favor of risky projects, unless relative risk grows linearly over time (Little and Mirrlees 1974).

This problem has become evident in the literature on sustainable growth. Consider a problem where there is some probability that a given growth path will prove unsustainable, yielding substantial reductions in income in the distant future. The use of an expected value approach will place inadequate weight on this outcome. Raising the discount rate to 'adjust' for risk will only exacerbate the problem. Under these circumstances, the discount rate is dominated by adverse outcomes, in the sense that much of the value of future consumption will arise where that consumption is available in states of nature where consumption is lower than at present. It follows that the distant future should be discounted at significantly lower rates than suggested by the current market discount rate (Newell and Pizer 2003).

---

[3] In the context of expected utility theory, the most natural concept of 'less risky' is that derived from the work of Rothschild and Stiglitz (1970). See also Quiggin (1991).

The expected utility approach and the rank-dependent utility approach

After subsuming and displacing mean-variance analysis in the 1960s, expected utility theory was the only framework used in economic analysis, to any significant extent, for several decades. However, the accuracy of expected utility theory, considered as a descriptive model of individual preferences came under increasing criticism during the 1970s, culminating in the critique of Kahneman and Tversky (1979). This criticism also renewed interest in the much earlier critique put forward by Allais (1953).

Allais had argued that the utility function was best understood in cardinal terms and that individuals might not choose to maximize expected utility, but might be concerned also with the variance of utility about its expected value. Defenders of the expected utility approach claimed that Allais had misunderstood the argument, and that risk aversion was entirely captured by the curvature of the utility function.

Disputes of this general form have been a recurring feature of the development of the theory of choice under uncertainty. Advocates of more general models have argued that existing models fail to capture important aspects of attitudes to risk and, in particular, of risk aversion. Supporters of the existing model argue that, to the extent they capture real phenomena, the supposedly new aspects of risk aversion are encompassed by the existing concepts.

The critiques of the 1970s led to the development of a wide range of alternatives to, and generalizations of, expected utility theory. Of these, the most significant and widely-used have been prospect theory (Kahneman and Tversky 1979) and rank-dependent utility models (Quiggin 1981, 1982; Segal 1987; Schmeidler 1989). The two approaches have been combined to yield cumulative prospect theory (Tversky and Kahneman 1990).

The central insight underlying rank-dependent approaches, going back to Allais (1953), is that risk attitudes may depend on the probability with which particular utility levels are realized and, more generally, on the entire probability distribution over utility levels, and not merely on the expectation of utility. Quiggin (1981, 1982) showed how this idea could be formalized using a rank-dependent transformation of the probability distribution. Schmeidler (1989) developed the same idea in the state–act framework due to Savage (1954), where probabilities are derived from preferences rather than being objectively given.

Quiggin (1981, 1982) argued for an 'S-shaped' probability distribution, which increased the weight on both the worst and the best outcomes in a symmetrical fashion. Other writers have examined the case of a concave transformation, embodying a concept of 'pessimism', in which lower-ranked outcomes are always given higher weight relative to their objective probabilities. Most empirical studies support a transformation that is S-shaped, but which is pessimistic over most of its range.

Bargiacci (2003) examines the relationship between rank-dependent utility and the precautionary principle, with specific application to climate change. With pessimistic preferences, the evaluation of risky outcomes is less favorable than under expected utility with the same utility function. With an S-shaped transformation function, impacts are ambiguous, though the general tendency is for less favorable evaluation of risky outcomes.

For random variables with moderate variance and symmetric distributions (such as normal distributions) the extension from expected utility to rank-dependent expected utility is unlikely to make much difference, even in the presence of systematic pessimism. This is because, if the risk preferences of a rank-dependent expected-utility maximizer are elicited under the incorrect assumption of expected-utility preferences, pessimism in the probability transformation will be reflected in additional concavity in the elicited utility function.

The importance of rank-dependent preferences is likely to be greatest when some choices yield distributions of outcomes that are skewed to the right (have a long left tail), so that there is a small probability of a severe adverse outcome. Given the existence of adverse low-probability events, rank-dependent expected-utility preferences displaying overweighting of extreme probabilities will yield more negative evaluations of the 'business as usual' strategy than will expected-utility preferences with the same utility function, even though on average, the two sets of preferences are about equally risk-averse.

Risk and ambiguity

In all of the models considered thus far, it has been assumed that, for any action under consideration, there exists a well-defined probability distribution over consequences. The usual way of formulating this assumption is the Savage (1954) framework in which acts are considered as a mapping from a space of states of nature to a space of outcomes. If preferences over actions display appropriate consistency properties, they are described as probabilistically sophisticated (Machina and Schmeidler 1992). That is, there exists a probability distribution over the state space with respect to which preferences respect first-order stochastic dominance. Both expected-utility and rank-dependent expected-utility preferences, as well as a large class of generalized expected-utility models, display probabilistic sophistication.

The assumption of probabilistic sophistication works well in decision problems where probabilities can be inferred from objective information, such as previous observations of the frequency of particular outcomes or knowledge about the physical characteristics of, say, a die. It may also be extended to problems involving subjective probabilities where decision-makers have extensive experience of judging problems of a given kind, such as the outcomes of horse races or marketing campaigns.

However, when faced with complex or unfamiliar problems, decision-makers often find probabilistic reasoning unhelpful. Keynes (1920) first formulated the objection that for many important decisions there was no basis for determining a reasonable probability distribution for outcomes. Experiments conducted by Ellsberg (1961) showed that in such situations, people preferred bets with known odds to either side of a bet in which there was no easy way to formulate probabilities. The latter situation is described as one of ambiguity.

The most successful approach to the analysis of problems of this kind has been the multiple-priors model of Gilboa and Schmeidler (1989). In this model, decision-makers are assumed to consider a set of possible probability distributions over states of nature. If decision-makers are averse to ambiguity, they will evaluate actions according to the probability distribution that is least favorable.

Compared to any assumption of fixed probabilities, the multiple-priors approach with ambiguity aversion will yield a lower evaluation of ambiguous actions whenever the given probabilities lie within the set of priors. The two approaches will coincide if probabilities are unambiguous, that is, if the set of priors has a single element. Thus, relative to the fixed probability approach, a multiple-priors approach will be less favorable to decisions involving high levels of uncertainty.

State spaces and proposition spaces

Although the multiple-priors model relaxes the unrealistic assumption that decision-makers have well-defined subjective probability distributions for all possible events, it still requires them to hold unrealistically precise beliefs about uncertain events. In particular, decision-makers are expected to be able to describe uncertainty in terms of a mutually exclusive and exhaustive set of possible states of the world. Although this is obviously implausible, simple modifications to the Savage framework (for example, the inclusion of a residual 'unspecified' event) do not appear to yield useful insights into the problem.

Grant and Quiggin (2004) adopt an alternative approach. Beliefs are described in terms of a finite set of propositions, which decision-makers can consider as true, false or possible. This finite set is assumed to be a proper subset of the set of propositions (assumed countably infinite) needed to characterize all possible states of the world.

The crucial contribution of this approach is that it provides a way of describing how new propositions, previously not considered, may enter the thinking of a decision model. Currently unconsidered elements of the set of all propositions are described as accessible if they have high information value with respect to the set of propositions currently under consideration by the decision-maker.

The prescriptive implications of the model are less clear-cut. Grant and Quiggin suggest that, in the absence of a complete description of the space of states of nature, a case-based approach similar to that advocated by Gilboa and Schmeidler (1995) may be appropriate. The case-based approach provides a rationale for the use of rules of thumb, like the precautionary principle, where these are supported by past experience.

Risk and rationality

Discussion of alternatives to, and generalizations of, expected utility theory have raised both positive and normative issues. It is fairly widely agreed that expected utility theory does not provide an adequate positive description of observed choices under uncertainty. On the other hand, many defenders of the expected utility model argue that its normative appeal as a guide to rational choices remains undiminished.

Supposing that policymakers accept the normative appeal of expected utility theory, but recognize that people do not, in general, act in accordance with its prescriptions, what should they do? It is commonly suggested that, in such circumstances, policymakers should disregard 'irrational' preferences, but this suggestion creates two major difficulties.

First, there are general issues of democratic process. If policymakers think that it is appropriate, as a general principle, to follow a rule of maximizing expected utility, they should seek to persuade the public of the desirability of this principle, rather than imposing it by fiat.

Second, in the absence of comprehensive central planning, a situation where policymakers implement preferences different from the aggregate preferences of the public is likely to generate welfare losses. Suppose, for example, that consumers place a high weight on the risk of harm from contaminants in water supplies, and that policymakers judge that a lower weight is appropriate. If policymakers follow this judgement and refraining from implementing improvements in water safety for which, on their assessment, costs exceed benefits, consumers may respond by switching to bottled water, incurring higher costs than would have been required to implement the improvements.

For both these reasons, it seems appropriate that, at least where individual preferences are consistent and permit aggregation, it seems better to follow these preferences in public decision-making, rather than substituting the judgements of policymakers. Hence, it seems likely that the expected utility model will not, in general, be the most appropriate approach.

An illustration: the Rasmussen Report

A noteworthy illustration of the problems in attempting a complete probabilistic specification of the state space is the Rasmussen report (US Nuclear Regulatory Commission 1974) on nuclear safety. Rasmussen and his colleagues attempted to estimate the probability of nuclear accidents using an event-tree analysis, and concluded that the probability of a serious meltdown was minuscule (one in 20,000 per reactor per year for a core meltdown)

The partial meltdown at Three Mile Island, which occurred in 1979, illustrated both the strengths and limitations of the event-tree approach. On the one hand, the general form of the accident was one that had been considered by Rasmussen. On the other hand, the chain of problems, including operator errors that aggravated the severity of the accident, was not. This implies that Rasmussen almost certainly underestimated the likelihood of more severe accidents. A study by Nordhaus (1979), using aggregate empirical evidence, rather than event-tree modeling concluded that

> Using the technique of maximum likelihood, our best guess estimate of the risk of accidents causing at least one fatality rises from the Reactor Safety Study's 32 per million reactor years to about 2000 per million reactor years.

3. The incompleteness meta-hypothesis and the precautionary principle

As noted above, in relation to any particular model of choice under uncertainty the incompleteness hypothesis asserts that, because the model fails to capture all relevant aspects of the problem, it will yield inaccurate estimates of the expected benefits of any given course of action. The discussion of the previous section suggests that the incompleteness hypothesis remains valid as we consider a sequence of increasingly general models, from the simplistic best-projection approach to models incorporating finite knowledge and multiple priors. Thus, we may consider encompassing these specific versions of the incompleteness hypothesis with a meta-hypothesis. The incompleteness meta-hypothesis states that:

Estimates of project outcomes derived from formal models of choice under uncertainty are inherently incomplete. Incomplete estimates will generally be over-optimistic. The errors will be greater, the less well-understood is the problem in question.

The second part of the hypothesis is crucial. If a problem is well understood within a given formal model, incompleteness will not be a serious issue. For example, casinos can normally rely on expected-value calculations, since the Law of Large Numbers ensures that, if a game has a positive expected value for the house, it will yield a positive average return over many plays, with probability close to one. Similarly, expected utility models appear to work well in the absence of 'edge effects' (low-probability events with extreme outcomes).

The incompleteness principle is a statement about the evaluation of prospects. By contrast, as noted above, the precautionary principle is typically presented as a guide to action in relation to proposed innovations. To link the two, it is necessary to add the auxiliary hypothesis that the consequences of innovations are less well-understood than the consequences of maintaining the *status quo*. In a situation where the *status quo* has been sustained for a long period, this hypothesis is not problematic. In many cases, however, there is no pre-existing equilibrium, but rather a set of 'business as usual' practices that may or may not be sustainable and for which the consequences of persisting with existing practice may or may not be well understood.

Assuming for the moment that the auxiliary hypothesis of an initial stable equilibrium is valid, the incompleteness hypothesis suggests that a formal evaluation within an incomplete model is likely to be biased in favor of innovation and against the *status quo*. This in turn implies that some sort of burden of proof should be placed on the advocates of innovation.

On the other hand, the incompleteness principle does not support strong versions of the precautionary principle in which the burden of proof is taken to mean 'proof beyond reasonable doubt' or something similar. Such an approach could be supported only by highly pessimistic decision criteria such as maximin, that is, maximizing the value of the worst possible outcome. Although maximin decision criteria have been advocated in various contexts, there is little empirical or theoretical support for the use of such criteria. The only context in which maximin is clearly justified is the case of zero-sum games, analysed by von Neumann and Morgenstern (1944). In relation to environmental problems, maximin would make sense only if 'Nature' were viewed as a malevolent opponent.

The auxiliary hypothesis that the consequences of innovations are less well-understood than the consequences of maintaining the *status quo* seems to work fairly well in the case of techniques associated with genetic engineering. The *status quo* in this case includes both long-standing traditional methods of crop and animal breeding, which have produced plants and animals with genetic endowments radically different from those of their wild ancestors, and more recent technical innovations such as artificial insemination and embryo transfer.

Advocates of the precautionary principle has been criticized for failing to provide a precise operational definition of the principle. The discussion presented above suggests that this criticism is misplaced. Any precise definition implies the existence of a well-defined formal analytical model within which the principle may be applied. But the incompleteness hypothesis states that any such model will exclude relevant factors. Hence, the precautionary principle must necessarily be considered as a heuristic check on formal decision-making procedures rather than as a rule to be applied within a given formal framework.

The standard practices of engineering provide a useful analogy. Particularly in critical applications, it is not good engineering practice to compute the optimal trade-off between cost and the risk of failure, using a standard model, and then to adopt the indicated solution. Rather the standard approach is to compute the optimal solution, then to allow a substantial safety margin, based on a combination of past experience and rules of thumb. This may be seen as an instance of the precautionary principle at work.

4. The precautionary principle and global warming

The problem of climate change may be used to illustrate the issues raised by the interpretation of the precautionary principle offered here.

Background

The global climate is determined, in large measure, by the 'greenhouse effect' of the earth's atmosphere, which reduces the extent to which heat is radiated into space, and thereby raises the global temperature. Without this greenhouse effect, life on earth would not be sustainable. On the other hand, rapid change in the magnitude of the effect will induce changes in global climate with generally adverse effects.

The magnitude of the greenhouse effect is primarily determined by the relative concentrations of different gases (and water vapor) in the atmosphere. Human activity has greatly increased the concentration of some gases, such as carbon dioxide (CO2), that promote the greenhouse effect.

Although there is general agreement on these basic points, almost every other aspect of the climate change problem is the subject of both disagreement and uncertainty. Most measures suggest that average global temperatures have increased over the past fifty years, and most climate models suggest that this increase is due, at least in part, to human activity. The evidence on these points has been summarized by the International Panel on Climate Change (IPCC) (2001a,b), which has also prepared a range of projections of changes in global climate, employing alternative models and a range of scenarios generating different time paths for emissions of greenhouse gases.

The main focus is on simulations using the Atmosphere-Ocean General Circulation Model (AOGCM), and modeling the change in global average surface air temperature, as measured by the difference between the average for the period

1961–1990 and the average for the period 2021–2050. In these simulations, the mean temperature increase is 1.3°C (IPCC 2001a).

Such an increase in temperature would damage some vulnerable ecosystems such as coral reefs, and might increase risks of flooding and storm damage in coastal areas and low-lying countries such as Bangladesh. However, for most countries, the effect on human activity would be modest. The IPCC also surveyed a number of studies using different climate models, with differing assumptions and over different time periods. The increases in average global temperature estimated in these studies range from 0.8°C to more than 5°C.

Responding to projections of climate change, a Climate Convention held in Kyoto, Japan, in 1997 agreed to the Kyoto Protocol (United Nations 1997) to the United Nations Framework Convention on Climate Change, originally adopted in 1992. Under the Protocol, developed nations agreed that, by 2012, they would reduce emissions of greenhouse gases to 1990 levels, subject to a complex set of adjustments for individual circumstances.

The United States and Australian governments subsequently announced that they would not ratify the protocol. The only other large country not to ratify the Protocol is Russia, where legislation to ratify the treaty is currently (October 2004) under consideration by the Duma (Parliament). Ratification by Russia would bring the treaty into force.

Implementation of the Kyoto Protocol would have only a modest effect on the rate of global warming, relative to business as usual. Hence, advocates of the Protocol normally regard it as a first step, preparatory to a broader agreement that would include less-developed countries, and would entail deeper cuts in emissions for developed countries.

A range of computable general equilibrium models have been used to model the economic costs and benefits of implementing the Kyoto protocol, with broadly consistent results. As would be expected on the basis of standard partial equilibrium analysis, targets for reductions in emissions can be achieved at a lower net cost through trade in emissions rights than through the imposition of quantitative restrictions on particular sources of emissions or source countries.

Estimates of the net cost of implementing the Kyoto Protocol through an emissions trading system range from 0.1 per cent of world product to 1.5 per cent. There has been less detailed modeling of the economic effects of policies to achieve the ultimate objective of stabilizing global concentrations of greenhouse gases. Such policies would require substantial reductions in the use of fossil fuels, and might be expected to lead to a doubling of the unit cost of energy. This would imply a welfare loss comparable in magnitude to the share of energy in world product, which is around 5 per cent. The adoption of mitigation strategies would yield benefits such as reductions in losses of coastal land and in biodiversity. There is no generally accepted monetary estimate of the value of these benefits.

Incompleteness

The climate change problem illustrates several aspects of the incompleteness hypothesis and the precautionary principle. Projections of the likely rate of climate change, and of its likely effects are incomplete in several important respects.

Reliance on a best estimate, such as the IPCC (1999) mean projection of a global temperature increase of 1.3°C, as opposed to a range of possible projections, could be misleading in a number of respects. First, consideration of a single projection may lead to the adoption of excessively inflexible policies for mitigation of climate change, without the capacity for adaptation to new information.

Second, many of the consequences of climate change are related nonlinearly to the rate of climate change. Although the consequences of an increase in global mean temperatures of 1.3°C over 50 years would be relatively modest, the consequences of an increase of 5°C over the same period could have catastrophic consequences, particularly if it led to large increases in sea levels.

Third, even if such catastrophic outcomes have low probability, many decision-makers might consider it appropriate to place a high weight on preventing them. Thus, an analysis based on expected costs and benefits would prove inappropriate.

Finally, it is important to consider the possibility of unforeseen developments that might radically alter the projections. By definition, such developments cannot be described in detail, but relevant possibilities include technological innovations (which might permit low-cost mitigation in future) or previously unknown climatic feedbacks (which might either mitigate or exacerbate climate change).

In summary, any formal approach to projecting climate change and its consequences is likely to be subject to the problem of incompleteness. There are also problems of incompleteness with respect to estimates of the costs of programs to mitigate climate change. However, as shown above, these problems are less severe than those of estimating the effects of climate change.

The precautionary principle

Before the precautionary principle can be applied, it is necessary to consider the nature of the 'innovation' under question. The answer to this question largely determines the way in which the policy debate is framed.

From the perspective of fossil fuel users, the introduction of restrictions on emissions of greenhouse gases is an innovation. On the other hand, doubling the concentration of carbon dioxide and other greenhouse gases is clearly an innovation as far as the global climate is concerned.

In this case, the consequences of continuing 'business as usual' are less well understood than the consequences of substantial reductions in emissions. Hence, the precautionary principle favors the adoption of measures to mitigate the danger of global warming, even if it is not clearly established that the benefits of those measures will exceed the costs.

In this context, the main benefit of measures undertaken to implement the Kyoto Protocol is that they provide the basis for a more extensive mitigation policy if the information that becomes available over the next decade confirms a relatively pessimistic assessment of the outlook for climate change. So, it is important that the approach taken to implementation should be consistent with the adoption of a broader agreement including developing as well as developed countries. In particular, it is desirable that mechanisms for global trade in emissions rights be developed as part of the implementation process.

It is, of course, possible that new information will indicate that concerns about climate change have been overstated. If so, resources devoted to implementation of the Kyoto Protocol will turn out, *ex post*, to have been wasted. But the discussion above indicates that, in problems of this kind, unpleasant surprises are more common than pleasant ones. It will be preferable to have devoted excessive resources to preparing for an outcome that turns out better-than-expected than to have devoted inadequate resources to preparation for a worse-than-expected outcome.

Concluding comments

Although the precautionary principle has played a prominent role in public debate for more than a decade, attempts to state the principle as an operational decision rule have produced no broad agreement either on the correct statement of the principle or on its validity as a guide to decisions. Disagreement of this kind suggests that the discussion is taking place at the wrong level of analysis. It may not be possible to state the precautionary principle as a formal decision rule. That does not diminish its importance as a guide to good decisions.

In this paper, it has been argued that the precautionary principle is best understood as a procedural 'burden of proof' constraint, requiring that arguments for risky innovations be held to a more stringent standard than that they are shown to be optimal by a (necessarily incomplete) decision-theoretic analysis. Incompleteness, and the associated bias towards poorly-understood options affects all formal decision procedures, from the commonplace best-projection approach to more sophisticated expected-utility analysis.

Viewed in this light, the precautionary principle provides a useful framework for the assessment of a range of policy problems, particularly environmental issues, where some components of the problem are well-known and amenable to formal analysis, while knowledge about other components of the problem is uncertain, ambiguous and incomplete. The problem of climate change is an ideal example.

References

Allais, M. (1953), 'Le comportement de l'homme rationnel devant le risque: critique des axioms et postulats de l'école Americaine', *Econometrica*, 21(4), 503–46.

Appell, D. (2001), 'The new uncertainty principle', *Scientific American*, January, http://www.biotech-info.net/uncertainty.html.

Arrow, K. and Fisher, A. (1974), 'Environmental preservation, uncertainty and irreversibility', *Quarterly Journal of Economics*, 2, 312—19.

Bargiacchi, R. (2003), 'Climate change scenarios and the precautionary principle', pp. 113–130 in (J.Wesseler, Weikard, H.P. & R.Weaver, eds) *Risk and Uncertainty in Environmental and Resource Economics*,, Edward Elgar, London.

Ellsberg, D. (1961), 'Risk, ambiguity and the Savage axioms', *Quarterly Journal of Economics*, 75(4), 643–69.

Flyvbjerg, B., Bruzeliu, N. & Rothengatter, W. (2003) *Megaprojects and Risk: an Anatomy of Ambition*, Cambridge University Press, Cambridge.

Ghirardato, P. (2001), 'Coping with ignorance: unforeseen contingencies and non-additive uncertainty', *Economic Theory*, *17*, 247–76.

Gilboa, I. and Schmeidler, D. (1989), 'Maxmin expected utility with non-unique prior', *Journal of Mathematical Economics*, 18(2), 141–53.

Gilboa, I. and Schmeidler, D. (1995), 'Case-based decision theory', *Quarterly Journal of Economics*, 110, 605–39.

Gollier, C., Jullien, B. and Treich, N. (2000), 'Scientific progress and irreversibility: an economic interpretation of the "Precautionary Principle"', *Journal of Public Economics*, 75, 229–53.

Gollier, C. and Treich, N. (2003), 'Decision-Making under Scientific Uncertainty: The Economics of the Precautionary Principle", *Journal of Risk and Uncertainty*, 27(1), 77-103

Grant, S. and Quiggin, J. (2004), 'Unforeseen contingencies: a propositional approach', Paper presented at FUR XI Conference, Paris, 3 July.

Heal, G. and Kriström, B. (2002), 'Uncertainty and climate change', *Environmental and Resource Economics*, 22, 3–39.

Henry, C. (1974), 'Investment decisions under uncertainty: the "Irreversibility Effect"', *American Economic Review*, 64(6), 1006–12.

Intergovernmental Panel on Climate Change (2001a), 'IPCC Third Assessment Report: Climate Change 2001',http://www.grida.no/climate/ipcc_tar/

IPCC(2001b), Climate Change 2001: Working Group I: The Scientific Basis, http://www.grida.no/climate/ipcc_tar/wg1/339.htm.

Kahn, H. (1965) *On Escalation: Metaphors and Scenarios*, Praeger, New York.

Kahneman, D. and Tversky, A. (1979), 'Prospect theory: an analysis of decision under risk', *Econometrica*, 47(2), 263–91.

Keynes, J.M. (1920) *A Treatise on Probability*, MacMillan, London.

Kinzig, A., Starrett, D. and others (2002), 'Coping With Uncertainty: A Call for a New Science-Policy Forum',Beijer Institute workshop at Askö marine laboratory Aug. 31 - Sep. 2

United Nations (1997), 'Kyoto Protocol to the United Nations Framework Convention on Climate Change', United Nations, New York, http://unfccc.int/resource/docs/convkp/kpeng.pdf .

Little, I. and Mirrlees, J. (1974) *Project Appraisal and Planning for Developing Countries,* Heinemann, London.

Machina, M. and Schmeidler, D. (1992), 'A more robust definition of subjective probability', *Econometrica*, 60(4), 745–80.

Newell, R. and Pizer, W. (2003), 'Discounting the distant future: how much do uncertain rates increase valuations?', *Journal of Environmental Economics and Management*, 46, 52–71.

Nordhaus, W. (1979), 'How should we revise our beliefs about nuclear power safety after Three Mile Island?"', Cowles Foundation Discussion Paper No. 532, Yale University.'

Quiggin, J. (1981), 'Risk perception and risk aversion among Australian farmers', *Australian Journal of Agricultural Economics*, 25(2), 160–69.

Quiggin, J. (1982), 'A theory of anticipated utility', *Journal of Economic Behavior and Organization*, 3(4), 323–43.

Quiggin, J. (1991), 'Increasing risk: another definition', In *Progress in Decision, Utility and Risk Theory,* (Ed, Chikan, A.) Kluwer, Amsterdam,

Quiggin, J. (2004), Risk and discounting in project evaluation, Report to Bureau of Transport Economics, Canberra.

Rothschild, M. and Stiglitz, J. (1970), 'Increasing risk: I. A definition', *Journal of Economic Theory*, 2(4), 225–43.

Savage, L.J. (1954) *Foundations of Statistics,* Wiley, New York.

Schmeidler, D. (1989), 'Subjective probability and expected utility without additivity', *Econometrica*, 57, 571–87.

Segal, U. (1987), 'Some remarks on Quiggin's Anticipated Utility', *Journal of Economic Behavior and Organisation*, 8(1), 145–54.

Tversky, A. and Kahneman, D. (1992), 'Cumulative prospect theory: an analysis of attitudes towards uncertainty and value', *Journal of Risk and Uncertainty*, 5(3), 297–323.United Nations (1997), 'Kyoto Protocol to the United Nations

Framework Convention on Climate Change', United Nations, New York, http://unfccc.int/resource/docs/convkp/kpeng.pdf .

US Nuclear Regulatory Commission (1974), 'Reactor Safety Study (WASH-1400) (Rasmussen Report)', Washington DC.

VanderZwaag, D. (1999), 'The precautionary principle in environmental law and policy: elusive rhetoric and first embraces', *Journal of Environmental Law and Practice*, 8, 355–75.

von Neumann, J. & Morgenstern, O. (1944) *Theory of Games and Economic Behavior*, Princeton University Press, Princeton.