

Provided for non-commercial research and education use.  
Not for reproduction, distribution or commercial use.



**This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.**

**Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.**

**In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:**

**<http://www.elsevier.com/copyright>**

Contents lists available at [SciVerse ScienceDirect](http://www.sciencedirect.com)

## Global Environmental Change

journal homepage: [www.elsevier.com/locate/gloenvcha](http://www.elsevier.com/locate/gloenvcha)

# Precaution and proportionality in the management of global environmental change<sup>☆</sup>

Terrence Iverson<sup>a,\*</sup>, Charles Perrings<sup>b</sup>

<sup>a</sup> Department of Economics, Colorado State University, United States

<sup>b</sup> School of Life Sciences, Arizona State University, United States

## ARTICLE INFO

## Article history:

Received 27 March 2010

Received in revised form 31 August 2011

Accepted 13 September 2011

Available online 12 October 2011

## Keywords:

Precautionary principle

Decisions under uncertainty

Global environmental change

Environmental externality

Multilateral environmental agreements

Climate change

## ABSTRACT

The precautionary principle is a mandate to tread cautiously when managing novel threats to the environment or human health. A major obstacle when applying the principle at the international level is disagreement about how precautionary efforts should be constrained to ensure that policy costs are proportional to the attained level of protection. Proportionality is an unresolved question when preliminary evidence precludes decision-makers from assigning probabilities over future events. The paper suggests practical analytical tools for communicating *ex ante* trade-offs when probabilities are unavailable. The tools could be used to facilitate discussion and compromise when implementing precautionary decisions in international settings where cooperation is important. The approach is demonstrated in an application to climate policy that uses the integrated assessment model DICE (Nordhaus, 2008). The paper also situates the task of precautionary decision-making within the broader context of implementing a precautionary response at the international level.

© 2011 Elsevier Ltd. All rights reserved.

## 1. Introduction

The growth and closer integration of the world economy have led to the emergence of novel environmental threats with potentially severe consequences for public health and for the provision of many ecosystem services. Familiar examples are climate change and the risk of emerging infectious diseases (IPCC, 2007; Perrings et al., 2010; Jones et al., 2008). In each case, there are few historical precedents for the threat, so future impacts are highly uncertain. The effects may also be cumulative and delayed, so an effective response cannot wait for scientific uncertainty to be resolved—policy must be anticipatory. Governments have often failed to act on the basis of early scientific warnings. For example, the “evil effects” of asbestos dust were first noted by a British medical inspector in 1898, yet asbestos was not banned in the UK until 1998 (Harremoës et al., 2001; UNESCO, 2005). Incidents like this have prompted public concern that conventional approaches to risk management may not adequately protect the public interest when there is potential for severe and irreversible harm. The precautionary principle has emerged as a focal point of this discussion.

<sup>☆</sup> The authors thank Buz Brock, Steven Durlauf, and Stephen Carpenter for helpful suggestions. They also thank the United Nations Environment Programme for funding; the funding source did not participate in the design or implementation of the research.

\* Corresponding author at: Department of Economics, Colorado State University, 1771 Campus Delivery, Fort Collins, CO 80523-1771, United States.

Tel.: +1 970 491 6079; fax: +1 970 491 2925.

E-mail address: [terry.iverson@colostate.edu](mailto:terry.iverson@colostate.edu) (T. Iverson).

The precautionary principle offers a guide for regulating novel threats under conditions of severe scientific uncertainty. It is argued to have originated in Germany, in the early 1970s, as the concept of ‘Vorsorgeprinzip’ (Haigh, 1993)—“foresight planning”—which emphasizes avoidance of potentially damaging actions even where there is uncertainty about the consequences of those actions. The principle has since been applied to a wide range of environmental threats. There is, however, little consistency in the way it has been defined. Alternative versions range from the Rio Declaration (UNCED, 1993), a relatively weak assertion that uncertainty should not be used as an excuse to justify inaction, to stronger assertions, such as the Wingspread Statement, which states: “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.”<sup>1</sup>

While efforts to establish a universal definition for the precautionary principle have been controversial, mandates to bias decisions toward caution when addressing novel threats have not. For many categories of novel risk, in many countries, a precautionary approach is well established in national law. The Hippocratic mandate “do no harm” is reflected in procedures for approving drugs and for responding to the side effects of recently released drugs in most developed countries. The precautionary principle has been adopted in regulatory statutes for health and the environment in Australia, Canada, New Zealand and the UK, and it is a founding component of European environmental law due to its inclusion in

<sup>1</sup> <http://www.commonweal.org/programs/wingspread-statement.html>.

the 1992 Maastricht Treaty. In the United States, a precautionary approach is embraced in numerous environmental and consumer statutes—the Delaney clauses to the Food, Drug, and Cosmetic Act prohibit use of food or chemical additives found to be carcinogenic in animals, the Endangered Species Act allows regulators to stop development activity when there is risk of irreversible species extinction, and The Clean Air Act Amendments of 1977 instruct regulators to ‘assess risk rather than wait for proof of actual harm’ when setting emissions standards (Vogel, 2002; Ashford, 2007).

At the international level, a precautionary approach was first endorsed in the World Charter for Nature, adopted by the UN general assembly in 1982, and the precautionary principle was first explicitly invoked in the Second International Conference on the Protection of the North Sea (1987). The principle was widely adopted in the 1990 Bergen Declaration (signed by 84 countries as a follow-up to the Brundtland Commission report) and in the 1992 Rio Summit. Versions of the principle have since appeared in a variety of multilateral agreements, such as the Biodiversity Convention’s Cartagena Protocol on Biosafety and the Sanitary and Phytosanitary Agreement. Nevertheless, there is little consistency across agreements, and a precautionary approach is not yet a rule of customary international law (Cooney and Dickson, 2005; Sadeleer, 2007).

A major obstacle when implementing the precautionary principle at the international level stems from the need for multilateral cooperation. Many international threats entail international public goods: the risks are shared while the costs of precautionary actions are incurred at a national level. Without a sovereign global authority to mobilize resources, efforts to provide international public goods demand cooperation among nation states. The challenge of agreeing on a common policy target is heightened when scientific evidence is too preliminary for decision-makers to confidently assign probabilities—either because historical data is limited or because causal mechanisms are not well understood. In this case, formal decision tools for implementing the precautionary principle leave substantial room for stakeholders to disagree. To demonstrate, we briefly review the literature on implementing the precautionary principle as a formal criterion for decision-making under uncertainty. We then build on this discussion to motivate the approach described here.

The decision-theoretic literature on implementing the precautionary principle has progressed in two directions. One emphasizes the dynamic nature of scientific uncertainty. The consequences of decisions are irreversible if they involve a change in the state of a system and it is impossible to recover the original state. If there is scope for acquiring more information about the original state of the system in the future, then it is optimal for an expected-welfare-maximizing decision maker to bias decisions toward maintaining the option to learn (Perrings and Brock, 2009). This result is a straightforward application of backward induction—future information has value only if the option to use it has been preserved. It implies that the opportunity cost of an irreversible development project should be augmented to include the potential value of future information. If the project were to develop a Redwood grove, future information would include the recreational value of the grove to future generations (Arrow and Fisher, 1974). Henry (1974) called this the irreversibility effect. Epstein (1980) provided a more general characterization of the effect, and subsequent work has applied variations of Epstein’s result to the question of optimal resource management (for example, Ulph and Ulph, 1997).

Gollier et al. (2000) and Gollier and Treich (2003) associate the precautionary principle with the irreversibility effect. They identify an important class of circumstances in which a precautionary bias is justified and thus ground this version of precaution on compelling theoretical foundations. Nevertheless, the approach provides little guidance for situations in which decision-makers cannot agree on a common probabilistic description of the phenomenon in question—

including both current probabilities and the path by which future learning is expected to unfold.

A second direction identifies the precautionary principle with criteria for decision-making under Knightian uncertainty (also called ambiguity or deep uncertainty). In this paper, we refer to Knightian uncertainty as just ‘uncertainty’. Uncertainty, as opposed to risk, refers to a situation in which there is no basis for assigning probabilities across alternative models (Knight, 1921). For climate policy the distinction is important because many components of an integrated scientific description of anthropogenic climate change defy probabilistic descriptions. For this reason, IPCC assessments have always made careful distinction between stating ranges for values of interest and stating degrees of confidence (or consensus) among experts (Manning, 2006).

The criterion for decision-making under uncertainty that most clearly captures the notion of conservatism in the face of potentially severe harm is *maximin*. By design, maximin policies are robust to catastrophe—they implement the optimal response to the worst-case scientific model. Maximin attains a high level of protection with respect to environmental threats, but the protection comes at a cost: if the true model turns out to be less dire than the worst possible case, resources will have been wasted (Harsanyi, 1977). The magnitude of this concern naturally depends on the span of model uncertainty. If the span of potential models is small, the risk of wasting resources would also be small; in this case, maximin may be compelling. Vardas and Xepapadeas (2010), focusing on an application in which the set of scientific models (priors) can be constructed as a local neighborhood of a well-defined reference model, associate the precautionary principle with a version of robust control that embeds maximin in a closed-loop dynamic control setting. These methods were introduced into the economics literature by Hansen and Sargent (2001, 2008).

In contrast, if the span of model uncertainty is large—as it is for many global change problems, including climate change (Morgan and Keith, 1995; Stainforth et al., 2005)—the magnitude of resources that may turn out to be wasted under maximin could be very large. In practice, when maximin is applied, it is typically the initial criterion in a strategy that evolves over time, moving toward an approach that looks more like expected utility maximization as decision-makers gain confidence in forecasts. For example, the International Health Regulations and the Sanitary and Phytosanitary Agreement authorizes trade interdictions as a first response to new information on infectious human, animal or plant diseases. They require such interdictions to be removed within a relatively short time unless there is statistically robust evidence of the risk posed by the commodities in question (World Health Organization, 2005; World Trade Organization, 1995). Similarly, the US Food and Drug Administration implements drug recalls in response to cases of adverse response. Such recalls are followed by hazard assessments that determine the risk associated with the recalled drug (US Food and Drug Administration, 2010). In this way, learning and adaptation decrease the potential magnitude of wasted resources.

Nevertheless, the opportunity cost of unnecessary protective measures during the learning phase under a maximin response could still be very large. Anticipating this concern, most statements of the precautionary principle recognize the need to constrain policy costs to be proportional—in some sense—to the value of protection attained.<sup>2</sup> For example, a recent formulation by UNESCO (2005) states: “When human activities may lead to morally unacceptable harm that is scientifically plausible but uncertain,

<sup>2</sup> In the paper, the term proportionality denotes a consideration of cost when choosing regulations to reduce exposure to particular risks. Our use of the term is distinct from its usage in European law, where the “principle of proportionality” invokes a generalized cost–benefit framework for constraining precaution (Wiener, 2003; Emilou, 1996).

actions shall be taken to avoid or diminish harm. . . . Actions should be chosen that are *proportional* to the seriousness of the potential harm” (emphasis added). An appealing feature of a criterion that maximizes expected utility is that it balances *ex ante* costs and *ex ante* benefits—i.e., expected costs and expected benefits—in a natural and compelling way. When a decision is to be made under (Knightian) uncertainty, however, the appropriate interpretation of a proportional response is less obvious.

A version of proportionality that is compelling when probabilities cannot be assigned across contending models—though “models” may themselves be probabilistic—is the decision criterion *minimax regret*. Minimax regret minimizes the worst-case mistake or *regret*, the difference between how well a policy does *ex post* and how well the decision-maker could have done had the true model been known in advance. As we show in Section 3, regret can be decomposed into additive components that align with the two main types of mistakes that decision-makers are apt to make when setting environmental policy under uncertainty: ‘false negatives’ (where scientific warnings are initially ignored but turn out to be valid) and ‘false positives’ (where scientific warnings are heeded but turn out to have been unnecessary). Minimax regret guards symmetrically against these mistakes and so implements a proportional response.

But a criterion that weights false negatives and false positives equally seems to miss part of the intended meaning of precaution—mainly, to protect in a heightened way against unacceptable outcomes or, conceived differently, to protect in a heightened way against the incidence of false negatives. Maximin achieves this conservatism in a strong way. In fact, under plausible assumptions on the structure of expert disagreement—specified in Section 3.1 (and in the appendix)—maximin can be shown to guard exclusively against the false negative mistake—it always avoids the mistake of damaging the environment excessively, but it also always makes the mistake of wasting resources (the false positive mistake). The sharp contrast between maximin and minimax regret suggests a useful structure within which to organize candidates for a precautionary response. If precaution is identified with the desire to avoid false negatives and proportionality with the desire to balance regard for false negatives and false positives equally, then maximin and minimax regret comprise the extremes in a continuum of plausible options for a precautionary response.

To characterize policies in this continuum, we introduce an asymmetric version of minimax regret that includes a weighting parameter to control the relative weight given to the false-negative mistake. By increasing the weight above one, the criterion traces a continuum of policies between minimax regret and maximin. Policies generated under intermediate weights prioritize the false-negative mistake while still preserving an explicit regard for proportionality. The social objective implemented in this way is consistent with a line of reasoning used by the European Commission to motivate precautionary thinking. In the report *Late Lessons from Early Warnings*, Harremoës et al. (2001) review case studies in which early scientific evidence of potentially severe threats to the environment or public health were ignored, with lamentable consequence. The report concludes that these lessons of history justify a guideline for regulating environmental threats that prioritizes concern for false negatives when severe and irreversible harm may result.

Given a range of precautionary alternatives generated in this way, we employ visual tools to convey the associated *ex ante* trade-offs—the extent to which higher protection against false negatives implies greater exposure to the incidence of false positives, and vice versa. Our aim is not to identify the precautionary principle with a particular value of the weighting parameter, but to communicate information about plausible alternatives for a precautionary response and about the trade-offs associated with choosing among them.

Other approaches for implementing the precautionary principle while preserving a notion of proportionality are suggested in the literature. Lempert and Collins (2007) compare options for implementing a precautionary response in a shallow lakes model where there is uncertainty about the regime-shift threshold. One of the criteria they consider takes a weighted average of expected welfare, evaluated with respect to a best-guess probability distribution over threshold values, and worst-case regret. Lange (2003), similarly, considers maximin expected utility—a weighted average of expected welfare and the worst-case outcome. Both approaches trade-off optimality under a reference model for reduced sensitivity to model misspecification, thus balancing precaution with proportionality. In a related approach, Barrieu and Sinclair-Dégagné (2006) provide an axiomatic formalization of a version of the precautionary principle that selects for robust policies across a set of scientific models. An advantage of our approach is that it can be implemented even when stakeholders disagree about the immediate implications of available scientific evidence.

To illustrate the suggested tools, we consider the problem of setting global climate policy using the Dynamic Integrated Model of Climate and the Economy (DICE) (Nordhaus, 2008). In our simple exercise, expert disagreement about climate science is assumed to be captured by disagreement about climate sensitivity. This disagreement is disciplined initially using the IPCC “plausible range” (IPCC, 2007). We subsequently consider the implications of relaxing the criterion for model selection. The application is also extended to accommodate learning.

## 2. Elements of a precautionary decision

Coordinating the actions needed to avoid novel environmental threats is complicated at the international level because appropriate institutions are often lacking (Walker et al., 2009). Whatever the institutional framework, however, the decision process involves four stages: identification of early warnings (the precursors of change) through some form of horizon-scanning activity, establishment of the case for action, and—after policy has been set—adaptation over time as the scientific basis for action changes. Fig. 1 situates the four components of a precautionary response.

### 2.1. Horizon scanning

Horizon scanning activities provide the evidence base for anticipatory decisions. Many national governments have offices responsible for horizon scanning. The UK Department for the Environment, Food and Rural Affairs, for example, defines the objectives of its own horizon scanning activities as “The systematic examination of potential threats, opportunities and likely future developments which are at the margins of current thinking and planning. Horizon scanning may explore novel and unexpected issues, as well as persistent problems or trends”.<sup>3</sup> But horizon scanning is a much more widespread dimension of strategic planning at all levels. In many cases it is based on breaking scientific results. Indeed, it may be interpreted as the activity by which the scientific community draws attention to results that potentially have significant consequences for human well-being, and that warrant action—whether research or policy.

Horizon scanning may be a precursor to targeted assessments, research initiatives, policy initiatives or to either private or public defensive responses. In the former US Office of Technology Assessment (OTA) it was linked to the preparation of targeted assessments. More generally, however, it is the basis for

<sup>3</sup> <http://horizonscanning.defra.gov.uk/>.

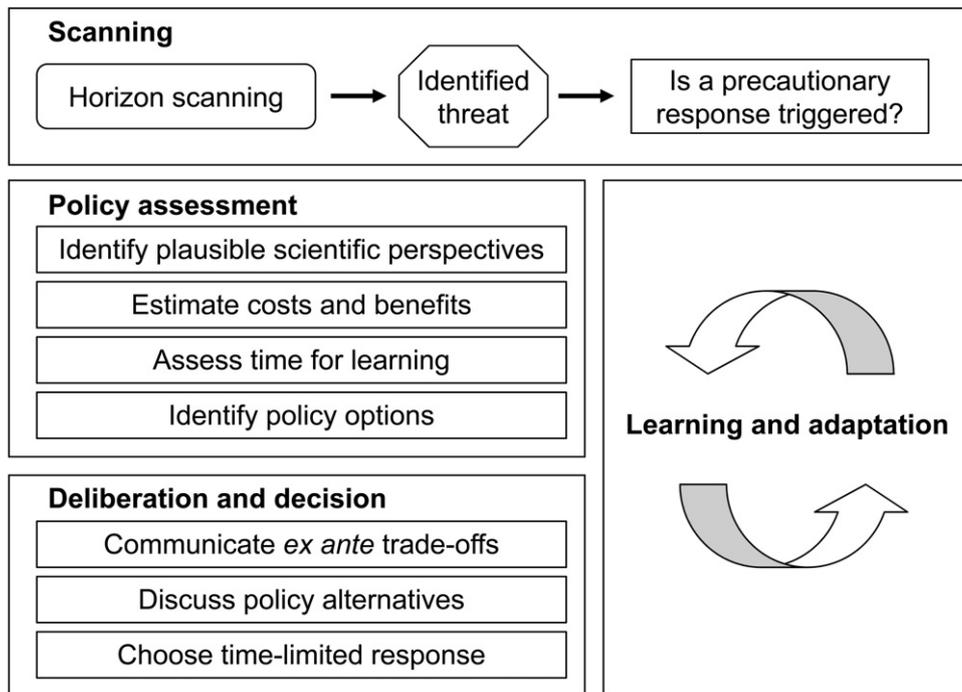


Fig. 1. Elements of a precautionary decision.

determining what forward action to take in cases where there is both novelty and uncertainty. At the international level there are few examples of horizon scanning activities. However, a number of intergovernmental organizations, international non-governmental organizations and multilateral environmental agreements have horizon scanning functions. In addition, some of the global assessment bodies, such as the Intergovernmental Panel on Climate Change and the Millennium Ecosystem Assessment, include a forward-looking component.

The proposed Intergovernmental Science–Policy Platform on Biodiversity and Ecosystem Services, recently approved by the UN General Assembly, will also have a horizon-scanning element. Its functions are expected to include, for example, a responsibility to proactively alert the international community to emerging issues (issues ‘new to science’) in order to allow timely responses.

For the information generated through horizon scanning to trigger a precautionary response, a number of attributes should be present:

- It should not support conventional risk analysis
- It should indicate the potential for severe and irreversible harm to the environment or human health—an outcome deemed “unacceptable” (UNESCO, 2005)
- It should be scientifically plausible

Extreme, unique, rare and irreversible events all have either no or only a few historical precedents. It may not be possible to predict a probability distribution of outcomes with confidence. If decision-makers are to act they need sufficient information. That is, they need to be able to recognize and respond to early warning signs. The review of the evidence on precaution by Harremoës et al. (2001) addresses the problem of early warnings. It considers how governments have assessed the likelihood of uncertain outcomes that might include irreversible, significant and widespread effects. This implies a weaker test than that required by standard scientific proof. Historic examples of such tests include ‘scientifically based suspicion’, ‘reasonable grounds for concern’ and the ‘balance of evidence’. All are based on data that fail standard tests of scientific

proof, but nevertheless provide a basis for decision-making. They are early warning signs rather than conclusive proof.

The source of ‘scientifically based suspicion’ or ‘reasonable grounds for concern’ is frequently a model of the natural or social system involved that is used to simulate possible futures on the basis of experimental research. Indeed model predictions provide the most important early warnings. Modeling offers a transparent basis for predictions. If preliminary theoretical research indicates the possibility of outcomes that involve significant irreversible cost, then the precautionary response involves actions that avoid those costs while using data to test the model predictions.

The early warnings in the form of isolated experiments or events do not provide a basis for estimating a probability distribution of outcomes. A common feature of the case studies reviewed is novelty. In the case of halocarbons, polychlorinated biphenyls (PCBs) and methyl tert-butyl ether (MTBE), the report argues that novelty, plus the persistence of the chemicals in the environment and their dispersion rate should have indicated a potential problem. Indeed, evidence of novelty, persistence and bioaccumulation are argued to provide a sufficient early warning to trigger precautionary action.

An equally common feature of the cases reviewed in that report is that early experimental results did not in fact trigger action. The results of animal tests of PCBs were available in 1937, for example, but were ignored. The contamination of the Great Lakes illustrates the point. Once it had been shown that the lakes contained organochlorine pesticides in the 1960s, research followed showing the bioaccumulation in birds of dichlorodiphenyltrichloroethane (DDT) and PCBs. The accumulation of evidence against DDT led to its banning both in the US and Canada by the early 1970s. But action against PCBs has been slower and weaker. A series of studies of the human development impacts of exposure to PCBs showed significant development deficits, but the evidence has been regarded as too weak to justify precautionary action. Current levels of PCBs in the Great Lakes are still two orders of magnitude higher than the level established to protect humans against cancer risk (Gilbertson, 2001).

Alongside grounds for action, a precautionary decision requires a legal and an institutional framework that allows action to be taken. This poses an acute problem at the international level where, even if the scientific community is persuaded that hypothesized global change outcomes are realistic, there is frequently no responsible authority to take precautionary action pending the research needed to test the hypotheses (Walker et al., 2009). While global human health hazards are covered by a multilateral agreement that authorizes precautionary action (the International Health Regulations), climate, biodiversity and sea areas beyond national jurisdiction are not.

## 2.2. Policy assessment

The policy assessment stage develops the inputs for decision analysis. This entails a description of scientific uncertainty, estimates for costs and benefits, and a forecast of anticipated prospects for learning. If scientific ‘uncertainty’ can be quantified probabilistically, then conventional risk analysis methods should be used. But if evidence and knowledge are insufficient for probabilities to be plausibly specified, then policy choice involves decision-making under uncertainty. To proceed, a set of qualifying scientific perspectives (models) must first be identified. Models themselves may be probabilistic even though uncertainty is assumed to hold across model alternatives. The qualifying set of models is constructed by identifying those that satisfy a specified threshold of plausibility.<sup>4</sup> This determines the extent to which minority scientific views are admitted into consideration and should balance scientific plausibility with the need to respond to potential early warnings.

The next step is to translate environmental effects and policy costs into commensurate values. This task raises a number of familiar issues in cost–benefit analysis. One difficult but essential challenge is to ensure that valuations reflect all direct and indirect impacts, whether reflected in market prices or not. Where novel change affects international public goods, for example, this requires assessment of impacts from a global perspective. Assessment of national impacts alone will understate the full collective benefits of the international public good. This may occur because the mandate of the treaty is too narrow or because the threat is not viewed with sufficient urgency to instigate collective action at a truly global scale.

Finally, the time needed for uncertainty to be resolved—or the expected rate of learning—determines the period over which precautionary policies restrict activities. This significantly affects the expected cost of policy. For threats like climate change, where policy analysis must forecast an evolving complex system and relevant outcomes unfold over centuries, severe uncertainty is likely to persist for decades (Allen and Frame, 2007). By contrast, for new infectious diseases, like SARS or Swine Flu, uncertainty may be resolved in a matter of weeks.

## 2.3. Learning

The last component of a precautionary approach (we discuss ‘deliberation and decision’ below) is a mechanism for learning. Since uncertainty implies that there is scope for learning, every problem that is a candidate for a precautionary approach is also a candidate for learning. A precautionary decision-making process requires a mechanism for updating the information on which decisions are taken. Indeed, the defining characteristic of a precautionary approach is that it enables learning without risking the most severe (but uncertain) consequences of an action.

Since precautionary action is triggered by the potentially severe costs of novel policies, activities or events, learning should establish whether an action does not involve such costs. An important question to ask of precautionary actions, therefore, is whether people undertaking activities with uncertain outcomes have an incentive to inform themselves of those outcomes. Harremoës et al. (2001) identify both institutional and disciplinary barriers. Particular problems may be ‘captured’ by particular disciplines or institutions, with the result that important aspects of the problem are simply ignored. Examples include asbestos and ionising radiation, captured by medical clinicians who focused only on the immediate acute effects of exposure, or bovine spongiform encephalopathy (BSE), captured by veterinarians who regarded the ‘risk’ of transmission to humans as ‘acceptably slight’.

But the incentive to learn about the consequences of novel activities is also closely related to the burden of proof. If someone undertakes an activity thought to have harmful effects on others, the legal burden of proof generally rests with those who maintain that the effect exists. The legal burden of proof does not rest with those who maintain that the effect does not exist. The link between smoking and lung cancer is a good example of this, but most environmental effects come into the same category. The burden of proof rests with society, rather than with those whose activities are the source of potential damage. As with grounds for action, therefore, it follows that learning also depends on the establishment of an appropriate institutional and legal structure.

## 3. Deliberation and decision

To clarify the options facing decision-makers, we think of policy choice as an ‘expert panel problem’ (Woodward and Bishop, 1997). Policy-makers face a panel of experts, and each expert prescribes a distinct scientific model. Models may themselves be probabilistic, but policy-makers have no basis for assigning probabilities across model alternatives. This gives rise to a decision under (Knightian) uncertainty. Without probabilities across alternative scientific models, the notion of proportionality is not obviously defined. This begs the question: What is meant by a precautionary and proportional response?

### 3.1. Trading off environmental harm and the costs of precaution

In the presence of uncertainty, a “proportional response” (intuitively) would balance an *ex ante* measure of policy benefits with an *ex ante* measure of policy costs. When probabilities are available, these *ex ante* considerations are captured by expected benefits and expected costs, defined with respect to the specified probabilities. But when probabilities are unavailable, the trade-off that a proportional response should balance is less obvious.

We propose a notion of proportionality that follows from orienting the decision around trade-offs between two competing objectives that decision-makers may pursue when setting environmental policy under uncertainty. One is to reduce exposure to the incidence of “false negatives”; the other is to reduce exposure to the incidence of “false positives”. False negatives occur when scientific warnings turn out to be valid but are initially ignored (or inadequately addressed). We measure their magnitude by the extent to which environmental damages turn out to exceed what would have been optimal (in cost–benefit terms) in hindsight. False positives, by contrast, occur when early scientific warnings are heeded but turn out to have been unwarranted. We measure their magnitude by the extent to which resources turn out to have been wasted on unnecessary protective measures. Both mistakes depend on the policy chosen *ex ante* and the scientific model that turns out to be true *ex post*. There are trade-offs because the objectives conflict: the more resources society commits to

<sup>4</sup> In the application, we make use of summary statements from the IPCC Fourth Assessment Report (IPCC, 2007) to “discipline” the set of admitted models.

protect the environment (and so decrease the incidence of false negatives), the more it increases the incidence of false positives—and vice versa.

This trade-off frames the tension between proportionality and precaution. A strictly proportional policy would weight the objectives equally, guarding in a symmetric way against false negatives as against false positives, while a strictly precautionary policy would guard foremost against false negatives. Under conditions elaborated below, two standard criteria for decision-making under uncertainty can be shown to implement these approaches directly. A strictly proportional approach is accomplished by *minimax regret*, while a strictly precautionary approach is accomplished by *maximin*. To keep the following discussion precise, we introduce some analytical structure. A more rigorous presentation of the propositions discussed in this section is provided in [Appendix A](#), and the interested reader is encouraged to visit the material there.

Policy-makers are assumed to face a finite set of experts—we denote this set  $M$ —and each expert prescribes a distinct scientific model,  $m$ . The true model is unknown at the time of the decision, so policy-makers choose policy *ex ante*; this policy is denoted  $a$ . We later interpret  $a$  as the level of abatement in a pollution control problem. If the problem has more than one period,  $a$  would denote a sequence of abatement levels, one for each period in the planning horizon. The *ex ante* policy  $a$  is distinguished from the optimal policy given a particular model  $m$ , with the latter denoted  $a^{*,m}$ .

Policy-makers evaluate outcomes with respect to a welfare function,  $W(a,m)$ . Welfare depends on environmental policy as well as on the scientific model that turns out to be true. These objects together yield a dynamic path of environmental variables. The DICE model used in the application provides an example. If expert models are probabilistic, then  $W(a,m)$  denotes expected welfare in the von Neumann–Morgenstern sense (von Neumann and Morgenstern, 1944).

A related outcome measure is regret, denoted  $R(a,m)$ . Regret is the welfare consequence of error: the difference between the best policy-makers could have done had they known the true model in advance and how well they actually do. It is a function of the policy  $a$  chosen *ex ante* and the model  $m$  that turns out to be true *ex post*:

$$R(a, m) = W(a^{*,m}, m) - W(a, m).$$

Regret is useful for our purpose because, as shown in sections one and four of the appendix, it can be decomposed into components that align with the false-negative and false-positive mistakes above. In particular, we can write

$$R(a, m) = R_1(a, m) + R_2(a, m),$$

where  $R_1(a,m) > 0$  implies that environmental damages turned out to exceed what would have been optimal in hindsight—the false-negative mistake—and  $R_2(a,m) > 0$  implies that resources were wasted on protective measures that turned out to have been unnecessary—the false-positive mistake.

The decision criterion ‘minimax regret’ selects a policy  $a$  in the following problem:

$$\min_a \max_{m \in M} R(a, m).$$

The criterion chooses  $a$  to minimize regret, but it does so anticipating that given whatever policy is chosen an “adversarial agent” will respond by choosing the regret-maximizing model from the panel of experts. If a low-abatement policy is chosen, the adversarial agent would respond by selecting a high-environmental-damage model, leading to a large false-negative

mistake, while if a high-abatement policy is chosen, the agent would select a low-damage model, leading to a large false-positive mistake. The criterion balances these contrasting effects, seeking the *ex ante* policy that minimizes maximal regret. Because regret is an equally weighted sum of the two regret components— $R = R_1 + R_2$ —the criterion guards against false negatives and false positives symmetrically. In this sense it implements a strictly proportional response.

In contrast, the decision criterion maximin maximizes welfare under the worst-case model. It selects a policy to solve<sup>5</sup>

$$\max_a \min_{m \in M} W(a, m).$$

Under conditions on the set of expert models described below, maximin guards exclusively against the false-negative mistake and so implements a strictly precautionary response in the following sense. With the maximin policy chosen *ex ante*, the false-negative mistake is always avoided ( $R_1 \leq 0$  for each expert model); at the same time, the false-positive mistake always occurs ( $R_2 > 0$  for each expert model), except when the worst-case model turns out to be true, in which case there is no mistake ( $R_2 = 0$ ). Panel 1 of [Fig. 3](#) demonstrates this property of maximin graphically. A proof of the result is provided in [Appendix A.2](#).

The above interpretation of maximin holds under two possible restrictions on the set of expert models.<sup>6</sup> First, it is sufficient for expert disagreement to be restricted to the damage side of the cost–benefit equation (including both the damage function and the dynamic scientific models of cumulative environmental effects). This does not require policy costs be known with certainty, only that policy-makers are able to quantify perceived “uncertainty” with respect to abatement costs by agreeing on a common probability distribution over contending models of abatement costs. We argue in the next section that this may be an appropriate assumption for climate policy.

A second sufficient condition applies when experts disagree about both sides of the cost–benefit equation, but the set of expert models is restricted in such a way that high-environmental-damage models are paired with low-abatement-cost models, while low-environmental-damage models are paired with high-abatement-cost models. We argue in the next section that this pairing may be plausible if the political process underlying the decision proceeds in such a way that interested stakeholders choose expert specifications endogenously. In this case, stakeholders inclined to protect the environment may naturally select models with high environmental damages and low abatement costs, while stakeholders inclined to protect industrial industries may favor models that reflect low environmental damages and high abatement costs.

Under either of the above restrictions, minimax regret and maximin can be regarded as the extremes in a range of plausible candidates for implementing a precautionary response to the expert panel problem. Intermediate to these, one can imagine a continuum of alternatives that guard against the incidence of false negatives more than equally, while still preserving some notion of proportionality. Criteria from this middle ground may offer more interesting candidates for a precautionary response than either maximin or minimax regret.

To operationalize an intermediate response, it is helpful to think of a decision criterion that includes a weighting parameter to control the relative weight on the competing objectives. We employ an asymmetric version of minimax regret, as suggested in

<sup>5</sup> If expert models are probabilistic, the criterion becomes maximin expected utility as prescribed in [Gilboa and Schmeidler \(1989\)](#).

<sup>6</sup> In fact, the interpretation holds under other circumstances also, but the property would need to be verified on a case-by-case basis if the sufficient conditions defined here are not satisfied. See the appendix for more details.

Iverson (2011). The criterion guards against the worst-case realization of ‘asymmetric regret’, a generalized version of regret that weights false negatives unequally. It solves

$$\min_a \max_{m \in M} R^A(a, m),$$

where  $R^A(a, m)$  equals<sup>7</sup>  $\kappa_E R_1(a, m) + R_2(a, m)$ , whenever  $R_1$  is positive, otherwise  $R^A(a, m)$  equals  $R_2(a, m)$ . We restrict  $\kappa_E$  to be greater than or equal to one as this captures the intended meaning of precaution. The described criterion weights  $R_1$  more than equally (by  $\kappa_E \geq 1$ ), but only when  $R_1$  turns out to be positive. The reason is because  $R_1$  negative is not in itself a mistake— $R_1$  negative means the *ex ante* policy protected the environment more than what turned out to have been optimal from a cost–benefit perspective; if this fact implies  $R_2 > 0$ , then that would be a mistake, but  $R_1$  negative on its own is not.

The criterion spans minimax regret and maximin. This means that it reduces to each criterion as a special case, while also identifying a continuum of intermediate alternatives. If decision-makers set  $\kappa_E$  to one, the criterion reduces to minimax regret. As  $\kappa_E$  increases, the abatement policy chosen by the criterion increases monotonically, and for  $\kappa_E$  sufficiently large the chosen policy eventually coincides with the maximin policy. A large  $\kappa_E$  tells the criterion to guard exclusively against the incidence of false negatives, which mimics the approach of maximin. These properties are demonstrated in Appendix A.

As mentioned already, the social objective implemented under intermediate weights on the false-negative mistake is consistent with a line of reasoning used by the European Commission to motivate precautionary thinking. Harremoës et al. (2001) emphasize that policy-makers should “focus on the potential costs of being wrong” and conclude that false negatives warrant greater concern than the possible incidence of false positives. Our goal is to communicate the options for trading-off exposure to these contrasting dimensions of “risk”. We use the suggested criterion to identify policies that balance the competing objectives in a range of possible ways. We then plot, for each policy option, the magnitude of each mistake across the range of models on the expert panel. These visual displays make the *ex ante* trade-offs across policy alternatives explicit.

To this point the decision has been posed as a static problem. Policy-makers choose the entire future path of abatement in the initial period given uncertainty specified at that time. This formulation is ultimately unsatisfactory because it ignores future learning, an essential feature of decisions involving scientific uncertainty. To accommodate learning, the application in Section 4 is extended to consider a sequential version of the expert panel problem in which decision-makers anticipate that future policy will respond to improvements in scientific knowledge that become available in the future.

Note that the weight attaching to the environmental outcome may be time varying. The most common examples of precaution in the protection of public health—product recalls—can be interpreted as decisions in which a novel outcome (e.g. product failure or health response) triggers a reaction that is initially highly precautionary, but is relaxed as information on the true likelihood of the outcome emerges. Indeed, the argument for prioritizing risks to the environment or human health is strongest when policies are viewed as temporary measures—providing an opportunity for learning to occur and preserving flexibility so future policy can adapt. Without learning and adaptation, maximin is hard to justify because it entails a significant risk that resources will turn out to be wasted. However, when learning is expected to unfold rapidly and there is risk of severe and irreversible harm, maximin may provide

a sensible rule of thumb for a time-limited response. Recent examples of temporary, maximin responses occurring under the International Health Regulations are the closure of Toronto in response to SARS and the movement restrictions in response to Swine Flu.

Of course, maximin is just one option from a range of alternatives. The extent to which the environmental objective should be prioritized depends on the identified case for precaution—for example, the risk of severe and irreversible harm—as well as the magnitude of trade-offs across the panel of expert models.

### 3.2. Preferences for a high level of protection

The decision to choose a high level of protection with respect to specific risks hinges on the preferences of the decision-maker (or the people they represent). Empirical evidence on choice behavior under uncertainty suggests at least two explanations for cautious behavior. One explanation is based on *loss aversion*, the tendency for people to weight “relative losses” more acutely than “relative gains”. Loss aversion is a component of Prospect Theory, a generalization of expected utility theory constructed to match empirical evidence that uncertain choice behavior is highly sensitive to the reference point of the decision maker (Kahneman and Tversky, 1979). Applied to long-term policy decisions like climate change policy, loss aversion suggests that people would tend to weight the risk of damaging the environment more heavily than the competing risk that resources might be wasted.

The argument presumes that the reference point of the decision maker setting climate policy is the current state of the system and current consumption. Relative to this reference point, climate change entails a relative loss, while increased abatement reduces future consumption growth, i.e. involves a diminished gain. In experimental studies, people weight relative losses about 2.25 times more than they do relative gains.

A second explanation is based on people's tendency to overweight low-probability, high-consequence events relative to what expected utility theory would suggest is rational (McDaniels et al., 1992). As a result, potentially catastrophic or ‘dreadful’ outcomes may receive disproportionate attention in the policy process, leading to a preference for precaution. Whether or not such fears are warranted by the data, they have an important place in collective decision-making and should be taken seriously. At the same time, where precaution is triggered by initially high weights accorded to the worst case, it is important to generate the information needed for people to adjust the weight they assign to the worst case.

## 4. Application to climate policy

Climate change is a global threat that satisfies many of the conditions for which a precautionary approach is indicated. The general circulation system is immensely complex and relevant data are limited, making it difficult to assign probabilities across alternative forecasting models (Manning, 2006). In addition, damages are potentially irreversible and severe. With a half-life over 80 years, carbon dioxide emissions are irreversible on a time frame that is economically significant. Moreover, a variety of irreversible environmental effects are possible including, for example, species extinctions. Even without environmental irreversibilities, the damages associated with climate change are expected to be significant. Nordhaus (2008), on the basis of conservative assumptions, estimates that unabated climate change would lead to a 2.5 percent annual loss in global output by the end of this century. These effects are widespread and public, with future generations bearing the brunt of damages. Use of a

<sup>7</sup> The “E” subscript on  $\kappa$  highlights the fact that the “environmental mistake” is the one weighted more heavily.

precautionary approach is also indicated by the terms of the UN Framework Convention on Climate Change.

#### 4.1. Communicating trade-offs in DICE

To illustrate a precautionary and proportional approach to climate policy decision-making, we employ Nordhaus's 2007 DICE model. DICE augments a neoclassical growth model to include a three-box model of the global carbon cycle, along with reduced-form, functional representations of climate damages and abatement costs. The model is relatively simple by climate science standards, but sufficiently rich to capture the essential features of the global policy decision.

To accommodate the proposed decision framework, we modify DICE in one important way. The standard version has two control variables: greenhouse-gas abatement and physical-capital investment. The rationale for controlling physical-capital investment is that it generates an investment sequence that is consistent with a perfect-foresight, competitive equilibrium with complete markets and internalized externalities. However, the interpretation hinges on the planner basing policy selection on a decision criterion that maximizes expected welfare. To preserve the intended interpretation when adopting alternative decision criteria, the model is modified so that abatement is the only control, while physical capital investment is fixed at the corresponding optimum for any given sequence of abatement.<sup>8</sup> Solving the model in this way preserves the competitive equilibrium interpretation of the standard model.

We calibrate the model following the baseline assumptions in Nordhaus (2008). Importantly, the time discount rate is jointly calibrated with the consumption elasticity to match historical data on the real return on capital. This is accomplished by setting the time discount rate equal to 1.5% and the consumption elasticity equal to 2; these parameters imply a real return on capital of about 5.5%. The choice of discount rate is controversial in applications to climate policy because damages are concentrated in the distant future. At a discount rate of 5.5%, \$1000 worth of damages occurring 100 years from now is worth \$5 today. Recent discussions of climate policy analysis that argue for a lower discount rate include Stern (2006), Heal (2008), and Gollier (2010). The model also does not explicitly account for the possibility of severe threshold events like the potential collapse of the Thermohaline Circulation in the North Atlantic. Augmented versions of DICE that explicitly model a catastrophe threshold include Keller et al. (2004) and Lemoine and Traeger (2010).

To illustrate our approach in a straightforward manner, we assume that scientific disagreement is restricted to disagreement regarding climate sensitivity, the equilibrium global-mean temperature change that would result from a doubled concentration of carbon dioxide relative to its pre-industrial level. This is of course a considerable simplification, and other forms of uncertainty could be added. To discipline the range of disagreement regarding climate sensitivity, we employ the summary conclusions from the Fourth IPCC Assessment Report (IPCC, 2007). The report defines the "plausible range" for climate sensitivity to be between 2.0 and 4.5 °C. This range reflects most-likely values, not the full set of possible values. In the next section, we consider the implications of relaxing the standard for model selection.

<sup>8</sup> We use an iterative approach to solve the model in this way. First, the abatement sequence that solves the (asymmetric) minimax regret problem is determined for a particular sequence of savings rates. Then the optimal savings rate is determined when this abatement sequence is fixed. In the next iteration, the optimal savings rate from the prior iteration is used to determine a new asymmetric-minimax-regret abatement sequence. The procedure is repeated until the savings rate converges.

As explained in the last section, abatement cost uncertainty could be included either by assuming a probability distribution over alternative abatement cost functions,<sup>9</sup> or by choosing expert models in a way that aligns high climate damage models with low abatement cost models (and vice versa). The former approach entails an asymmetric treatment of uncertainties with damage-side uncertainty assumed to be more severe. This is a plausible assumption in climate policy for at least three reasons. First, the chain of reasoning needed to specify uncertainties is longer on the damage side where it is necessary to link greenhouse gas concentrations to the value of climate-change impacts net of adaptation. Second, while abatement costs are naturally denominated in monetary units, climate damages are very hard to quantify in terms of willingness to pay. Among the many well-known difficulties in this area, valuation of distant future climate damages depends on speculative assumptions about the preferences of generations not yet born. Finally, abatement costs are front-loaded relative to climate damages, and uncertainty compounds as the distance in the future increases.

Even though damage-side uncertainty is more severe, it may sometimes be appealing to treat both categories of uncertainty without probabilities. This is especially true when confronting the challenge of negotiating a compromise among stakeholders with conflicting interests and values. Coincidentally, the model restrictions needed for our framework to accommodate Knightian uncertainty with respect to both damages and abatement costs are easiest to defend in this contested-decision setting. In particular, if combined expert opinions are chosen (or recommended) by interested stakeholders, then one would expect "pro environment" stakeholders to select models with both high environmental damages and low abatement costs (perhaps reflected in an emphasis on green jobs) and for "pro industry" stakeholders to select models with both low environmental damages and high abatement costs (perhaps reflected in an emphasis on the economic risks of government regulation). Consequently, endogenous model selection would lead to precisely the pairing of damage and cost models needed for our framework to apply. The treatment of abatement cost uncertainty without probabilities may be appropriate when aggressive abatement paths are considered. In DICE, optimal policies approach 100% abatement sometime next century. Reductions of this scale will require large-scale transformations of our energy and transportation systems, and the cost of this transformation will depend on institutional, technological, and human behavioral changes that are difficult if not impossible to predict.<sup>10</sup>

Our specification allows for an arbitrary set of experts with climate sensitivity between 2.0 and 4.5 °C. Nevertheless, we find numerically that interior experts are redundant for each of the considered criteria in the sense that the recommended action is the same when experts interior to the model space extremes are excluded. This simplifies the analysis by allowing us to focus on the model extremes; in this case, the 2.0° model and the 4.5° model. Redundancy of interior models is a common property of minimax regret (and asymmetric minimax regret) though it is easy to find examples in which it does not hold. To decompose regret, we use the approach described in Appendix A.4. We numerically solve the

<sup>9</sup> The abatement cost function in DICE takes the form  $\theta_1 \mu^{\theta_2}$ , where  $\mu$  is the fraction of emissions to be abated. Allowing for (independent) uncertainty with respect to the linear parameter  $\theta_1$  would not affect the results relative to the case in which  $\theta_1$  is fixed at its expected value. This follows because marginal abatement costs are linear in this parameter, so the expected value of abatement costs would equal abatement cost evaluated at the expected value of the parameter. In contrast, incorporating probabilistic uncertainty with respect to the exponential parameter  $\theta_2$  could have a quantitatively significant impact on the results. The nature of the effect would depend on the specification of uncertainty.

<sup>10</sup> We are thankful to an anonymous referee for recommending this argument.

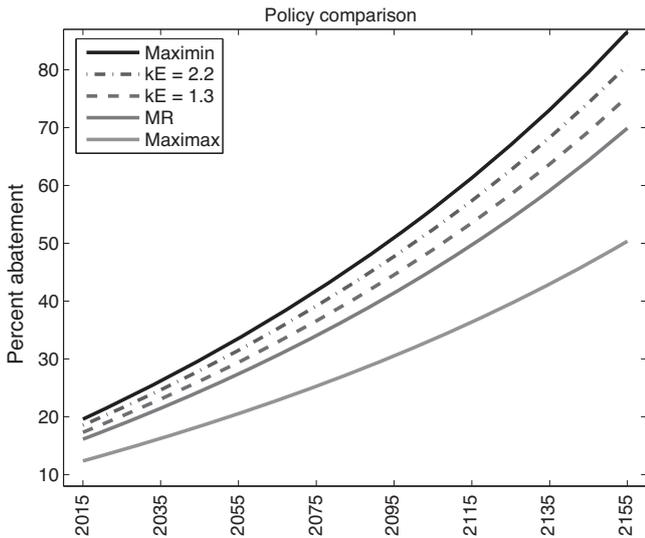


Fig. 2. Abatement paths for a range of asymmetric minimax regret policies.

optimization problem corresponding to each criterion using a sequential quadratic programming algorithm—specifically, we use the ‘trust-region-reflective’ portion of the ‘fmincon’ algorithm in MATLAB. We combine this nonlinear, constrained optimization algorithm with a best-of-N restarts strategy to promote convergence to a global maximizer. We found that the optimization problem for each decision criterion was well behaved and numerically stable. As noted above, the asymmetric minimax regret problem is solved for a fixed path of savings rate, and we use an iterative approach to fix the savings rates at the corresponding optimal path for each sequence of abatement. The iterations in this process converged easily.

With the set of plausible scientific models and relationships for costs and benefits in place, the decision tools suggested in Section 3 are applied under two possible scenarios for future learning. Figs. 2

and 3 show asymmetric-minimax-regret policy alternatives and the corresponding *ex ante* trade-offs constructed under a scenario of no learning—policies chosen today, amid current uncertainty, remain in place forever. Figs. 4 and 5 repeat the exercise assuming that the true scientific model will be learned perfectly at some future date (2075 in the exercise depicted here) with policies adjusted optimally thereafter. As expected, the mistakes are smaller when future policy adjusts in response to future learning. Nevertheless, the policies chosen by the asymmetric minimax regret criterion turn out to be nearly the same under each scenario.

Fig. 2 displays abatement policies generated under a range of values for the relative weight on the false-negative mistake ( $\kappa_E$ —“E” stands for environmental). Decades are shown on the horizontal axis and “fraction of global emissions to be abated” is depicted on the vertical axis. As in the calibrated DICE model, abatement in the initial “period,” the decade centered around 2005, is fixed at 0.5%. The top solid line is the policy generated by maximin—it is the optimal policy given temperature sensitivity of 4.5 °C. The bottom solid line is the policy generated by maximax—the optimal policy given temperature sensitivity of 2.0 °C. These extremes display the range of policies that can be justified, on cost-benefit terms, given the considered range of expert disagreement. The middle solid line is the policy generated by minimax regret, and the dashed lines indicate intermediate policies that weight false negatives more than false positives. The weighting parameters of 1.3 and 2.2 were chosen to generate equally spaced policies. By varying the weighting parameter, the asymmetric criterion traces a continuum of policies between minimax regret and maximin.

Taking the four precautionary policies in Figs. 2 and 3 display the corresponding exposure to the two dimensions of “risk”. Mistakes are shown on the vertical axes in trillions of 2005 US dollars. White bars denote false negatives (or environmental mistakes); black bars denote false positives (the mistake of wasting resources). A negative false negative means that the chosen policy overcompensated, protecting the environment more than would have been cost-effective in hindsight. A negative false positive

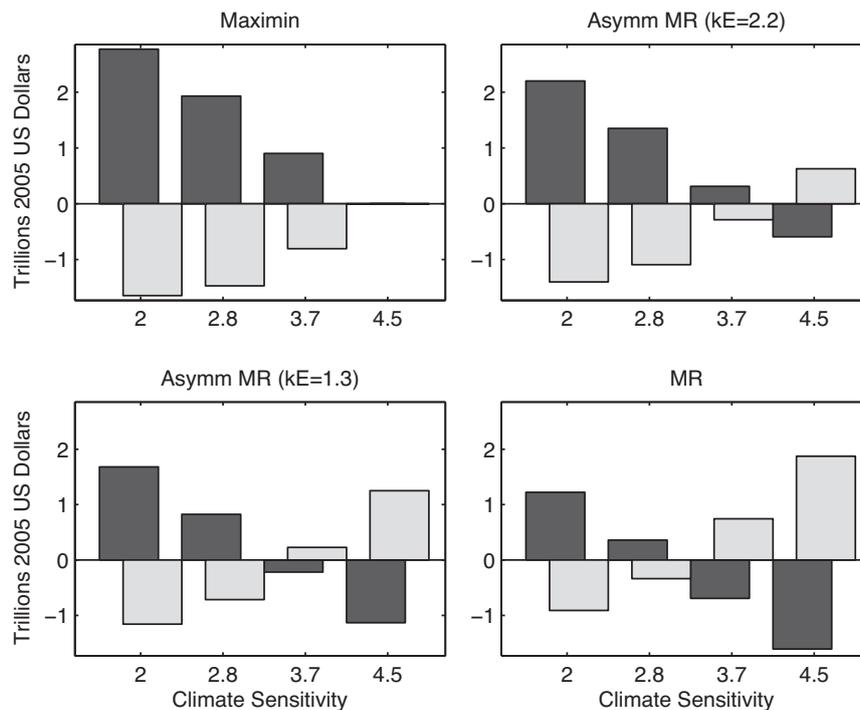


Fig. 3. Component mistakes on vertical axes in trillions of 2005 US dollars; climate sensitivity on horizontal axes. White bars indicate false negatives; black bars indicate false positives. Each panel presents the profile of risks for the indicated policy from Fig. 2.

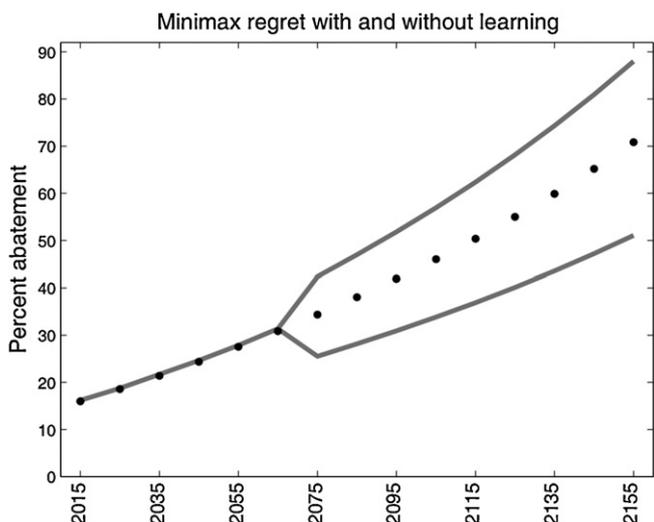


Fig. 4. Black dots indicate the minimax regret policy without learning. The grey fork indicates the minimax regret policy with perfect learning in 2075. The upper branch is the optimal continuation under low damages; the lower branch is the optimal continuation under high damages.

means that the chosen policy spent less on protection than would have been optimal in hindsight. Each panel displays the incidence of the two mistakes across a representative sample of models from the plausible range.

The first panel in Fig. 3 demonstrates the sense in which maximin guards exclusively against the environmental concern. It always avoids the environmental mistake (white bars are never positive), while it also always makes the mistake—at least weakly—of having wasted resources (black bars are never negative). According to the exercise, if the maximin policy is kept in place

from now into the distant future, and if the true value for climate sensitivity turns out to be 2.0 °C, then the world community will have spent, in present value terms, about 2.5 trillion dollars more on greenhouse gas abatement than it would have under the optimal policy. At the same time, it will have reduced climate damages by about 1.5 trillion dollars more than it would have under the optimal policy—again in present value terms. As a point of comparison, world GDP in 2008, calculated at purchasing power parity exchange rates, was about 70 trillion dollars.

Moving to the right in Fig. 3 through less aggressive policies—stepping down the ladder of policies in Fig. 2—the black bars get pushed down and the white bars get pushed up; less aggressive policies reduce the incidence of the false positives, while increasing the incidence of false negatives. This reflects the trade-offs at the heart of the policy decision. Higher levels of environmental protection come at a cost—represented in *ex ante* terms as higher exposure to the incidence of the mistake that resources could turn out to have been wasted.

The trade-offs in Fig. 3 reflect once-and-for-all mistakes—the policy path chosen in the initial period is assumed to remain in place for the indefinite future. In reality, future decision makers will benefit from better information as it becomes available over time. To accommodate learning in a non-Bayesian setting, we assume that the true scientific model becomes known (perfectly) in some future period and that policy adjusts optimally thereafter. This formulation has the merit of being easily interpretable and computationally manageable, though it is at best an approximation to the actual process by which learning can be expected to occur. The decision problem with anticipated learning is described in more detail in Appendix A.4.

We assume in the exercise that learning occurs in 2075. For comparison, the uncertainty range for climate sensitivity has not changed substantially in 30 years (NRC, 1979; IPCC, 2007). Moreover, Kelly and Kolstad (1999) estimate that 9–16 decades are needed to reject alternative climate sensitivities with 95%

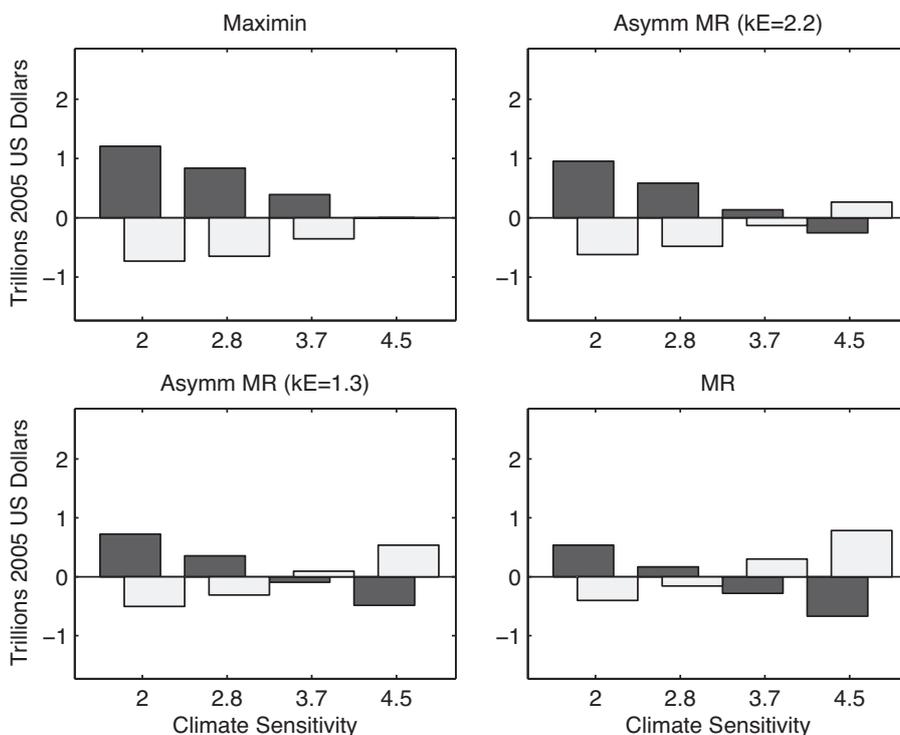


Fig. 5. Component mistakes on vertical axes in trillions of 2005 US dollars; climate sensitivity on horizontal axes. White bars indicate false negatives; black bars indicate false positives. Policies and regrets assume perfect learning in 2075. The vertical lines distinguish between models inside the plausible range and models outside the plausible range.

confidence. Our use of 2075 is meant to illustrate a range of possible effects. Increasing the time to learning would increase the implied magnitude of mistakes; in the limit, as the time to learning is pushed to the distant future, the magnitude of mistakes approaches the case in which no learning occurs—the case depicted in Fig. 3. So trade-offs for the 2075-learning case might be seen as a possible lower bound on the magnitude of trade-offs, to be contrasted with the upper bound from Fig. 3. More appropriately, we would report results for a variety of learning dates, allowing decision-makers to compare the sensitivity of trade-offs to the learning date. We avoid this elaboration due to space constraints, though the exercise can be easily repeated.

With the decision reformulated, one might expect that policies chosen by the various decision criteria for the run-up period before learning occurs would be significantly affected by the anticipation of future adjustments. We find in the context of the application, that this effect (the irreversibility effect) is very small. Fig. 4 compares for minimax regret the generated policy under the learning-in-2075 assumption with the policy generated under the no-learning assumption. The dotted line denotes the policy without learning, and the solid line indicates the policy with perfect learning in 2075. The learning policy forks in the period in which the true model is learned: the upper branch indicates the optimal continuation policy when the true model turns out to have temperature sensitivity of 2.0 °C, and the lower branch indicates the optimal continuation when the true model turns out to have temperature sensitivity of 4.5 °C. So the short-run policy under minimax regret is largely insensitive to the two learning assumptions. Similar results are found for the other decision criteria considered in the paper, and the effect holds even more strongly for learning dates beyond 2075.

A substantial literature has considered the effect of expected learning on optimal climate policy. Hammit et al. (1992) were the first to emphasize the need for a sequential decision framework when comparing near term options. Kolstad (1996) evaluates the irreversibility effect by incorporating an explicit learning process into an optimal growth economy-climate model. He finds that the stock effect for greenhouse gas accumulation is too small to generate a significant learning effect, though irreversibility in control capital investments is significant, causing the optimal abatement trajectory to decrease somewhat relative to the case without learning. Webster (2002) draws similar conclusions, but also shows that the optimal learning effect can be significantly influenced by interactions between near term policy and future marginal costs and marginal damages; for example, incorporating induced innovation or threshold effects for ecosystem dynamics could substantially change the effect of learning. Moving beyond expected utility, Lange (2003) characterizes the irreversibility effect for maximin-expected utility, a criterion that employs a weighted average of maximin and expected utility. The current work differs from this literature by evaluating the effect of learning on minimax regret (and related criteria) in a climate policy decision under Knightian uncertainty. The result in Fig. 4 can be interpreted as saying that the irreversibility effect is insignificant in DICE when minimax regret is used and physical capital investments are reversible. This finding is consistent with the findings in Kolstad (1996) and Webster (2002).

*Ex ante* trade-offs for the exercise with learning in 2075 are shown in Fig. 5. The vertical axes denote trillions of 2005 US dollars, and the horizontal axes indicate climate sensitivity. Comparing Fig. 5 with Fig. 3 indicates the extent to which future adjustments bring the realized policy trajectory closer to the *ex post* model-specific optima, thus decreasing the magnitude of both mistakes. The mistakes when learning occurs in 2075 are decreased by roughly half relative to the case in which no learning occurs. If one expects climate science to reduce uncertainties

significantly within the next century—and for global environmental policy to adjust nimbly in response—then the trade-offs displayed in Fig. 5 are more plausible.

#### 4.2. Relaxing the standard for model selection

Formulating the decision as an expert panel problem relaxes the strict information requirements needed to specify a decision under risk, but restricting attention to a narrowed set of potential models could leave decision-makers blind to early signs of surprise. Where the set of possible outcomes is initially unknown, we refer to 'ignorance' as distinct from uncertainty. Ignorance is an important feature of the decision landscape for novel risks, and dealing with it is an important part of a precautionary response. As emphasized by the editors of the *Late Lessons* report, No matter how sophisticated knowledge is it will always be subject to some degree of ignorance. To be alert to—and humble about—the potential gaps in those bodies of knowledge that are included in our decision-making is fundamental. Surprise is inevitable. ... Thus a key element in a precautionary approach to regulation involves a greater willingness to acknowledge the possibility of surprise. ... [This] leads to greater humility about the status of available science, requiring greater care and deliberation in making the ensuing decisions (Harremoës et al., 2002).

To enable a more effective response to potential surprise, we consider the implications of relaxing the standard for admitting expert models. The scientific literature estimates subjective probabilities for climate sensitivity either by restricting the range of probability density functions consistent with historical observations or by comparing the spread of results across model ensembles. The Fourth Assessment Report presents 5–95% ranges for a variety of such studies, implying a composite range of published values between 1.0 °C and 9.1 °C (IPCC, 2007). This composite range groups together scientific opinions, including minority perspectives, that satisfy the minimal constraint of one-time publication in the peer-reviewed literature along with the 5–95% truncation rule. This is a weaker standard of scientific plausibility than imposed by the IPCC plausible range.

As shown in Fig. 6, relaxing the standard for model selection has a much larger effect on the model space upper bound than it does on the model space lower bound. This is a consequence of the well-known fact that observable quantities, such as transient temperature response, are inversely proportional to climate sensitivity, which is unobservable. It follows that fluctuations in historical data provide less information about the relative likelihood of high climate sensitivity values than about the relative likelihood of low climate sensitivity values, making it hard to rule out very high values (Allen et al., 1996).

Fig. 7 compares precautionary policies generated under the composite range with precautionary policies generated under the plausible range. The dark-grey and mid-grey region between 'MR, plausible range' and 'maximin, plausible range' indicates the set of precautionary policies generated when the decision under uncertainty is posed with respect to the plausible range. This

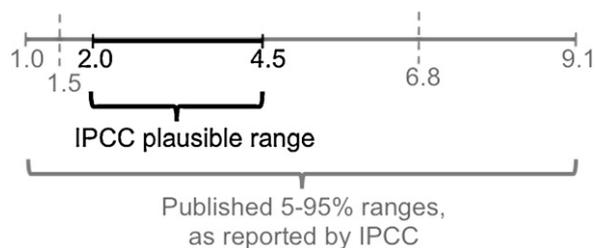


Fig. 6. Expanding sets of climate sensitivity values.

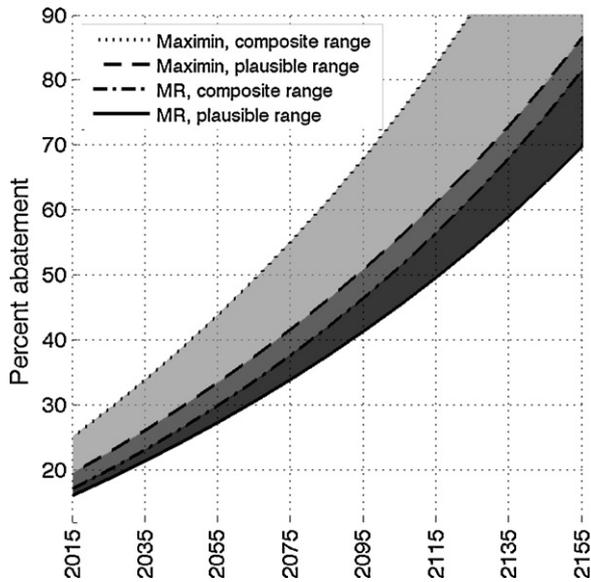


Fig. 7. Overlapping regions show ranges of precautionary policies generated by minimax regret, asymmetric minimax regret, and maximin given different model space assumptions. The dark-grey and mid-grey region indicates the plausible-range policies, while the light-grey and mid-grey region indicates the composite-range policies.

region covers the set of policies shown in Fig. 2. The light-grey and mid-grey region between 'maximin, composite range' and 'MR, composite range' indicates the range of precautionary policies generated when the decision is posed with respect to the composite range. The mid-grey wedge common to both regions indicates overlap.

For any given decision criterion, the composite-range policies are more aggressive than the corresponding plausible-range policies. This is what one would expect since the worst-case environmental mistake increases more than the worst-case growth mistake as the model set is expanded. The two policy sets provide an interesting reference for discussion, though there are good reasons why one might find the plausible-range policies more compelling—at least in the context of the policy exercise posed above. Most importantly, the normative justification for maximin and minimax regret presume that decision-makers have absolutely no basis for assessing the relative likelihood of different models in the model set; without this assumption, the axiomatic foundations for asserting that these criteria provide a “rational” approach to the stated decision break down (Milnor, 1954; Arrow and Hurwicz, 1972). For the IPCC plausible range, this assumption is compelling, but not so for the composite range. While some reputable scientists argue that climate sensitivity values near the extremes of the composite range are possible, none would seriously contend that such values are as likely as values from

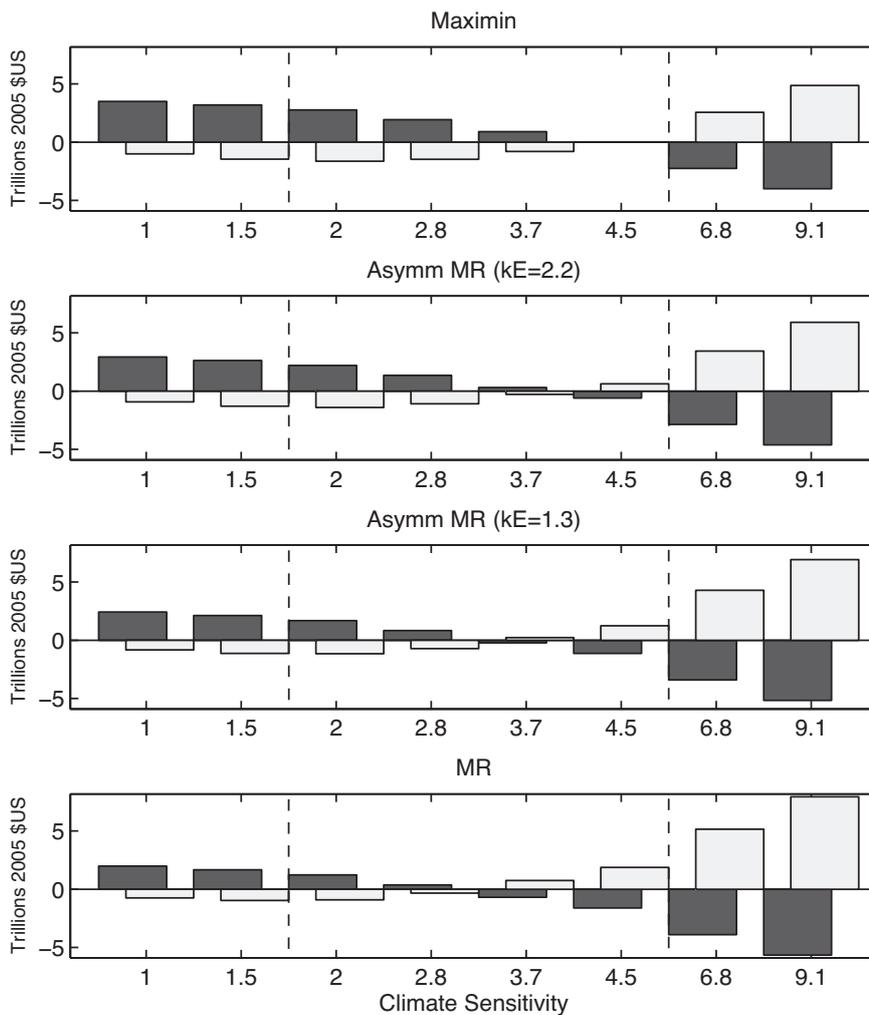


Fig. 8. Component mistakes on vertical axes in trillions of 2005 US dollars; climate sensitivity on horizontal axes. White bars indicate false negatives; black bars indicate false positives. Specified policies correspond to those in Fig. 2. Climate sensitivity models correspond to those shown in Fig. 4. The vertical lines distinguish between models inside the plausible range and models outside the plausible range.

the plausible range (or, more precisely, an interval of values of comparable length).

A further reason for preferring the plausible range policies—at least as a basis for defining initial alternatives and considering the associated trade-offs—has to do with the global nature of the policies under consideration. Policies in DICE prescribe the fraction of global emissions to be abated. Implementation of global abatement targets will require a high degree of international cooperation. But the global public good nature of climate change mitigation, together with the high cost of aggressive abatement strategies, will make achieving this level of cooperation very difficult. For this reason, it may be more realistic to entertain a discussion around the plausible-range alternatives and to find another way to heighten awareness of potential surprise that does not condition immediate abatement targets on highly preliminary evidence.

To this end, we modify our DICE application by expanding the set of models across which trade-offs are communicated. This conveys the implications of preliminary scientific perspectives, while restricting attention to the plausible-range policies. Fig. 8 presents the results. The plots show the implied incidence of the contrasting mistakes across the set of considered models for each considered policy. The inner four bars display trade-offs across the plausible range and thus convey the same information as Fig. 3. The next bars display the implied mistakes for model realizations intermediate to the extremes of the plausible and composite ranges, while the outermost bars convey the implied mistakes for realizations at the extremes of the composite range.

The expanded information enhances the appeal of a precautionary response. Because the expanded model sets are right-skewed, the exercise makes policies that seem highly precautionary when attention is restricted to the plausible range models appear less severe. For example, maximin's exclusive emphasis on avoiding the environmental mistake falls away when more extreme possibilities are entertained. Of course, because the plots are intended only as a device to communicate trade-offs, individual decision-makers would naturally interpret the results through their own prior beliefs about the relative validity of different models. Those who deem models outside the plausible range as suspect may choose to give these outcomes only passing account, perhaps warranting further monitoring and research, but no weight in setting current policy targets. In contrast, those predisposed to consider these models more likely, may choose to evaluate policy alternatives giving these models at least some weight. We believe the trade-offs conveyed in this exercise deserve center stage in a discussion about precautionary alternatives, though the actual decision would ultimately hinge on a negotiated compromise among decision-makers with diverse interpretations of this information.

An alternative approach for incorporating preliminary perspectives within the proposed framework would be to formulate expert models as distinct probability distributions over climate sensitivity values, with uncertainty assumed to hold across these probabilistic models. This approach would have the merit of more-accurately accommodating expert views that put positive (if small) weight on low-probability, tail events; such views could be considered without extreme outcomes overwhelming the *ex ante* decisions. It would also allow for careful distinction between thin- and fat-tailed distributions, a distinction that has been given attention in recent work by Weitzman (2009). Unfortunately, computational limitations make implementing this approach quite difficult in DICE, and we leave this extension to future research.

## 5. Conclusions

There is pressing need for international decision-makers to anticipate and adequately respond to novel risks to the environ-

ment and public health. The precautionary principle is an important element of anticipatory policy where scientific knowledge is incomplete: i.e. where preliminary evidence precludes use of probabilistic methods for assessing risk. Yet the principle has been controversial, in part because of disagreement about how to implement it in the face of preliminary scientific understanding of a novel threat. One source of disagreement is the need for proportionality—the extent to which the *ex ante* consideration of costs should be admitted alongside the competing desire to achieve a high level of protection with respect to an identified threat. The paper develops practical analytical tools that frame the tension between precaution and proportionality in a way that makes sense without probabilities and that enables decision makers to identify and evaluate policies that balance these competing objectives in desired ways.

We show, in plausible circumstances, that a strictly precautionary response is accomplished by the decision criterion maximin. Such an approach may offer a reasonable, time-limited response in situations where the cost of prevention is low, where there is a plausible threat of severe and irreversible harm, and where learning is expected to proceed rapidly, but it is important to also consider less-conservative options. A contrasting approach is accomplished by the decision criterion minimax regret, which can be interpreted as implementing a strictly proportional response.

Framing the decision around these alternatives, we further consider an asymmetric version of minimax regret that flexibly varies the relative weight on the competing objectives of reducing society's exposure to the two mistakes. The criterion spans maximin and minimax regret, while also identifying intermediate alternatives. The intermediate alternatives implement a version of precaution that guards more than equally—though not exclusively—against false negative mistakes. To evaluate identified alternatives, we employ visual tools that convey the associated *ex ante* trade-offs. Taken together, the decision tools can be used to orient difficult intergovernmental decisions around trade-offs when scientific information is preliminary and contentious.

## Appendix A

The following sections provide a more rigorous foundation for the discussion in the text. Sections 2 and 3 build on the presentation in Section 1, so they should be read sequentially. The proofs mirror a number of results presented in the related paper Iverson (2011).

### A.1. Regret decomposition: separable case

First, we decompose regret into false-negative and false-positive mistakes under the assumption that welfare is additively separable in the welfare cost of abatement and the welfare cost of environmental damages.<sup>11</sup> In Appendix A.3, we provide an approximate decomposition to decompose regret in the DICE application.

Consider a set of expert models for a stock-pollutant abatement problem. We first present the case in which expert models are deterministic. We then allow for “probabilistic” expert models. Each model  $m \in M$  specifies an abatement cost function, an environmental damage function, and a mapping from sequences of abatement into sequences of environmental variables. The social objective is to minimize the present discounted sum of abatement costs and

<sup>11</sup> A possible interpretation of this separability condition is that the discount factor is a consumption discount factor used to discount the monetary value of future costs.

environmental damages, or equivalently, the objective is to maximize welfare defined to take the following “additively separable” form:

$$W(\{a_t\}, m) = \sum_{t=0}^T \beta^t [-C^{(m)}(a_t) - D^{(m)}(x_t)].$$

$C(m)(\bullet)$  is the welfare cost of abatement under model  $m$ ,  $D(m)(\bullet)$  is the welfare cost of environmental damages under model  $m$ , and  $\beta \in (0,1)$  is a constant discount factor. In addition, for a given path of (scalar) abatement levels  $\{a_t\}$ ,  $m$  defines the transition dynamics for a corresponding environmental state variable  $x$ :

$$x_{t+1} = G^{(m)}(x_t, a_t, t).$$

It is assumed that cost and damage functions are strictly increasing and strictly convex in their respective arguments. Moreover, it is assumed that the transition function  $G$  is strictly increasing in the level of the prior period stock variable and strictly decreasing in abatement. As a matter of notation, we let  $\{x_t^m\}$  denote the sequence of environmental variables that would result, given an arbitrary abatement sequence  $\{a_t\}$  and model  $m$ ; while  $\{x_t^{*,m}\}$  denotes the sequence of environmental variables that would result, given the optimal abatement path under  $m$ — $\{a_t^{*,m}\}$ —and model  $m$ .

It follows that

$$\begin{aligned} R(\{a_t\}, m) &= W(\{a_t^{*,m}\}, m) - W(\{a_t\}, m) \\ &= \sum_{t=0}^T \beta^t [-C^{(m)}(a_t^{*,m}, m) - D^{(m)}(x_t^{*,m})] - \sum_{t=0}^T \beta^t [-C^{(m)}(a_t) \\ &\quad - D^{(m)}(x_t^m)] \\ &= \sum_{t=0}^T \beta^t [-D^{(m)}(x_t^m) - D^{(m)}(x_t^{*,m})] - \sum_{t=0}^T \beta^t [-C^{(m)}(a_t) \\ &\quad - C^{(m)}(a_t^{*,m})] \equiv R_1(\{a_t\}, m) + R_2(\{a_t\}, m). \end{aligned}$$

Moreover,  $R_1(\{a_t\}, m) > 0$  implies

$$\sum_{t=0}^T \beta^t D^{(m)}(x_t^m) > \sum_{t=0}^T \beta^t D^{(m)}(x_t^{*,m})$$

which says the present discounted value of damages under  $\{a_t\}$  exceeds what would have been optimal—the false-negative mistake. Similarly,  $R_2(\{a_t\}, m)$  implies

$$\sum_{t=0}^T \beta^t C^{(m)}(x_t^m) > \sum_{t=0}^T \beta^t C^{(m)}(x_t^{*,m})$$

which says the present discounted value of abatement costs under  $\{a_t\}$  exceeds what would have been optimal—the false-positive mistake.

To further allow for expert models that accommodate risk, suppose that each expert  $\gamma \in \Gamma$  defines a probability distribution over a finite number of abatement cost functions  $C^{(m)}(\bullet)$ , a finite number of environmental damage functions  $D^{(m)}(\bullet)$ , and a finite number of transition equations  $G^{(m)}(\bullet)$ . This formulation allows for accumulation laws that are stochastic in the “static” sense of mapping abatement sequences into finite support probability distributions over sequences of state variables:

$$\gamma : \{a_t\} \rightarrow \Delta\{x_t\}$$

The structure of uncertainty is consistent with a situation in which incomplete knowledge—as opposed to randomness—comprises the dominant source of uncertainty and where there is no learning.

Let  $\gamma = (\gamma_1, \dots, \gamma_n)$  denote a probability distribution over models, and suppose that the social objective is to maximize expected welfare:

$$\begin{aligned} W(\{a_t\}, \gamma) &= E^\gamma \left\{ \sum_{t=0}^T \beta^t [-C^{(m)}(a_t) - D^{(m)}(x_t^m)] \right\} \\ &= \sum_{m_i \in M} \left\{ \sum_{t=0}^T \beta^t [-C^{(m_i)}(a_t) - D^{(m_i)}(x_t^{m_i})] \right\} \gamma_i. \end{aligned}$$

Let  $\{a_t^{*,\gamma}\}$  denote the optimal abatement sequence under expert model  $\gamma$ , and let  $\{x_t^{\gamma, m_i}\}$  denote the (deterministic) path of state variables that would ensue under model  $m$  given abatement sequence  $\{a_t^{*,\gamma}\}$ . Then

$$\begin{aligned} R(\{a_t\}, \gamma) &= W(\{a_t^{*,\gamma}\}, \gamma) - W(\{a_t\}, \gamma) \\ &= \sum_{m_i \in M} \left\{ \sum_{t=0}^T \beta^t [-C^{(m_i)}(a_t^{*,\gamma}) - D^{(m_i)}(x_t^{\gamma, m_i})] \right\} \gamma_i \\ &\quad + \sum_{m_i \in M} \left\{ \sum_{t=0}^T \beta^t [-C^{(m_i)}(a_t) - D^{(m_i)}(x_t^{m_i})] \right\} \gamma_i \\ &= \sum_{m_i \in M} \left\{ \sum_{t=0}^T \beta^t [D^{(m_i)}(x_t^{m_i}) - D^{(m_i)}(x_t^{\gamma, m_i})] \right\} \gamma_i \\ &= \sum_{m_i \in M} \left\{ \sum_{t=0}^T \beta^t [C^{(m_i)}(a_t) - C^{(m_i)}(a_t^{*,\gamma})] \right\} \gamma_i \equiv R_1(\{a_t\}, \gamma) \\ &\quad + R_2(\{a_t\}, \gamma). \end{aligned} \tag{1}$$

The third equality follows because all objects in curly brackets are deterministic for fixed  $m_i$ .

It follows that  $R_1(\{a_t\}, m) > 0$  implies  $E^\gamma \left[ \sum_{t=0}^T \beta^t D^{(m)}(x_t^m) \right] > E^\gamma \left[ \sum_{t=0}^T \beta^t D^{(m)}(x_t^{\gamma, m}) \right]$ , which says that realized damages exceed what they optimally would have been in an expected present discounted value sense. Similarly,  $R_2(\{a_t\}, m) > 0$  implies  $E^\gamma \left[ \sum_{t=0}^T \beta^t C^{(m)}(a_t) \right] > E^\gamma \left[ \sum_{t=0}^T \beta^t C^{(m)}(a_t^{*,\gamma}) \right]$

### A.2. Maximin “unilateral”

As discussed in the text, under specific conditions, the maximin criterion guards “unilaterally” against the environmental hazard in that it always, across the range of expert models, avoids the environmental mistake, while it also always makes the mistake (at least weakly) of wasting resources. The result is demonstrated numerically for the application. In particular, the maximin plot in Fig. 3 shows that positive realizations of the environmental mistake are always avoided under the maximin policy, while simultaneously, the growth mistake is always made. The following proposition establishes conditions under which this relationship can be expected to hold, and the subsequent corollary and discussion defines sufficient conditions on the model fundamentals for the proposition’s hypothesis to apply.

**Proposition.** *Suppose there exists an expert model  $\tilde{\gamma}$  that is “worst-case” in the sense that, for all  $\gamma \neq \tilde{\gamma}$  and all  $\{a_t\}$  feasible,*

$$W(\{a_t\}, \tilde{\gamma}) < W(\{a_t\}, \gamma) \tag{2}$$

and

$$\{a_t^{*\bar{\gamma}}\} > \{a_t^{*\gamma}\}, \tag{3}$$

where the latter means that for all  $t$ ,  $a_t^{*\bar{\gamma}} \geq a_t^{*\gamma}$ , while for some  $t$ ,  $a_t^{*\bar{\gamma}} > a_t^{*\gamma}$ . Then the maximin policy guards “unilaterally” against the environmental mistake in the sense that, for all  $\gamma \neq \bar{\gamma}$

$$R_1(\{a_t^{\text{maximin}}\}, \gamma) < 0$$

and

$$R_2(\{a_t^{\text{maximin}}\}, \bar{\gamma}) = 0$$

In addition,  $R_1(\{a_t^{\text{maximin}}\}, \bar{\gamma}) = R_2(\{a_t^{\text{maximin}}\}, \bar{\gamma}) = 0$ .

The conditions (2) and (3) will often go together, but they do not have to. It is easy to verify for a two-period version of the problem that sufficient conditions for (2) and (3) to hold jointly include either (a) that (all else equal) the pollutant accumulates faster in expectation under  $\bar{\gamma}$  in that, for all  $x, a, t$  and  $\gamma \neq \bar{\gamma}$ ,  $E^{\bar{\gamma}}[(d/dx)D^{\bar{\gamma}}(g^m(x, a, t))] > E^{\gamma}[(d/dx)D^{\gamma}(g^m(x, a, t))]$ , or (b) that (all else equal) marginal damages are higher in expectation: for all  $x$  and  $\gamma \neq \bar{\gamma}$ ,  $E^{\bar{\gamma}}[D_x^{\bar{\gamma}}(x)] > E^{\gamma}[D_x^{\gamma}(x)]$ . In contrast, the conditions would not hold if experts disagreed about abatement costs only. In this case, relatively high abatement costs would lower welfare for any given sequence of abatement, but high abatement cost would also lead to a lower optimal path of abatement, causing the two conditions to move in opposite directions. The following corollary extends these conditions for a deterministic, linear quadratic version of the problem where the time horizon is infinite. This problem admits an explicit analytical solution for the abatement path, so the relevant comparative static results are straightforward. Of course, the hypothesis of the proposition could also hold under many possible combinations of expert models mixing disagreements about damages, stock accumulation models, and abatement costs, and the hypothesis can be tested in a straightforward way within any particular model specification.

In addition, as we argue in the paper, severe uncertainty in decision problems regarding novel environmental threats is likely to be a more appropriate modeling assumption for describing uncertainty with respect to environmental models (including both the accumulation dynamics and the damage function) then it would be for modeling uncertainty about abatement costs. This is because precautionary decisions for novel threats typically involve preliminary scientific evidence with limited historical data and incomplete knowledge. In contrast, estimating costs, while no doubt uncertain, would typically lend themselves to adequate comparables in historical experience that could be used to estimate probabilities over risky possibilities. For this reason, we imagine the framework being used most often in a situation where policy-makers regard uncertainty with respect to abatement costs as being adequately described by risk. Translated into the expert panel problem, this means that different experts would share the same (probabilistic) model of abatement costs. In this case, conditions (2) and (3) from the proposition would be compatible.

**Corollary.** Consider a deterministic, linear-quadratic version of the abatement problem above. We formulate the problem in continuous-time and take the time horizon to be infinite. In particular, the social planner chooses the path of abatement to maximize

$$\int_0^{\infty} [-\alpha a(t)^2 - \gamma x(t)^2] e^{-\rho t} dt$$

subject to

$$\begin{aligned} \dot{x}(t) &= -\beta x(t) + \bar{E} - a(t), & 0 < \beta < 1, \bar{E} \geq 0, \\ x(0) &= x_0 \text{ given} \\ x(t) &\geq 0, \end{aligned}$$

where  $x$  is the stock of pollution,  $\beta$  is the natural rate of decay, and  $\bar{E}$  is the per-period flow of emissions along the business as usual path. Suppose moreover that the initial pollution level is below the long run steady state. Then if experts disagree about  $\gamma$  and  $\beta$  the worst-case model according to condition (2) would have the highest value of  $\gamma$  and the lowest value of  $\beta$ . In this case, conditions (2) and (3) would be satisfied jointly. In contrast, if experts disagree about the cost parameter  $\alpha$  (only), then the worst-case model according to condition (2) would have the highest value of  $\alpha$ . But the corresponding abatement sequence under this worst-case model would then be lower than for alternative models, causing condition (3) to be violated.

**Proof of proposition.** From the first part of the hypothesis, the maximin adversarial agent will always play  $\bar{\gamma}$  in the maximin zero-sum game, so  $\{a_t^{\text{maximin}}\} = \{a_t^{*\bar{\gamma}}\}$ . Fix  $\gamma \neq \bar{\gamma}$ . Then  $\{a_t^{*\bar{\gamma}}\} > \{a_t^{*\gamma}\}$  from the hypothesis, together with  $C' > 0$ , implies that for all  $m$ ,

$$\begin{aligned} R_2(\{a_t^{*\bar{\gamma}}\}, m) &= \sum \left(\frac{1}{1+\rho}\right)^t C(a_t^{*\bar{\gamma}}) \\ &\quad - \sum \left(\frac{1}{1+\rho}\right)^t C(a_t^{*m}) > \sum \left(\frac{1}{1+\rho}\right)^t C(a_t^{*\gamma}) \\ &\quad - \sum \left(\frac{1}{1+\rho}\right)^t C(a_t^{*m}) \\ &= R_1(\{a_t^{*\gamma}\}, m). \end{aligned}$$

Therefore,

$$\begin{aligned} R_2(\{a_t^{*\bar{\gamma}}\}, \gamma) &= \sum_i \gamma_i R_2(\{a_t^{*\bar{\gamma}}\}, m_i) > \sum_i \gamma_i R_2(\{a_t^{*\gamma}\}, m_i) \\ &= R_2(\{a_t^{*\gamma}\}, \gamma) = 0. \end{aligned}$$

The first equality follows from Eq. (1). The last equality follows because  $\{a_t^{*\gamma}\}$  is optimal given  $\gamma$ . Similarly, the assumption  $g_a < 0$  implies that  $\{x_t^{*\bar{\gamma}}\} < \{x_t^{*\gamma}\}$ . This follows from the second part of the hypothesis and can be proven by induction. Consequently,  $D' > 0$  implies that for all  $m$ ,  $R_1(\{a_t^{*\bar{\gamma}}\}, \gamma) = \sum_i \gamma_i R_1(\{a_t^{*\bar{\gamma}}\}, m_i) < \sum_i \gamma_i R_1(\{a_t^{*\gamma}\}, m_i) = R_1(\{a_t^{*\gamma}\}, \gamma) = 0$

Finally, when  $\gamma = \bar{\gamma}$ ,  $\{a_t^{*\bar{\gamma}}\}$  is optimal so  $R_1(\{a_t^{*\bar{\gamma}}\}, \bar{\gamma}) = R_2(\{a_t^{*\bar{\gamma}}\}, \bar{\gamma}) = 0$ .

**Proof of corollary.** Suppose experts disagree about  $\gamma$  and  $\beta$  and let the worst-case expert model correspond to the high  $\gamma$ , low  $\beta$  expert. Then it is straightforward to verify that for all abatement sequences, welfare is lowest for the worst case model. Similarly, if experts disagree about the cost parameter, then for an arbitrary abatement sequence, welfare is lowest when the cost parameter is highest. We next show that when the worst-case model is defined to have high  $\gamma$  and low  $\beta$ , then condition (3) will also be satisfied. In contrast, if the worst-case model is defined to have high  $\alpha$  then condition (3) will be violated. Both results follow from taking the derivative of the abatement sequence with respect to the relevant parameters at each point in time.

Using the maximum principle, one can show that the optimal path of the environmental state variable converges in the long run to the steady state

$$\hat{x} = \frac{(\beta + \rho)\bar{E}}{\beta^2 + (\gamma/\alpha) + \beta\rho}.$$

Moreover, the optimal path of  $x$  is given by

$$x(t) = \hat{x} + (x_0 - \hat{x})e^{-rt},$$

where  $r = -(\rho/2) + \sqrt{(\rho/2)^2 + \beta^2 + (\gamma/\alpha) + \beta\rho}$ .

So the state variable converges at the exponential rate  $r$  from  $x_0$  to  $\hat{x}$ . Meanwhile, the optimal path of abatement is

$$a(t) = -(r + \beta)(x_0 - \hat{x})e^{-rt} - \beta\hat{x} + \bar{E}.$$

It is then straightforward to verify that, for all  $x_0 \geq 0$  and all  $t \geq 1$ ,  $(d/d\gamma)a(t) > 0$  and  $(d/d\alpha)a(t) < 0$ . Moreover, if  $x_0 \leq \hat{x}$  then it is straightforward to verify that, for all  $t \geq 1$ ,  $(d/d\beta)a(t) < 0$ .

### A.3. Approximate decomposition, non-separable case

For many environmental policy models of interest, the welfare function does not take the separable form assumed above. This is the case, for example, in DICE. We employ an approximation that mimics the decomposition above, and that works well in the context of the DICE application in the paper.

Welfare in DICE can be written as a function of a sequence of abatement costs  $\{AC_t\}$  (measured in monetary units) and a sequence of environmental damages  $\{D_t\}$  (measured in monetary units):

$$W(\{a_t\}, m) = W(\{AC_t^m\}, \{D_t^m\}).$$

This is interesting because it implies one can mix and match. For example,  $W(\{AC_t^m\}, \{D_t^{*m}\})$  would denote the hypothetical welfare that would accrue if abatement were as under and damages (adjoined to a corresponding sequence of atmospheric temperatures) were as under  $\{a_t^{*m}\}$ .

The false-negative component of regret,  $R_1(\{a_t^{*m}\})$ , entails the extent to which the welfare cost of environmental damages exceeds what would have been optimal. This corresponds to the welfare difference that would result under the alternate damage sequences  $\{D_t^{*m}\}$  and  $\{D_t^m\}$ , holding *all else equal*—and in particular, holding equal the sequence of abatement costs. Consequently, two natural candidates for  $R_1$  are

$$r_{11} = W(\{AC_t^m\}, \{D_t^{*m}\}) - W(\{AC_t^m\}, \{D_t^m\}),$$

which holds abatement cost fixed at the “unstarred” level, and

$$r_{12} = W(\{AC_t^{*m}\}, \{D_t^{*m}\}) - W(\{AC_t^{*m}\}, \{D_t^m\}),$$

which holds abatement cost fixed at the “starred” level.

Similarly, two natural candidates for  $R_2$  are

$$r_{21} = W(\{AC_t^{*m}\}, \{D_t^m\}) - W(\{AC_t^m\}, \{D_t^m\}),$$

which holds damages fixed at the “unstarred” level, and

$$r_{22} = W(\{AC_t^{*m}\}, \{D_t^{*m}\}) - W(\{AC_t^m\}, \{D_t^{*m}\}),$$

which holds damages fixed at the “starred” level.

Of course, one could also hold the “all else equal” sequences fixed at intermediate values, so the approximate decomposition that we are about to suggest would be suspect if the candidates differed by a lot. In our example, they differ in magnitude by less than a quarter of a percent.

While the suggested candidates capture the intended meaning of the regret components, they do not in fact decompose regret. But if one defines the approximate components as a simple average of the candidates above, then the decomposition is exact. In particular, define

$$\tilde{R}_1(\{a_t\}, m) = \frac{1}{2}r_{11} + \frac{1}{2}r_{12}.$$

and

$$\tilde{R}_2(\{a_t\}, m) = \frac{1}{2}r_{21} + \frac{1}{2}r_{22}.$$

Then

$$\tilde{R}_1(\{a_t\}, m) + \tilde{R}_2(\{a_t\}, m) = R(\{a_t\}, m),$$

as confirmed by simple arithmetic. These approximations are used to generate the results presented in the paper.

### A.4. The decision problem with anticipated learning

The application considers an extension in which the decision maker anticipates learning the true model in some future period  $T_L$ . We describe the problem structure in more detail below. To simplify the exposition, we focus on the case in which expert models are deterministic, though the problem could be similarly specified for the case in which expert models include risk.

Working backwards, the continuation value starting in period  $T_L$  when the true model  $\gamma$  is known, is given by  $V(x; \gamma)$ , a function of the state variable in period  $T_L$  that equals the value of the following optimization problem:

$$\begin{aligned} & \max_{\{a_s\}_{s=T_L}^{T_L}} \sum_{s=T_L}^T \beta^s u(a_s, x_s, \gamma; t) \\ & \text{to} \\ & \text{subject} \\ & \quad x_{s+1} = g^\gamma(x_s, a_s) \\ & \quad x_{T_L} = x. \end{aligned}$$

The decision then is to choose a policy path for the run-up period before learning occurs. For a given abatement sequence, ex post welfare is given by

$$\begin{aligned} & \left\{ W(\{a_t\}_{t=0}^{T_L-1}, \gamma) = \max_{\{a_t\}_{t=0}^{T_L-1}} \sum_{t=0}^{T_L-1} \beta^t u(a_t, x_t, \gamma; t) + V(x_{T_L}; \gamma) \right\} \\ & \text{subject to} \\ & \quad x_{t+1} = g^\gamma(x_t, a_t) \\ & \quad x_{T_L} = x. \end{aligned}$$

With welfare so defined, the corresponding definitions for the considered non-Bayesian decision criteria apply as before.

## References

- Allen, M., et al., 1996. Observational constraints on climate sensitivity. In: Schellhuber, J.S., et al. (Eds.), *In Avoiding Dangerous Climate Change*. Cambridge University Press, pp. 281–289.
- Allen, M., Frame, D., 2007. Call off the quest. *Science* 318, 582–583.
- Arrow, K., Fisher, A., 1974. Environmental preservation, uncertainty and irreversibility. *Quarterly Journal of Economics* 88, 312–319.
- Arrow, K., Hurwicz, L., 1972. An optimality criterion for decision-making under ignorance. In: Carter, C., Ford, J. (Eds.), *Uncertainty and Expectations in Economics: Essays in Honour of G.L.S. Shackle*. Basil Blackwell, Oxford.
- Ashford, N., 2007. The legacy of the precautionary principle in US Law: the rise of cost-benefit analysis and risk assessment as undermining factors in health, safety and environmental protection. In: Sadeleer, N. (Ed.), *Implementing the Precautionary Principle: Approaches From the Nordic Countries, the EU and the USA*. Earthscan, London, pp. 352–378.
- Barrieu, P., Sinclair-Dégagné, B., 2006. On precautionary policies. *Management Science* 52, 1145–1154.
- Cooney, R., Dickson, B., 2005. Biodiversity and the Precautionary Principle: Risk and Uncertainty in Conservation and Sustainable Use. Earthscan, London.
- Epstein, L., 1980. Decision-making and temporal resolution of uncertainty. *International Economic Review* 21, 269–284.
- Gilbertson, M., 2001. The precautionary principle and early warnings of chemical contamination of the Great Lakes. In: Harremoës, P., Gee, D., MacGarvin, M., Stirling, A., Keys, J., Wynne, B., Guedes Vaz, S. (Eds.), *Late Lessons from Early Warnings: The Precautionary Principle*, Environment Issue Report No. 22. European Environment Agency, Copenhagen, pp. 126–134.
- Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18 (2), 141–153.
- Gollier, C., Bruno, J., Treich, N., 2000. Scientific progress and irreversibility: an economic interpretation of the precautionary principle. *Journal of Public Economics* 75, 229–253.
- Gollier, C., Treich, N., 2003. Decision-making under scientific uncertainty: the economics of the precautionary principle. *The Journal of Risk and Uncertainty* 27 (1), 77–103.
- Gollier, C., 2010. Ecological discounting. *Journal of Economic Theory* 145 (2), 812–829.
- Haigh, N., 1993. *The Precautionary Principle in British Environmental Policy*. Institute for European Environmental Policy, London.
- Hammit, J., Lempert, R., Schlesinger, M., 1992. A sequential-decision strategy for abating climate change. *Nature* 357, 315–318.
- Hansen, L.P., Sargent, T.J., 2008. *Robustness*. Princeton University Press, Princeton, NJ.
- Harremoës, P., et al. (Eds.), 2001. *Late lessons from early warning: the precautionary principle 1896–2000*. European Environment Agency, Copenhagen.
- Harsanyi, J., 1977. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge University Press.
- Heal, G., 2008. Climate economics: a meta review and some suggestions for future research. *Review of Environmental Economics and Policy* 3 (1), 4–28.
- Henry, C., 1974. Investment decisions under uncertainty: the irreversibility effect. *American Economic Review* 64, 1006–1012.
- IPCC, 2007. *Climate Change 2007: The Physical Science Basis*. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press.
- Iverson, T., 2011. *Communicating Trade-offs Amid Controversial Science*. Working Paper.
- Jones, K., Patel, N., Levy, M., Storeygard, A., Balk, D., Gittleman, J., Daszak, P., 2008. Global trends in emerging infectious diseases. *Nature* 451 (7181), 990–993.
- Kahneman, Tversky, A., 1979. Prospect theory an analysis of decision under risk. *Econometrica* 47 (2), 263–292.
- Keller, K., Bolker, B., Bradford, B., 2004. Uncertain climate thresholds and optimal economic growth. *Journal of Environmental Economics and Management* 48 (1), 723–741.
- Kelly, D., Kolstad, C., 1999. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control* 23, 491–518.
- Knight, F., 1921. *Risk, Uncertainty, and Profit*. Houghton Mifflin, Boston.
- Kolstad, C., 1996. Learning and stock effects in environmental regulation: the case of greenhouse gas emissions. *Journal of Environmental Economics and Management* 31 (1), 1–18.
- Lange, A., 2003. Climate change and the irreversibility effect – combining expected utility and maximin. *Environmental and Resource Economics* 25 (4), 417–434.
- Lemoine, D., Traeger, C., 2010. *Tipping Points and Ambiguity in the Integrated Assessment of Climate Change*. CUDARE Working Paper.
- Lempert, R., Collins, M., 2007. Managing the risk of uncertain threshold responses: comparison of robust, optimum, and precautionary approaches. *Risk Analysis* 37 (4), 1009–1026.
- Manning, M., 2006. The treatment of uncertainties in the fourth IPCC assessment report. *Advances in Climate Change Research* 2 (1), 13–21.
- McDaniels, T., Kamlet, M., Fischer, G., 1992. Risk perception and the value of safety. *Risk Analysis* 12 (4), 495–503.
- Milnor, J., 1954. *Games against nature*. In: Thrall, R., Coombs, C., Davis, R. (Eds.), *Decision Processes*. Wiley, New York.
- Morgan, M., Keith, D., 1995. Subjective judgements by climate experts. *Environmental Science & Technology* 29 (10), 468–476.
- Nordhaus, W., 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.
- National Research Council (NRC), 1979. *Carbon Dioxide and Climate: A Scientific Assessment (Report of an Ad Hoc Study group on Carbon Dioxide and Climate)*. Natl. Acad. Sci., Washington, DC, 22 pp.
- Perrings, C., Brock, W., 2009. Irreversibility in economics. *Annual Review of Resource Economics* 1, 219–238.
- Sadeleer, N. (Ed.), 2007. *Implementing the Precautionary Principle: Approaches from the Nordic Countries, the EU and the USA*. Earthscan, London.
- Stainforth, D., Aina, T., Christensen, C., Collins, M., Faull, N., Frame, D., Kettleborough, J., Knight, S., Martin, A., Murphy, J., 2005. Uncertainty in predictions of the climate response to rising levels of greenhouse gases. *Nature* 433 (7024), 403–406.
- Stern, N., 2006. *The Stern Review of the Economics of Climate Change*. Cambridge University Press, Cambridge, UK.
- Ulph, A., Ulph, D., 1997. Global warming, irreversibility and learning. *The Economic Journal* 107, 636–650.
- UNCED, 1993. *The Earth summit: the United Nations conference on environment and development (1992: Rio De Janeiro)*. In: Johnson Stanley, P. (Ed.), *Introduction and commentary*. Graham & Trotman, London.
- UNESCO, 2005. *Precautionary Principle World Commission on the Ethics of Scientific Knowledge and Technology*. United Nations Educational Scientific and Cultural Organization, France.
- US Food and Drug Administration, 2010. *Regulatory Procedures Manual*. US Food and Drug Administration, Silver Spring, MD.
- Vardas, G., Xepapadeas, A., 2010. Model uncertainty, ambiguity and the precautionary principle: implications for biodiversity management. *Environmental and Resource Economics* 45 (3), 379–404.
- Vogel, D., 2002. The politics of risk regulation in Europe and the United States. *Yearbook of European Environmental Law* 3, 1–42.
- von Neumann, J., Morgenstern, O., 1944. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.
- Walker, B., Barrett, S., Polasky, S., Galaz, V., Folke, C., Engström, G., Ackerman, F., Arrow, K., Carpenter, S., Chopra, K., Daily, G., Ehrlich, P., Hughes, T., Kautsky, N., Levin, S., Mäler, K., Shogren, J., Vincent, J., Xepapadeas, T., de Zeeuw, A., 2009. Looming global-scale failures and missing institutions. *Science* 325 (5946), 1345–1346.
- Webster, M., 2002. The curious role of learning in climate policy: should we wait for more data? *Energy Journal* 23 (2), 97–119.
- Weitzman, M., 2009. On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics* 91 (1), 1–19.
- Woodward, R., Bishop, R., 1997. How to decide when experts disagree: uncertainty-based choice rules in environmental policy. *Land Economics* 73 (4), 492–507.
- World Health Organization, 2005. *International Health Regulations*. World Health Organization, Geneva.
- World Trade Organization, 1995. *The WTO Agreement on the Application of Sanitary and Phytosanitary Measures (SPS Agreement)*. World Trade Organization, Geneva.