

The Precautionary Approach and the Role of Scientists in Environmental Decision-Making

Jan Sprenger*

January 9, 2011

1. Introduction

The role of scientists in advising public policy has recently been the subject of vivid discussion, particularly with respect to the value-ladenness of the scientists' epistemic analysis. The typical example for discussing this question is the problem of accepting vs. rejecting a particular hypothesis of interest (Rudner 1953; Levi 1960): setting thresholds for the probability of false positives and false negatives requires considerations about the impact and consequences of wrong decisions. An allegedly objective, impartial analysis is thus inevitably shaped a scientist's personal view on the significance of various kinds of damage (Douglas 2009, ch. 5). A similar case can be made for inferences in a Bayesian framework, where the choice of a Bayes estimator of a quantity of interest – the estimator that minimizes the expected risk, based on one's posterior distribution – does not only depend on that posterior distribution, but also on the context-dependent loss function that has been adopted. In that sense, scientists who advise a Government agency on the values of a quantity of interest or the truth of a certain hypothesis go beyond a purely epistemic analysis, and become implicit policy-makers themselves.

It might be surmised that once those epistemic decision problems have been solved, scientists have played their part in policy advice. A hypothesis has been established or rejected, a best estimate has been given, a probabilistic assessment of uncertainty has been made. For example, the Fourth Assessment Report of the IPCC (2007) summarizes its conclusions on global warming by qualifying various tenets that it investigated as “virtually certain” (>99% probability of occurrence), “extremely likely” (>95%), “very likely” (>90%), etc. Policy-makers are supposed to take it from there and to make their decisions by calibrating these probabilistic assessments with their subjective utility functions.

However, it has been neglected is that many cases of precautionary environmental management or risk analysis don't work like this. Here I don't refer to the problem of whether it makes any sense at all to make probabilistic forecasts for the extent of global warming (Frame et al. 2007; Stainforth et al. 2007).

*Contact information: Tilburg Center for Logic and Philosophy of Science, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands. Email: j.sprenger@uvt.nl. Webpage: www.laeuferpaar.de.

Rather, I mean that such probabilistic assessments are no sufficient basis for actual decision-making since the scope and extent of the underlying uncertainty is neglected. Supporting Douglas' (2009, ch. 7) claim that a neat separation of risk analysis into a scientific and a policy-related part (risk assessment and risk management, respectively) is impossible, I argue that a precautionary approach to risk analysis needs scientific expertise *at all stages of the analysis*, including the proper decision-making. The argument proceeds in two steps: First, I elaborate decision-theoretic implications of the Precautionary Principle (PP), a widely endorsed guideline for policy-making and legal decision-making in the face of environmental hazards. While this result may be considered to have some value in itself, I am particularly interested in its implications: I contend that adhering to the decision rules compatible with the PP is not feasible without thorough scientific understanding of the underlying environmental system, vindicating the indispensability of scientists in environmental policy-making.

2. The Precautionary Principle

The PP is, as witnessed by its inclusion in international treaties, declarations and legal systems, a major guide for decision-making when dealing with environmental hazards in the face of severe uncertainty. In its most prominent wording, the principle states that

In order to protect the environment, the precautionary approach shall be widely adopted by States according to their capabilities. Where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation. (UNCED 1993, cf. EU 2000)

These formulations straightforwardly apply to assessing environmental hazards, such as allowing for the exploitations of natural resources, building a pipeline through a natural reserve, constructing atomic power plants, etc. The PP can, more generally, be understood as the requirement to refrain from actions and policies that run a suspected risk of causing harm to the public or to the environment, even if science has not yet established that such an action would indeed be harmful (Sandin 1999).

It is important to recognize that the Precautionary Principle is not designed for situations where we *can* quantify the risk associated with a certain action, e.g. by assigning a probability of a harmful outcome: the scope is a scientifically plausible, but highly uncertain environmental hazard where we lack a sensible way to quantify the risk involved. Indeed, in environmental decision-making, theoretical understanding is very often scarce, models may be biased in various directions, and systems are complex to an extent that we can hardly assess and quantify hazards that might occur (Halpern et al. 2006; Tebaldi et al. 2005; Tebaldi and Knutti 2007).

An ostensibly natural interpretation of the PP focuses on actions with the *greatest probability* to cause fatal harm to the public. Possible explications along these lines are examined by Peterson (2006, cf. Hansson (1997)):

PP_α If one act is more likely to give rise to a fatal outcome than another, then the latter should be preferred to the former. (*de minimis rule*)

PP_β If one act is more likely to give rise to a fatal outcome than another, then the latter should be preferred to the former, given that both fatal outcomes are equally undesirable.

In his contribution, Peterson shows convincingly that any of these decision rules conflicts with uncontentious principles of rational choice. If we assume (1) that preferences are totally ordered, that is, they are complete, asymmetric and transitive, (2) that our decision rules must obey dominance reasoning, (3) that shifting probability from bad to good outcomes can only increase the desirability of an action, then the above interpretation of the PP is not an admissible decision rule.

Apart from these decision-theoretic concerns, such a reading of the PP is also unrealistically conservative. That is, it focuses exclusively on the possibility of disasters and makes effective action virtually impossible (Harris and Holm 1999). For some discoveries, like new medical drugs, it is impossible to rule out adverse effects a priori; nevertheless, the benefits of a successful innovation can sometimes be so great that it seems to be reasonable to take the risk. Indeed, the consequences of such a reading of the PP would be so absurd that one may charge the critics of basing their points on a very uncharitable interpretation.

Therefore, Peterson's conclusion that "no version of the PP can be reasonably applied to decisions that may lead to fatal outcomes" (Peterson 2006, 595) is premature. He does not explore other ways of making sense of the principle, apart from the extremely risk-averse decision rules PP_α and PP_β. But if that cannot be the solution, what guidance is the Precautionary Principle supposed to provide in the face of uncertainty?

Some indications are given by the report of the World Commission on the Ethics of Scientific Knowledge and Technology (COMEST) on how to interpret and to apply the Precautionary Principle. The potential harms have to be *morally* unacceptable in order for the PP to apply: "When human activities may lead to morally unacceptable harm that is scientifically plausible but uncertain, actions shall be taken to avoid or diminish that harm." (COMEST 2005, 14) Such morally unacceptable harms may include irreversible damage – this is often true of ecological, but less frequently of economic harms –, issues of intergenerational justice, or a dismal effect on third parties (e.g. repeated floodings in Bangladesh due to anthropogenic global warming), etc. Thus, if the PP requires that public policy include measures to avoid or diminish those morally unacceptable harms, it directly aligns with the ideal of sustainable development.

In the remainder, I am going to show how the COMEST thesis can be used to flesh out decision-theoretic implications of the PP, and how these implications necessitate the engagement of scientists in the actual decision-making.

3. Decision-theoretic implications of the Precautionary Principle

The different ways of understanding the PP in environmental contexts can be divided into four groups:

1. As suggesting a concrete decision rule, e.g. to avoid any action where we cannot rule out potentially disastrous consequences beyond reasonable doubt.
2. As specifying a particular type of decision-rule, or approach to decision-making that should be adopted.
3. As guiding the specification and setup of a decision problem.
4. As determining which values should be adopted, and which goals should be pursued in an environmental decision problem.

These understandings can be mutually compatible, that is, it is possible to adopt two or more of the suggested readings of the PP. The first interpretation is clearly the most concrete, demanding and controversial. It seems to suggest something like a *maximin rule*, namely to adopt the option that ensures, if nature is malevolent, the least evil. Another option would be a *de minimis* rule: unless the probability of serious harm can be bounded below a certain threshold, we should abstain from a certain action. But we have already seen that such an understanding makes the job very easy for opponents of the principle. Since the only decision rules that could possibly be identified with the Precautionary Principle are so extreme, it seems better to reject the first option and not to identify the PP with a particular decision rule.

This road is also taken by Steele (2006) who remains skeptical about concrete decision-theoretic implications of the PP. Noting that the COMEST reading of the PP naturally aligns with the ideal of sustainable development, she surmises instead that the PP is best understood as urging us to conserve a wide range of possible options to future generations, something that, e.g., unlimited greenhouse gas emission is unlikely to achieve. In other words, not all welfare losses are on a par with respect to their moral status, and the PP asks us to favor the morally least suspicious option. Steele conceives of the PP as a general guideline, not as a concrete decision rule.

It could also be required to expand the space of possible states of nature that could obtain as a consequence of our actions (Resnik 1987). This would be the third option: according to the PP, a well-specified decision problem should include unlikely (but perhaps plausible) high-impact scenarios in order to better assess and to control the risks of actions that we set out to take.

The question is whether this is all that we can extract from the principle. Remember that the PP is not only a guideline for policy-making – it has been implemented in legal systems and international treaties, and its scope, weight and validity have been the subject of international lawsuits (Trouwborst 2006). For instance, when the EU banned the importation of US beef that had been

treated with hormones, the US filed a complaint at the WTO. If we take the principle as merely providing recommendations for *shaping* a decision problem, rather than solving it, it could be seen as an empty shell with little content. The COMEST report does not think so either, by stating that

When ethical dimensions of inter- and intragenerational equity are at stake, the other decision principles [such as Expected Utility Theory, J.S.] fail to satisfactorily address these problems characteristics [...] Because the PP applies to these cases where serious adverse effects/surprises can occur with unknown probability, it is rational to follow a “better safe than sorry” strategy. (COMEST 2005)

The emphasis on decision-theoretic rationality in this prominent report on the PP suggests to explore the second option. While accepting the more general guidelines listed in 3. and 4., I would contend that we can nevertheless extract some decision-theoretic implications from the PP. As argued above, it would be misguided to bet one’s money on a concrete decision rule, but I think that the PP favors a particular *type* of decision rules. This is similar to the distinction between rules and principles in legal theory: a legal rule is either valid or invalid and cannot have exceptions, whereas a legal principle can be exemplified by different rules in different contexts. I think that decision-theoretically, the PP should be understood as a broad principle comprising an arsenal of different, but related decision rules.

It goes beyond the scope of this paper to provide a neat taxonomy of the rules that exemplify the PP, but I would like to contend that these rules should satisfy two properties: first, they should focus on *performance thresholds* rather than the search for an optimal outcome,¹ second, they should be *robust to error* in the underlying scientific model.

The first tenet is motivated from the idea that we have to avert morally *unacceptable* harms. Interpreting the PP along these lines accounts for the intuition that there are certain kinds of harm that have to be avoided at almost all costs. Examples are leaving our descendants with completely uncontrolled global warming, irreversible destruction of vital ecosystems, a rise of the sea level that would flood overpopulated Bangladesh, and so on. In other words, efforts have to be directed towards ensuring that a certain acceptance threshold for the consequences of an action be met, and this should be the *goal* of our decision-making. This line of reasoning is different from the optimizing perspective of Subjective Expected Utility Theory (Savage 1972) and can be seen as a concrete decision-theoretic implication of the PP, pertaining to the type of reasoning that we have to adopt in environmental decision-making. We will give a case study in the next section.

Now I would like to defend the second decision-theoretic implication of the PP: *robust* decision-making. This is, to my mind, derivative on the PP’s stated goals to avoid morally unacceptable outcomes. Especially when we are eager to

¹As the reader may have noticed, this distinction resembles the well-known optimizing/satisficing dichotomy, but that terminology is slightly misleading since a satisficing rule can be represented as an optimizing rule as well.

avoid disastrous outcomes, and when our knowledge of the underlying systems is feeble, we will be inclined to take measures to ensure ourselves against slight deviations of the model that do not satisfy our acceptance thresholds any more. Such an interpretation naturally applies to, e.g., the climatic change debate: While there are economic analyses showing that combatting climate change may be more costly than just adapting to it, e.g. Nordhaus (1991), such analyses are based on a lot of contentious assumptions with respect to quantifying the costs of climate change, and assessing the plausibility of catastrophic global warming. In this example, cutting down carbon emissions and restructuring our energy production has, compared to pure adaptation, the immense advantage of being *robust*: we may, in the short and middle run, considerably slow down the growth of global economy, but we immunize the well-being of future generations against the disastrous consequences of a trivial modeling error in our more optimistic analyses. Without prejudicing complex decision problems in favor of environment-friendly policies, I think it is fair to say that decision-theoretic robustness is their main advantage.

It could now be argued that robustness is something that we would like to pursue anyway, that it is not specific to environmental decision-making. Then it would stand rather orthogonal to the Precautionary Principle. While I agree with the first claim, I disagree with the second. This standpoint has a descriptive side – as a matter of fact, robustness is more important in environmental contexts than in other branches of science and policy-making. But it also has normative foundations, because I believe that robustness has *extra value* in environmental decision problems. First, since experience teaches that the knowledge basis for our decisions can change quickly, responsible decision-making involves some healthy skepticism about our current state of knowledge, even if it is informed by our best science. Robust decisions are more likely to sustain. Second, policy decisions are, especially if they involve international collaboration, difficult to overturn, given the lengthy and cost-intensive deliberations needed to achieve an agreement in the first place. For those reasons, robustness should play a central role in environmental policy decisions.

4. Precautionary Reasoning in Practice

Accepting that the PP favors decision rules that pursue the goal of robust satisficing, I return to the original question: the interplay of science and public policy in environmental decision-making. I am arguing that choosing an appropriate decision rule is so complex a matter, and so sensitive to the epistemic situation we are in, that scientific advice is indispensable. In this sense, scientists extend their policy advice from the epistemic questions discussed in the introduction to the proper decision problems.

Note that this result could not occur if we had interpreted the PP as a concrete decision rule, like maximin or de minimis. In that case, the conceptual explication of the PP would have gone far enough to leave the decision to policy-makers alone on the basis of the scientists' epistemic assessments. Scientists could have contented themselves with specifying the possible scenarios, assessing

Option/State Probability of State	Poaching 0.1	Loss of Habitat 0.3	Dem. Accidents 0.5	Disease 0.01
Translocation	0.3	0.1	0.05	0.1
New Reserve	0.3	0.1	0.05	0.1
Captive Breeding	0.9	0.2	0.01	0.4

Table 1: A utility matrix of the Sumatran rhino conservation problem, following the analysis of Maguire et al. (1987).

the plausibility of the outcomes and leaving the decisions to the policy-makers. Similarly, if the PP had no decision-theoretic implications at all, it would just affect our goals and guide the specification of a decision problem, but it would have given policy-makers all freedom to choose a particular decision rule. In our understanding, however, the principle asks the decision-makers to choose a decision rule that aligns with robust satisficing, and to properly accomplish this goal, they need to rely on the help of scientists.

A good illustration of this thesis is the problem of conserving the endangered species of the Sumatran rhino. The original analysis (Maguire et al. 1987) evaluated the available options by means of ranking them according to their expected utility, defined as a decreasing linear function of the probability of extinction in various states of the world. Those probabilities can either be obtained from stochastic simulations or by subjective judgment. A simplified summary of such an analysis is given in table 1.²

The act “captive breeding” has the highest expected utility and was, in fact, the one recommended by the scientists and implemented by the policy-makers, with disastrous consequences. This is not surprising since captive breeding is especially vulnerable to demographic accidents in the captived population. Underestimating the probability of that event can have far-reaching effects for the expected utility of captive breeding. In other words, a straightforward, unreflected expected utility analysis – and which policy-makers like to reflect on what they are doing? – neglects the manifold sources of error and bias that affect numbers such as those in table 1. Unfortunately, environmental management problems are often characterized by severe uncertainty, paving the way for similar dismal failures.

I am now contrasting an expected utility analysis to Yakov Ben-Haim’s *information-gap approach* (in short: info-gap) that aims at decision-making in the spirit of robust satisficing (Ben-Haim 2006). Notably, the robustness of info-gap differs from the standard conception of robustness in philosophy of science, as the validation of a certain conclusion by means of various models or experiments. Rather, a robust decision allows for some violation in the model/estimates the while still satisfying a set of performance criteria. Ben-Haim’s info-gap approach exemplifies this idea. Assume that we want to estimate an unknown function of the variable t by our estimate $\hat{u}(t)$. This estimate

²Table 1 is no representation of the actual results of Maguire et al. (1987), but Regan et al. (2005) use this table to compare an expected utility perspective to an info-gap perspective. I am following their analysis here.

may, or may not, substantially miss the true function. The core idea is to define a *info-gap model* $\mathcal{U}(\hat{u}, \alpha)$ of our estimate: namely the set of functions $u(t)$ such that the fractional deviation of \hat{u} from u (or another reasonable divergence measure, such as the L^2 norm) does not exceed a threshold α :

$$\mathcal{U}(\hat{u}, \alpha) = \left\{ u(t) : \left| \frac{u(t) - \hat{u}(t)}{\hat{u}(t)} \right| \leq \alpha \right\} \quad (1)$$

The threshold value α is here called the *horizon of uncertainty*: it reflects the quality of our initial estimate and determines which functions $u(t)$ are taken into account, dependent on their distance from the initial estimate. Any available option q in the original decision problem can now be assessed in terms of the minimal return it ensures, for a given horizon of uncertainty α . This assessment is conducted by giving the *robustness function* of a decision q : the maximal horizon of uncertainty for which q still yields an acceptable outcome, that is, a loss $L(q, u)$ below the critical value l_c .

$$\hat{\alpha}(q, l_c) = \max \left\{ \alpha : \left(\min_{u \in \mathcal{U}(\hat{u}, \alpha)} L(q, u) \right) \leq l_c \right\} \quad (2)$$

In other words, $\hat{\alpha}$ quantifies the *robustness* of a decision q regarding uncertainty in our estimate $\hat{u}(t)$, by indicating the highest horizon of uncertainty for which a satisfactory decision (as expressed by payoff function L) is always ensured. By focussing on achieving a result above the acceptance threshold, the info-gap approach naturally implements the satisficing perspective that we have extracted from the PP.

Dependent on the robustness that we demand, the expected gains and losses can differ. If we fully trust our model, we can just take the decision that is optimal for the case of zero uncertainty. However, the more uncertainty we have, the more do we have to consider cautious options that may be unnecessarily cautious if our actual estimate was indeed reliable, but that ensure an acceptable performance if our estimate was unluckily mistaken. So typically, there will be an adverse relationship between robustness and expected utility, because in most info-gap models, robustness comes at a price. The decision-maker is then invited to choose an option that best balances the need for robustness and the opportuneness of expected utility calculus.

In our example of the Sumatran rhino, the posthumous analysis of Regan et al. (2005) shows that such a perspective would lead to a different decision. When the probabilities and utilities in our scientific analysis are subjected to an info-gap analysis, we realize that captive breeding is only for a narrow horizon of uncertainty superior to the other options. See figure 1. As soon as the uncertainty attached to our probabilistic estimates increases, captive breeding becomes more vulnerable to error, and we may be better advised to establish a new reserve. This decision may then be justifiably called a more robust choice.

This approach has a variety of virtues. By means of the horizon of uncertainty, robustness considerations are directly built into the entire model. Info-gap may therefore be called a “truly precautionary approach to management

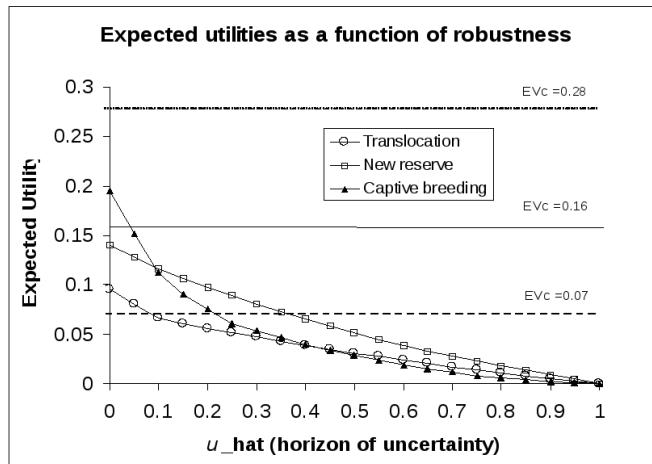


Figure 1: The minimal expected utility of the three management options in table 1, as a function of the horizon of uncertainty. Source: Regan et al. (2005).

and conservation” (Halpern et al. 2006). Moreover, it can be applied without using complicated model averaging procedures: we can directly plug-in our best scientific model (the most comprehensive one, the most successful one, the one most closely related to scientific theory, etc.) and evaluate the estimates based on that model. At least from a practical point of view, that is a great simplification.

On the other hand, this simplicity of info-gap is also a weakness. On various occasions (e.g., Sniedovich (2008)), the info-gap approach has been criticized for relying on an inferential procedure on a specific estimate, generated from a specific model. What justification do we have for basing our robustness analysis on a single model estimate? This model may be highly biased, as argued in the introduction. Much depends on whether our best estimate is reasonable at all. This presumption does not seem to do proper justice to a situation of radical uncertainty that we often encounter in practice. Similarly, the choice of an appropriate metric such as the fractional deviation in (1) is likely to affect the outcome. These non-trivial decisions are unfortunately often passed over in silence, even in standard literature on the topic (Ben-Haim 2006). It is also unclear whether it makes sense to apply the approach when we face non-linear dynamics, e.g. in climate science, where the regions of acceptance for any actions can be ill-behaved, and the robustness functions of the available options can be misleading.

Whether these objections can, in a concrete case, be answered satisfactorily is a question that only scientists can decide. They are best equipped to evaluate

the soundness of a peculiar decision-theoretic approach against the epistemic situation we are in, and the likely sources of error and bias. To strengthen this point further, we note that there are lots of conceptually different alternatives to achieve the goal of higher robustness. Whereas most approaches, such as bounding probability densities, or working with probability intervals rather than precise numbers, specify the amount of uncertainty in advance, info-gap applies a sort of reverse engineering by deliberately leaving open the amount of uncertainty that we could possibly encounter (Halpern et al. 2006). Scientists need to choose a conceptual framework on the basis of their expertise in and experience with the relevant environmental management problems.

5. Conclusions

The interplay of science and public policy is an issue that gets more and more relevant in the world we live in. Environmental policy-making is a case at hand that has attracted particular interest over the last years, and it is a field where the precautionary approach to decision-making is most pronounced. According to this paper, spelling out a precautionary approach in practice is as much a scientific task as it is a task for the proper policy-makers in a society, contradicting the view that scientific policy advice is restricted to making some epistemic claims and assessing the uncertainty attached to them.

This claim has been defended by means of extracting decision-theoretic implications from the Precautionary Principle. On the interpretation proposed in this paper, the PP asks us to switch from an optimizing to a robust satisficing perspective: a perspective where the robust satisfaction of certain performance thresholds is more important than striving for the optimal outcome. Due to the multitude of decision-theoretic approaches that pursue the goal of robust satisficing, and the need to calibrate the epistemic characteristics of a given problem with the chosen decision rule, the expertise of scientists is indispensable in the actual decision-making. Therefore I conclude that scientists have to play an active role in policy-making (rather than just “providing the numbers”) and that responsible, precautionary environmental decisions demand a close interplay of science and public policy at all stages of the analysis.

References

- Ben-Haim, Yakov (2006): *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*. Second edition. San Diego: Academic Press.
- Douglas, Heather (2009): *Science, Policy, and the Value-free Ideal*. Pittsburgh: University of Pittsburgh Press.
- European Commission Directorate-General for the Environment (2000): *White Paper on Environmental Liability*. Luxemburg: European Commission.
- Frame, D.J., N.E. Faull, M.M. Moshi, and M.R. Allen (2007): “Probabilistic climate forecasts and inductive problems”, *Philosophical Transactions of the Royal Society A* 365, 1971–1992.

- Halpern, Benjamin S., Regan, Helen M., Possingham, Hugh P., and Michael A. McCarthy (2006): “Accounting for uncertainty in marine reserve design”, *Ecology Letters* 9, 2-11.
- Hansson, Sven Ove (1997): “The Limits of Precaution”, *Foundations of Science* 2, 293–306.
- Holm, Søren, and John Harris (1999): “Precautionary Principle Stifles Discovery”, *Nature* 400, 398.
- The Intergovernmental Panel on Climate Change (2007): *4th Assessment Report*. Retrieved from www.ipcc.ch on the 6th of January 2010.
- Maguire, L.A., U.S. Seal, and P.F. Brussard (1987): “Managing critically endangered species: a case study of the Sumatran rhino”, in: M. Soulé (ed.), *Viable populations for conservation*, 141–158. Cambridge: Cambridge University Press.
- Nordhaus, William (1991): “To Slow or Not to Slow: The Economics of the Greenhouse Effect”, *The Economic Journal* 101, 920–37.
- Levi, Isaac (1960): “Must the Scientist Make Value Judgments?”, *Journal of Philosophy* 57, 345–357.
- Peterson, Martin (2006): “The Precautionary Principle Is Incoherent”, *Risk Analysis* 26, 595–601.
- Regan, H.M., Y. Ben-Haim, B. Langford, W.G. Wilson, P. Lundberg, S.J. Andelman, and M.A. Burgman (2005): “Robust decision making under severe uncertainty for conservation management”. *Ecological Applications* 15: 1471–1477.
- Resnik, Michael D. (1987): *Choices: An Introduction to Decision Theory*. Minneapolis: University of Minnesota Press.
- Rudner, Richard (1953): “The Scientist qua Scientist Makes Value Judgments”, *Philosophy of Science* 20, 1–6.
- Sandin, Per (1999): “Dimensions of the Precautionary Principle”, *Human and Ecological Risk Assessment* 5, 889–907.
- Savage, Leonard J. (1972): *The Foundations of Statistics*. New York: Wiley and Sons.
- Sniedovich, Moshe (2008): “Wald’s Maximin model: a treasure in disguise!”, *Journal of Risk Finance* 9: 287–291.
- Stainforth, D. A., M. R. Allen, E. R. Tredger and L. A. Smith (2007): “Confidence, uncertainty and decision-support relevance in climate predictions”, *Philosophical Transactions of the Royal Society A* 365, 2145–2161.

- Steele, Katie (2006): “The precautionary principle: a new approach to public decision-making?”, *Law, Probability and Risk* 5, 19–31.
- Tebaldi, Claudia, Richard L. Smith, Doug Nychka, and Linda O. Mearns (2005): “Quantifying Uncertainty in Projections of Regional Climate Change: A Bayesian Approach to the Analysis of Multimodel Ensembles”, *Journal of Climate* 18, 1524–1540.
- Tebaldi, Claudia, and Reto Knutti (2007): “The use of the multi-model ensemble in probabilistic climate projections”, *Philosophical Transactions of the Royal Society A* 365, 2053–2075.
- Trouwborst, Arie (2006): *Precautionary Rights and Duties of States*. Ph.D. Thesis: University of Utrecht.
- United Nations (1993): *The Earth Summit: The United Nations Conference on Environment and Development*. (Rio de Janeiro 1992). London: Graham & Trotman.
- World Commission on the Ethics of Scientific Knowledge and Technology (2005): *The Precautionary Principle*. Paris: UNESCO.