CHAPTER 1

The precaution controversy

> In order to protect the environment, the precautionary approach shall be widely applied 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.
>
> *Principle 15 of the 1992 Rio Declaration on Environment and Development*[1]

> 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.
>
> *Wingspread Statement on the Precautionary Principle*[2]

> The precautionary principle may well be the most innovative, pervasive, and significant new concept in environmental policy over the past quarter century. It may also be the most reckless, arbitrary, and ill-advised.
>
> Gary Marchant and Kenneth Mossman[3]

1.1 Introduction

Since the 1980s, the precautionary principle (PP) has become an increasingly prevalent fixture of international environmental agreements, from chlorofluorocarbons to biodiversity to climate change.[4] But despite – or perhaps because of – its prominence, PP is also extremely controversial. While PP is the subject of a massive academic literature, it remains notoriously difficult to define and responses to basic objections remain unclear.

[1] See the United Nations Framework Convention on Climate Change (article 3.3) for a very similar statement.
[2] See Raffensberger and Tickner (1999, pp. 353–4).
[3] See Marchant and Mossman (2004, p. 1).
[4] See (Raffensberger and Tickner 1999; Fischer, Jones, and von Schomberg 2006; Foster 2011; Trouwborst 2006; Whiteside 2006).

This book presents and defends an interpretation of PP from the perspective of philosophy of science. As a philosopher, I am concerned with basic problems of interpretation, logical coherence, and rationale. And as a philosopher of science, I am particularly concerned with how PP connects to scientific fields, such as climate science and toxicology, whose research comes in close contact with controversial environmental issues. Although PP has been approached from a variety of perspectives, I believe that philosophy generally and philosophy of science specifically has something valuable and important to offer. Objections to PP typically boil down to issues of a deeply philosophical nature, a number of which pertain to basic questions concerning values and scientific research. Can PP be formulated so it is both a distinctive and rational approach to environmental policy issues? What is its basic rationale in comparison to other approaches? How does it differ from approaches with which it is often contrasted, such as cost–benefit analysis? How does PP interact with policy relevant science? And can it do so in a way that does not threaten the integrity and reliability of scientific research? This book develops an interpretation of PP that aims to provide more adequate answers to these questions.

Much valuable work has already been devoted to achieving a better understanding of PP. A number of authors have examined the logical structure of the principle (Manson 2002; Sandin 1999), dissecting it into conditions concerning harm, knowledge, and a proposed remedy. Various perspectives or possible types of interpretation of PP have been distinguished (Ahteensuu and Sandin 2012; Sandin 2006). For instance, PP might be construed as a meta-rule that imposes general constraints on how decisions about environmental policy are made, as a decision rule that selects among concrete policy options, or as an epistemic rule requiring that a high standard of evidence be satisfied before a new technology is accepted as safe. Others have explored the relationship between PP and related concepts, such as the maximin rule (Ackerman 2008a; Gardiner 2006, 2010, 2011; Hansson 1997), robust adaptive planning (Doyen and Pereau 2009; Johnson 2012; Mitchell 2009; Sprenger 2012), and alternatives assessment (O'Brien 2000). In addition, much work has been done on explicating the role of PP in international law (Foster 2011; Trouwborst 2006).

Nevertheless, I believe that some fundamental challenges have not been adequately addressed in previous literature on PP. I discuss three of these now to provide a foretaste of the line of argument that this book will pursue: (1) the lack of an adequate response to the objection that, depending on

how it is interpreted, PP is either vacuous or irrational; (2) absence of an explanation of how the many ideas and perspectives associated with PP fit together as a coherent whole; and (3) an adequate account of the relation between PP and policy-relevant scientific research. Let us consider these issues in turn.

Critics of PP often charge that it can be given either a weak interpretation – according to which uncertainty does not justify inaction in the face of serious threats – or a strong interpretation – according to which precaution is required in the face of any scientifically plausible and serious environmental hazard. Weak interpretations are said to be true but trivial, since no reasonable person would demand complete certainty as a requisite for taking precautions. On the other hand, strong interpretations are claimed to be incoherent, and hence irrational, because environmental regulations themselves come with some risk of harmful effects and hence PP often precludes the very steps it recommends. I refer to this argument against PP as the dilemma objection.

Most responses to the dilemma objection focus on the second horn. One common reply is to propose that PP should qualified by a *de minimis* condition, which specifies a fixed evidential threshold that must be crossed before the principle is triggered (see Peterson 2002; Sandin 2005; Sandin *et al.* 2002, pp. 291–2). However, this reply is inadequate, because incoherence still arises whenever the harmful effects of the precaution themselves attain the evidential standard set out in the *de minimis* condition. In such circumstances, PP would recommend both for and against the precaution, the *de minimis* condition notwithstanding.

A number of other responses to the incoherence horn of the dilemma objection can be found in the literature. For example, Per Sandin (2006, pp. 179–80) suggests that applications of PP must be understood in relation to a context in which a particular type of danger is salient. To illustrate this idea, Sandin considers the practice of prescribing antibiotics as a precautionary measure for patients undergoing surgery. In this context, prescribing the antibiotic is precautionary only with respect to a possible infection, and not with respect to possible harms of excessive antibiotic use such as the evolution of resistant strains of bacteria. Thus, Sandin argues, PP does not generate contradictory recommendations if it is understood in its proper context. However, this is a problematic response, because there is no discernible justification for making a decision in a context wherein plausible and significant harmful effects of a proposed action are disregarded. In Sandin's example, it is entirely reasonable to insist that antibiotic resistance be relevant to decisions concerning best practices

for prescribing antibiotics. To claim that PP is applied in contexts in which negative effects of the recommended precaution are "off screen" seems only to reinforce the position of critics who assert that the incoherence of PP is often overlooked because of the regrettable human tendency to fixate on a single threat at a time (see Sunstein 2001, chapter 2; Sunstein 2005, chapters 2, 3, 4).

Sandin's "context" response might be read as suggesting that the incoherence objection is only an example of the general problem of local versus global framing of decisions (see Sandin *et al.* 2002, p. 293). For example, the decision above could be framed locally as whether to prescribe antibiotics as a prophylactic *for this particular patient* rather than more globally as whether such prescriptions should be made *generally* to patients in similar circumstances, and a decision rule might lead to opposite results in the two cases. If this were all the incoherence objection amounted to, then defenders of PP could easily reply that the problem is not specific to PP but confronts all decision rules. However, there is no need to interpret the incoherence objection in this manner. The objection is most naturally construed as charging that PP can lead to incoherent results *within a single framing of a decision problem*. For instance, critics would assert that PP leads to incoherent results when applied to questions about best practices of antibiotic prescription, recommending both for and against.

Another line of response to the incoherence horn of the dilemma turns on a deontological distinction between positive and negative duties (John 2007, p. 222; Weckert and Moor 2006, p. 199). A positive duty is an obligation to do good things, while a negative duty is an obligation to refrain from doing harm. It is commonly thought that negative duties are weightier than positive duties: for instance, that it is worse to murder a person than to fail to rescue her. However, this response is also problematic because PP is primarily intended to justify regulations, for example, that restrict the use of a toxic chemical. But regulations cannot be justified by appeal to negative duties, because enacting a regulation is not an omission but instead an action implemented by a government agency (see Munthe 2011, p. 71). The obligation to avoid harms resulting from regulations, then, would be a negative duty, and the argument in favor of a regulation would be grounded in a positive duty of the government to protect the environment or public health.

Another response is that the incoherence horn of the dilemma fails due to not noticing the role of proportionality in applications of PP

(see Fischer, Jones, and von Schomberg 2006; Whiteside 2006). Proportionality recommends that "precautionary responses ought to correspond to the perceived dimensions of the risks involved" (Trouwborst 2006, p. 150). Hence, "ban it!" is far from the only policy option available to PP, as some versions of the second horn of the dilemma seem to presume. I think that the concept of proportionality is in fact the key to turning back the second horn of the dilemma. However, this response is not adequate without some further elaboration. Granted, outright bans are not the only type of policy PP can recommend. But that alone is not an adequate answer, because the second horn of the dilemma does not require the assumption that precautions always take the form of absolute prohibitions. Instead, it turns on the possibility that the proposed precaution has potentially harmful effects that would be sufficient to trigger an application of PP, which would in turn recommend that the precaution itself be avoided or substantially restricted. Whether or not proportionality effectively addresses this issue is unclear given the rather vague terms in which it has been formulated. So the concept stands in need of further development if it is to serve as an adequate answer to the charge of incoherence.

I turn now to a second major challenge confronting interpretations of PP. This challenge is the great multiplicity of sometimes apparently conflicting ideas associated with it. To take just one example, consider the maximin rule, which several authors have suggested as a basis for interpreting PP (Ackerman 2008a; Gardiner 2006; Hansson 1997). The maximin rule recommends that one select the policy option that has the least bad worst-case outcome. However, other authors have suggested that PP should be understood in relation to the concept of minimax regret, according to which one should choose the action that minimizes the maximum shortfall from the best that could have been achieved (Chisholm and Clarke 1993). Yet minimax regret and the maximin rule can easily lead to conflicting results (see Hansson 1997). If PP is construed as a principle that aims to generate useful policy guidance, this simply will not do. It cannot be identical to a pair of contradictory principles. Of course, the same point holds for any other conflicting ideas associated with PP.

Some advocates of PP do not regard the absence of a unified account of PP as problematic. For instance, Lauren Hartzell-Nichols (2012, p. 160; 2013) proposes that, rather than one precautionary principle, there are many, each designed for a distinct set of circumstances. However, I do not think this is a stable position. For what makes all of these different

precautionary principles instances of the same general type? On the one hand, if a substantive answer can be provided to this question, then it seems that some unification of PP is afoot after all. On the other, if no substantive answer can be provided, then "PP" would be little more than an empty label that can be applied to almost anything one likes. But such a situation would render calls to adhere to PP found in international environmental agreements vacuous. Another disunified approach suggests that, despite its name, PP is in fact not a principle at all but rather a "repository" in which to deposit "adventurous" ideas that challenge conventional approaches to environmental policy (Jordan and O'Riordan 1999, p. 16). But such an approach faces obvious difficulties. For how do we decide which ideas may be dropped into the precautionary grab bag? And what should we do when those ideas conflict with one another? Answering such questions would require articulating some general conception of what PP does and does not assert, which the "repository" approach to PP explicitly disavows. But without answering such questions, the "repository" approach merely lends support to critics who charge that PP is no more than empty rhetoric masquerading as a serious approach to environmental issues.[5]

Some advocates of PP have attempted general interpretations that encompass a wide range of approaches. Arie Trouwborst (2006) attempts to distill the central elements of PP on the basis of an extensive review of formulations of the principle found in international law. This effort is, I think, extremely valuable insofar as providing a sense of what is generally meant by PP in a wide range of international agreements on environmental issues. As such, it imposes some general constraints on what a philosophical interpretation of PP should look like. In Trouwborst's account, international environmental law treats PP as a genuine principle – contrary to the "repository" approach described in the previous paragraph – involving several components, such as proportionality and the "tripod" of a knowledge condition, harm condition, and recommended precaution. But these general outlines are not sufficiently specific to resolve either concerns about the multitude of potentially conflicting ideas associated with PP or the dilemma objection. For instance, questions about the relationship of maximin and minimax regret in regard to PP are not answered, and while Trouwborst describes the incoherence horn of the dilemma objection, no answer to it is proposed (2006, pp. 184–7).

[5] Indeed, Marchant and Mossman (2004) cite Jordan and O'Riordan (1999) to support just such claims.

Another notable effort to articulate a unified perspective on PP is due to Stephen Gardiner (2006, 2010, 2011), who proposes a restricted version of the maximin rule as a "core" of PP. This rule recommends precaution when four conditions are satisfied: (1) nothing is known about the probabilities of the possible outcomes, (2) the precaution assures that catastrophe will be avoided, (3) the costs of enacting the precaution are minimal, and (4) any alternative action to the precaution may result in catastrophe. The thought is that, by working outwards from this core, it may be possible to attain a general understanding of PP. While I find much to admire in Gardiner's discussion, I think that his strategy for attempting to achieve a unified conception of PP is an unpromising one. The difficulty has to do with the nature of the conditions to which maximin is restricted. These conditions are such as to make the decision relatively easy. Hence, a number of decision rules that often conflict with maximin in other circumstances agree with it in the special case Gardiner examines. As a result, the "core" case provides very little indication of which direction to go when those restrictive conditions are relaxed. Should one continue to follow the maximin rule, or some other principle, and on what basis should such decisions be made? Gardiner's approach is also problematic if construed as an answer to the dilemma objection. In particular, critics assert that it is obvious that a precaution should be enacted if it is assured of preventing a potential catastrophe at practically no cost but claim that this observation is unhelpful since real issues of environmental policy generally involve hard trade-offs (see Sunstein 2005, p. 112). Thus, although Gardiner's restricted maximin interpretation of PP avoids incoherence, it is arguably skewered on the horn of triviality.[6]

Finally, let us turn to the third challenge concerning the relationship between PP and policy-relevant science. Although PP is most commonly discussed as a decision rule, it is not unusual for advocates to propose that it also has methodological implications for policy-relevant scientific research, sometimes under the banner of "precautionary science" (Barrett and Raffensperger 1999; Kriebel *et al.* 2001; Sachs 2011; Tickner and Kriebel 2006). One such implication of PP is the rejection of the ideal of value-free science. According to the value-free ideal, scientific research should be kept as separate as possible from ethical and political value judgments that inevitably influence policy decisions on environmental and human health issues (Douglas 2009; Lacey 1999; Proctor 1991). An epistemic PP conflicts with the value-free ideal by suggesting that the aims of protecting human

[6] See section 3.3.2 for a more detailed discussion of Gardiner's proposal.

health and the environment can legitimately influence methodological decisions in policy-relevant science. For example, it suggests that what should count as sufficient evidence that a new technology does not pose undue risks reflects a value judgment concerning the relative costs of unnecessary regulation versus harmful environmental or human health impacts. The argument from inductive risk is one classic and influential critique of the value-free ideal that is motivated by value judgments such as these (see Braithwaite 1953; Churchman 1948; Cranor 1993; Douglas 2009; Hempel 1965; Lemons, Shrader-Frechette, and Cranor 1997; Nagel 1961; Rudner 1953; Shrader-Frechette 1991; Steel 2010). According to this argument, the decision to accept a hypothesis involves a value judgment about what should count as sufficient evidence, a judgment that may depend on ethical considerations about the seriousness of distinct types of error. Versions of the argument from inductive risk are often encountered in discussions of the epistemic implications of PP (John 2007, 223; Kriebel *et al.* 2001, pp. 873–4; Peterson 2007, pp. 7–8; Sachs 2011, pp. 1302–3; Sandin *et al.* 2002, pp. 294–5).[7] However, these proponents of epistemic precaution do not engage with the philosophical literature criticizing the argument (Dorato 2004; Jeffrey 1956; Lacey 1999, 2004; Levi 1960, 1962, 1967; McMullin 1982; Mitchell 2004). The most common objection is that the argument from inductive risk relies on an outmoded behaviorist conception of acceptance, according to which to accept a hypothesis is to undertake some act that would be appropriate if the hypothesis were true. A defense of an epistemic PP, then, requires answering such charges. In addition, rejecting the value-free ideal requires proposing some alternative standard for distinguishing between legitimate and illegitimate influences of values in scientific research. An emerging literature on this topic exists in philosophy of science (see Douglas 2000, 2009; Elliott 2011a, 2011b, 2013; Elliott and McKaughan 2014; Kitcher 2001, 2011; Kourany 2010; Longino 2002; Steel 2010; Steel and Whyte 2012) but again has mostly been neglected in discussions of epistemic aspects of PP.

So I claim that the three challenges described above – the dilemma objection, the multitude of potentially conflicting ideas associated with PP, and the relation between PP and policy-relevant science – remain live concerns. This book is written with the firm conviction that they are not independent. Answering these challenges requires carefully examining how the several elements of PP interconnect with one another.

[7] In fact, Sandin *et al.* (2002, pp. 294–5) quote Rudner's classic (1953) statement of the argument.

In the next section, I sketch the outlines of how I believe this can be done.

1.2 Putting the pieces together

Questions about the interpretation of PP are an example of a classic type of philosophical puzzle. One is faced with an important, interesting, and yet tantalizingly unclear concept, and the problem is to provide a clear and coherent account of that concept, its rationale, and its logical implications. In the abstract, this is just the sort of puzzle that arises repeatedly in the Platonic dialogues with such questions as "What is justice?" or "What is love?" And as with any kind of puzzle, a solution depends on two factors: having the right pieces and putting them together in the right way. So what are the pieces to this puzzle, and how should they be assembled?

Paradigm applications of PP involve a trade-off between short-term gain, often for an influential party, against a harm that is uncertain or spatially or temporally distant. I propose that PP recommends the following three "core themes" for such decisions:

1. **The Meta-Precautionary Principle (MPP)**: The MPP asserts that uncertainty should not be a reason for inaction in the face of serious environmental threats. This principle is called "meta" because it is not a rule that indicates which of several environmental policy options to select – for instance, whether to set the allowable level of arsenic in drinking water at 50, 10, or 5 parts per billion. Instead, it places a restriction on what sorts of rules should be used for that purpose, namely decision rules that are susceptible to paralysis by scientific uncertainty should be avoided. In my interpretation, MPP is the most fundamental piece of PP insofar as it imposes constraints on the operations of the other two elements.
2. **The "Tripod"**: The term "tripod" refers to the knowledge condition, harm condition, and recommended precaution involved in any application of PP (see Trouwborst 2006). Like several other authors (Manson 2002; Munthe 2011; Sandin 1999), I take the elements of the tripod to be adjustable rather than fixed.[8] This means that there are

[8] Sandin (1999) adds a fourth element, namely how strongly the precaution is recommended – for instance, whether it is mandatory or merely permissible. In my approach, degrees of obligation have to do with the extent to which, in a particular context, adherence to PP is required. In other words, I think it is helpful to distinguish between (a) what PP recommends in a context and (b) whether one should act in accordance with PP in that context. Clearly, answering (a) is prerequisite for answering (b).

multiple ways of specifying the knowledge condition, harm condition, and recommended precaution, and careful consideration of the particulars of each application are relevant to deciding how to fill in the blanks. I will use the expression *version of PP* to refer to any statement according to which satisfying a specific knowledge and harm condition is sufficient to justify a specific precaution. For example, "If it is possible that an activity might lead to irreversible harm, then that activity should be banned" is one version of PP, while "If there is some scientific evidence that an activity will lead to irreversible harm, then an alternative should be substituted for that activity if feasible" is another. Which version of PP should be used in a given application is influenced by MPP – we should avoid versions that turn scientific uncertainty into paralysis – as well as by the next and final component.

3. **Proportionality**: Roughly, proportionality is the idea that the aggressiveness of the precaution should correspond to the plausibility and severity of the threat. I propose that proportionality be defined more precisely in terms of two subsidiary principles that I call consistency and efficiency. Consistency requires that the precaution not be recommended against by the same version of PP that was used to justify it. Efficiency states that, among those precautions that can be consistently recommended, the less costly should be preferred. Consistency and efficiency place important constraints on what can be justified by PP in a given context. For example, if there is no version of PP that can consistently recommend an action (say, preemptive war) in a given context (say, a US invasion of Iraq in 2003), then PP cannot justify that action in those circumstances. Finally, MPP affects how proportionality is applied. For example, comparisons of the relative efficiency of policy options should not be done in a way that makes scientific uncertainty grounds for continual delay.

Although these three core themes include many familiar elements of PP, as they should, the resulting proposal is distinctive in several important respects.

The most fundamental distinctive feature of the interpretation proposed here is the extent to which it ties together aspects of PP that are usually treated as separate or even conflicting. To explain this point more fully, it will be helpful to explicitly consider three possible ways of characterizing the role of PP (see Ahteensuu and Sandin 2012, pp. 971–2). One could view PP as either a

- *procedural requirement* that places some general constraints how decisions should be made; or a