University of Calgary

PRISM Repository **https://prism.ucalgary.ca**

The Vault Open Theses and Dissertations

2018-04-30

The Precautionary Principle: Concepts, Tools, and Strategies for Analysis

Leyenaar, James Andrew

Leyenaar, J. A. (2018). The Precautionary Principle: Concepts, Tools, and Strategies for Analysis (Master's thesis, University of Calgary, Calgary, Canada). Retrieved from https://prism.ucalgary.ca. doi:10.11575/PRISM/31885

http://hdl.handle.net/1880/106599

Downloaded from PRISM Repository, University of Calgary

UNIVERSITY OF CALGARY

The Precautionary Principle: Concepts, Tools, and Strategies for Analysis

by

James Andrew Leyenaar

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE

DEGREE OF MASTER OF ARTS

GRADUATE PROGRAM IN PHILOSOPHY

CALGARY, ALBERTA

APRIL, 2018

© James Andrew Leyenaar 2018

Abstract

In this thesis, I assess the Precautionary Principle (PP) as a guide to decision making in environmental policy. The PP says, roughly, that decision makers should act in advance of scientific certainty to protect against harms to the environment and public health. I argue that the main objections to the PP in the literature can be defended against and, further, that there is a need for such a principle due to the shortcomings of popular quantitative approaches to decision making. I also examine Stephen Gardiner's influential interpretation of the PP as a version of the maximin decision rule. I argue that Gardiner's interpretation makes the PP inapplicable in the epistemic situation of ignorance, a situation which most accurately characterizes the nature of our knowledge in many important environmental policy problems. I conclude that the PP should not be interpreted as a version of the maximin rule.

Acknowledgements

I would like to thank my thesis supervisor, Allen Habib, for guiding me through the development of the project, and for the many enjoyable conversations along the way. I thank Denise Retzlaff for her many words of encouragement. I also thank my family and friends for their generous support.

Table of Contents

Abstract ii
Acknowledgements iii
Table of Contents iv
List of Tables vi
List of Abbreviations vii

Chapter 1 The Precaution Debate 1
1.1 Overview of the Precautionary Principle 1
1.2 Problems facing the Precautionary Principle 2
1.3 Goals and methodology 3
1.4 Weak and strong versions 5

Chapter 2 Decision Theory Background 9
2.1 Introduction 9
2.2 Formal decision theory 9
2.3 Decision-theoretic and economic definitions 12
2.4 Interpretations of Probability 15

Chapter 3 Assessing Quantitative Approaches 18
3.1 Introduction 18
3.2 Overview of quantitative approaches 18
3.3 The applicability of quantitative approaches 19
3.4 Principle of indifference 20
3.5 Probability ranges, distributions, and second-order probabilities 21

Chapter 4 Incoherence Arguments 26
4.1 Introduction 26
4.2 Too strong - Peterson's incoherence argument 26
4.3 Too strong II - Sunstein's incoherence argument 30
4.4 Challenging the incoherence argument 33
4.5 The Precautionary Principle Schema 35
4.6 Limiting Strategies 36
4.7 Conditioning Strategies 40

Chapter 5 Gardiner's Proposal 44
5.1 Introduction 44
5.2 The intuitive argument for maximin 44
5.3 A formal grounding 48
5.4 The Rawlsian Core Precautionary Principle 49
5.5 Gardiner's uncertainty condition 54

Chapter 6 Assessing the Care Little for Gains Condition 62
6.1 Introduction 62
6.2 Assessing the Premises of Gardiner's Argument 64

6.2.1 Argument against Premise 1 64
6.2.2 Argument against Premise 3 65
6.2.3 Argument against Premise 2 66
6.3 Assessing the conclusion of Gardiner's argument.....73
6.4 Costly precautions.....75
6.5 Geoengineering overview 77
6.6 Geoengineering and the Precautionary Principle.....81
6.7 Typology of catastrophe cases 84

Conclusion 86

References.....88

List of Tables

Umbrella Decision Matrix 10
Umbrella Utility Matrix 11
Wager 1 45
Wager 2 46
Wager 3 46
Harsanyi's Job Matrix 50
Catastrophe Matrix 1 66
Catastrophe Matrix 2 67
Fatal Illness Matrix 70
Catastrophe Matrix 3 72

List of Abbreviations

Symbol	Definition
CBA	Cost-Benefit Analysis
CDR	Carbon Dioxide Removal
CELA	Canadian Environmental Law Association
CFCs	Chlorofluorocarbons
DM	Decision-Maker
ECCC	Environment and Climate Change Canada
EEA	European Environment Agency
GMOs	Genetically Modified Organisms
ICC	Intentional Climate Change
IPCC	Intergovernmental Panel on Climate Change
LEA	Lesser Evil Argument
MEU	Maximizing Expected Utility
PP	Precautionary Principle
PPS	Precautionary Principle Schema
RA	Risk Analysis
RBA	Risk-Benefit Analysis
SAG	Sulphate Aerosol Geoengineering
SRM	Solar Radiation Management
UNFCC	United Nations Framework Convention on Climate Change

Chapter 1 The Precaution Debate

1.1 Overview of the Precautionary Principle

The Precautionary Principle (PP) is a guiding principle of international environmental law and policy (Cameron & Abouchar 1991, pp.1-2; Trouwborst 2006, p. 6; Farber 2015, p. 1673; Scott 2013, p. 341), and is meant to aid decision-making in contexts in which technological and economic activities pose threats of harm to the environment and public health. One of the central components of the PP is the notion of acting in advance of scientific certainty; decision-makers should not wait until the causal relationships between an activity or technology and its possible harmful environmental effects have been established with certainty before taking preventive action. (Hartzell-Nichols 2012, p. 159; 2013, p. 308; Jordan & O'Riordan 1999, p. 23; Soule 2000, p. 312; Sunstein 2005, p. 4; Thalos 2012, p. 171). The primary motivation for the emergence and adoption of the PP is the apparent ineffectiveness of past environmental regulatory policy, which has for the most part been based on quantitative approaches to decision-making such as risk assessment and cost-benefit analysis. A report by the European Environment Agency (EEA) discusses several cases over the period from 1896-2000 in which regulators' delayed action on "early warnings" - credible preliminary scientific evidence of harm (though not proof of harm) - resulted in significant costs to human and environmental well-being (Harremoës et al. 2002). Such cases include: overfishing and the collapse of the California sardine fishery, benzene and leukemia, asbestos and lung cancer, polychlorinated biphenyls and liver damage, chlorofluorocarbons and the destruction of the ozone layer, antimicrobials as growth promoters in animal feed and the rise of antimicrobial resistant bacteria, and sulphur dioxide emissions and acid rain, among others (Ibid.). The current lack of action on the problem of global climate change is further evidence of this trend, as we have known of possible harmful

climate change effects due to greenhouse gas emissions since at least the 1990s, though there is uncertainty about the exact nature and likelihood of future climate scenarios. The PP is meant to buck this trend by making precautionary action the default response in situations like these, rather than costly delay. The idea is that it is better to be *safe*, i.e., to take early action to reduce or eliminate the risks associated with an activity or technology, than to be *sorry*, i.e., to delay preventive action and incur harm.

1.2 Problems facing the Precautionary Principle

The Precautionary Principle faces challenges of its own, however. First, there is currently no specific, universally agreed upon formulation of the PP; instead, there are several different versions to be found in the literature that range widely in their implications for environmental policy. The result is that it isn't clear which version of the principle decision-makers should use. Second, the PP has been the target of intense criticism. Opponents have accused various versions of the principle of being too vague (Bodansky 1991, p. 5; Peterson 2007, p. 306), vacuous and uninformative (Sunstein 2005), and even logically incoherent (Sunstein 2005; Manson 2002; Peterson 2006, 2007, 2017a). In general, critics who are skeptical of the PP (and who by and large favour economic cost-benefit analysis as the appropriate guide to environmental policy decision-making) view the principle as excessively pessimistic, anti-science, anti-evidence, anti-progress, and contradictory and perverse in its recommendations – in a word, *irrational*.[1] These challenges call into question the PP's legitimacy as a guiding principle of international

[1] According to Martin Peterson, the PP "tells us to replace traditional cost–benefit analyses with a more imprecise reasoning that focuses on possible negative effects. The precautionary principle therefore replaces the balancing of risks and benefits with what might best be described as pure pessimism" (2007, p. 306).

environmental law and policy, having lead some to lament that, "to date, precaution provides few if any operable guidelines for policy-makers nor does it constitute a rigorous analytical schema" (Jordan & O'Riordan 1995, p. 192).

1.3 Goals and methodology

There is a large and diverse literature on the Precautionary Principle, primarily in the fields of law, risk analysis, economics, public policy, and philosophy, and there are a multitude of proposals for interpreting and formulating the PP, as well as many ideas associated with precautionary reasoning in general. The main goal of the present paper is not to offer a new theory of the PP, or even to argue for the truth of an existing one, but is far more modest in ambition – to show that there *is* a need for the PP, or something like it, in environmental law and policy. One reason for this is that many of the most common objections to the PP (vagueness, vacuity, incoherence, irrationality) can be defended against by making sensible modifications to the principle, as well as by adding conditions for appropriate precautionary measures (argued for in chapter 4). Another reason is that there are real limitations to popular quantitative principles for decision-making, such as maximizing expected utility and cost-benefit analysis (which is an applied version of maximizing expected utility). Though they do very well, and should be used, in certain contexts, there is a lack of sensitivity of such approaches to the uncertainties that are present in many environmental and technology decision-problems, which can ultimately lead to an obfuscation and misrepresentation of the decision (argued for in chapter 3). The PP, as a *qualitative*, not quantitative, decision principle, can perhaps do better.

As for methodology, much of the paper (as indicated by the title) is devoted to identifying and developing relevant analytical machinery, such as the concepts, tools, strategies,

frameworks, and taxonomies, that can be used to understand the component parts of the Precautionary Principle and their interplay, the situations and contexts in which the PP should and should not apply, how to defend against various criticisms, and how to formulate promising versions of the principle. Conceptual ground-clearing of this sort is essential to understanding the controversy and for navigating a complex debate that spans several academic fields with varying methodologies and approaches. Another significant portion of the paper involves examining and analyzing specific proposals for formulating the PP, identifying where they succeed and where they do not, and making targeted recommendations for modification and improvement. Much of this aspect of the analysis focuses on Stephen Gardiner's (2006) proposal for interpreting the PP as a constrained version of the maximin decision rule (explained in chapter 5). There are two reasons for doing so. One is that Gardiner's proposal is probably the most influential work in the philosophy of the PP and has since generated a wealth of philosophical debate over the principle. There is thus much that can be learned by using Gardiner's proposal as a focal point for structuring discussion and analysis of different aspects of the PP. Another reason is that the maximin rule appears to be a clear first choice for interpreting the PP. It is a qualitative decision rule, it is meant to apply in the presence of uncertainty, and it is very risk-averse, as it is concerned with avoiding worst-case outcomes, and with making sure that the worst that can happen is as good as possible.

In addition to the two conclusions already mentioned, namely, that many of the common objections to the PP can be defended against, and that there are significant limitations to quantitative approaches to decision-making in the context of environmental policy, I also forward three arguments against a maximin interpretation of the PP. The first is that the nature of the uncertainties in such contexts also poses problems for applying the maximin rule. Maximin

requires that the full range of the possible outcomes of our actions be known in order to be applied, but many environmental policy decision problems are such that there is significant uncertainty about the full range of possible outcomes (argued for in chapter 5). Second, maximin requires that precautionary measures be guaranteed to prevent threats of harm in order to be enacted, but this is often unknown as well (argued for in chapter 6). Third, Gardiner's maximin principle only applies in cases where the costs of implementing precautionary measures are relatively low, and where decision-makers do not care much about the foregone benefits of regulation, thus leaving out many cases where precautions are likely to be moderate or high cost (argued for in chapter 6). Precautionary measures can be justified in each case, but the maximin rule will be unable to recommend them. These three conclusions strongly suggest that the PP should not be interpreted as a version of the maximin decision rule, despite the initial plausibility of that project. Furthermore, they imply that any plausible formulation of the PP needs to be applicable in situations in which there is uncertainty about the range of possible outcomes, about the effectiveness of precautionary measures, and where precautionary measures are potentially costly. I also argue (in chapter 6) that a common condition on how costly precautionary measures can be in order to prevent threats of harm – the proportionality condition – should be recognized as defeasible in certain kinds of cases involving catastrophic possible outcomes, and thus should be amended. In the remainder of the current chapter, I explain the difference between "weak" and "strong" versions of the PP, and offer examples from various legal and policy declarations. I explain that a strong version of the PP is what is needed to address the problem of costly delays.

1.4 Weak and strong versions

One problem facing the Precautionary Principle, as was mentioned, is that there are several different versions of the PP to be found in the environmental law and policy literature, not all of which are compatible with one another, and which range widely in their implications for environmental policy.[2] For example, the "Rio PP", which was agreed upon at the 1992 United Nations Conference on Environment and Development held in Rio de Janeiro, states that, "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." The Rio PP prohibits proponents of a potentially harmful activity, be they governments, companies, or persons, from claiming that precautionary measures are unjustified because it isn't scientifically certain that the activity poses a risk. This line of reasoning is commonplace among those with vested interests in deregulation (think tobacco companies claiming that it isn't certain that smoking increases the risk of lung cancer, or oil companies stressing the uncertainties in climate science as a justification for continuing on with business as usual) and so the Rio PP would likely prevent one way in which precautionary action can be delayed. But it does not *require* that precautionary action be taken when there are threats of "serious or irreversible" harm. It implies, rather, that postponing precautions is permissible so long as the reason identified for postponement is not a lack of full scientific certainty (Hartzell 2009, p. 29). Opponents of precaution could well cite some other reason against taking precautionary action. As it happens, the Rio PP also suggests that precautionary measures need to be "cost-effective" in order to be taken. Thus, opponents of precaution could argue that proposed precautionary measures are too expensive or are inefficient from the point of view of

[2]Strictly speaking, the phrase "the Precautionary Principle" is inaccurate because there is no single, universally agreed upon version of the PP (Hartzell-Nichols 2013).

cost-benefit analysis, in the sense that the monetary benefits of precautionary measures do not outweigh their costs, and thus should not be taken (Ibid.).[3] Thus the Rio PP is widely viewed as a "weak" principle because it still leaves plenty of leeway for delaying or avoiding precautionary action – it does not ensure that regulatory policy will be more risk-averse (Soule 2000, p. 315). Other examples of so-called weak versions of the PP are the following:

> 1990 Bergen Declaration on Sustainable Development [Bergen PP]
> In order to achieve sustainable development, policies must be based on the precautionary principle. Environmental measures must anticipate, prevent and attack the causes of environmental degradation. Where there are threats of serious or irreversible damage, lack of full scientific certainty should not be used as a reason for postponing measures to prevent environmental degradation.
>
> 1992 United Nations Framework Convention on Climate Change [UNFCC]
> The Parties should take precautionary measures to anticipate, prevent or minimize the causes of climate change and mitigate its adverse effects. Where there are threats of serious or irreversible damage, lack of full scientific certainty should not be used as a reason for postponing such measures, taking into account that policies and measures to deal with climate change should be cost-effective so as to ensure global benefits at the lowest possible cost.

It is useful to contrast the Rio PP with a so-called "strong" version of the Precautionary Principle that was agreed upon at the 1998 Wingspread Conference held by the Science and Environmental Health Network. The "Wingspread PP" states that, "When an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause and effect relationships are not fully established scientifically." The Wingspread PP is more demanding than the Rio PP because it actually recommends precautionary action when there is a lack of full scientific certainty and does not tie the justification of precautionary measures to considerations of cost-effectiveness, which is more in

---

[3] This would depend on what is meant by "cost-effective." I discuss different interpretations of the cost-effectiveness requirement for precautionary measures in section 2.4.

keeping with the concerns raised by the EEA report. In general, it is common to differentiate strong and weak interpretations of the PP in the following way: weak precautionary principles claim that a lack of full scientific certainty is *not a reason against* taking precautionary action when faced with threats of harm (or, in other words, that scientific certainty is not a requirement for taking precautions) (Sunstein 2005, p. 18; Steel 2014, pp. 3, 19), while strong precautionary principles claim that precautionary action is *required* when faced with threats of harm even if there is a lack of full scientific certainty regarding causes and effects (Sunstein 2005, p. 24; Soule 2000, p. 316; McKinnon 2009, p. 190; Neumayer 2010, p. 104). A strong Precautionary Principle that makes precautionary action mandatory instead of optional is what is needed to address the problem of costly delays, i.e., to ensure that environmental regulatory policy will be more risk-averse. Other examples of so-called strong versions of the PP are the following:

> 1990 Third North Sea Conference [North Sea PP]
> The participants… will continue to apply the precautionary principle, that is to take action to avoid potentially damaging impacts of substances that are persistent, toxic, and liable to bioaccummulate even where there is no scientific evidence to prove a causal link between emissions and effects.
>
> Canadian Environmental Law Association [CELA]
> The precautionary principle denotes a duty to prevent harm, when it is within our power to do so, *even when all the evidence is not in*. This principle has been codified in several international treaties to which Canada is a signatory. Domestic law makes reference to this principle but implementation remains limited.
>
> Environment and Climate Change Canada [ECCC]
> The Federal Government has developed a document entitled "A Framework for the Application of Precaution in Science-based Decision Making about Risk" [Privy Council Office (PCO), 2003]. This paper addresses the application of precaution in its various forms - "precaution", "the Precautionary Principle" or "the precautionary approach" - all of which have three basic components: the need for a decision; a risk of serious or irreversible harm; and a lack of full scientific certainty.

Chapter 2 Decision Theory Background

2.1 Introduction

Many of the arguments in the literature on the Precautionary Principle require some understanding of decision theory, and so I will first introduce the relevant framework and concepts from that field. In this chapter, I explain "alternatives," "outcomes," "states of the world," "decision rules," "decision matrices," three normative decision rules (maximin, maximax, and maximizing expected utility), as well as the concepts of "risk," "certainty," "uncertainty," "ignorance," and "probability." Analyzing these concepts and settling on some definitions will also help clarify the issues and make them more tractable.

2.2 Formal decision theory[4]

In formal decision theory, decision-makers are presented with *alternatives*, or actions available to them, possible *states of the world*, or ways the world may turn out to be, and possible *outcomes*, which are a function of the decision-maker's alternatives and the possible states of the world. *Decision rules* are meant to tell the decision-maker which alternative(s) they ought to choose, or that they are rationally required to choose, given their goals, values, preferences, etc. A simple example that will help make the decision theory framework just described more concrete is the decision whether to bring an umbrella with you when you leave the house. In this example, suppose that you have two and only two alternatives, either to (A) bring an umbrella or to (B) not bring an umbrella. Suppose also that there are two and only two possible states of the world, either (S1) it rains or (S2) it does not rain. Further, suppose that if

---

[4] The presentation of the material in in this section draws heavily on Sven Hansson's (1994) introduction to decision theory, specifically, pp. 23-9, 60-1.

you choose to bring an umbrella and it rains, the outcome is that (A1) you stay dry but must carry an umbrella, and if it does not rain the outcome is the same. If you choose not to bring an umbrella and it rains then the outcome is that (B1) you are soaked but do not have to carry an umbrella, and if it does not rain the outcome is that (B2) you stay dry and do not have to carry an umbrella. Finally, suppose that you rank the possible outcomes, from best to worst, as follows: (B2), (A1), (B1). The following *decision matrix* represents the decision whether to bring an umbrella as it has just been characterized:

Umbrella Decision Matrix

	S1: it rains	S2: it does not rain
A: bring umbrella	A1: dry/encumbered	A2 = A1: dry/encumbered
B: don't bring umbrella	B1: soaked/unencumbered	B2: dry/unencumbered

The *maximin decision rule* tells the decision-maker to choose the alternative that has the least bad (i.e. best) worst possible outcome. In other words, the maximin rule tells you to maximize the minimum outcome that it is possible to achieve. In order to apply the maximin rule you need only compare the worst possible outcomes of the alternatives, and choose the alternative with the best one. The worst possible outcome of alternative (A) is (A1) and the worst possible outcome of alternative (B) is (B1). Since (A1) is better than (B1), and (A) and (B) are the only alternatives, alternative (A) has the least bad worst possible outcome. Thus the maximin rule tells you to choose to (A) bring an umbrella. We can contrast the maximin rule with the *maximax decision rule*, which tells you to choose the alternative with the best possible outcome. Since (B2) is the best possible outcome of the alternatives, the maximax rule tells you to choose to (B) not bring an umbrella.

It is also useful to contrast the maximin rule with *maximizing expected utility* (MEU), which requires you to calculate the expected utility (or expected value) of each alternative, and then choose the alternative with the highest expected utility (value). In order to calculate the expected utility of an alternative you must calculate the expected utility of its possible outcomes and then add them together. In order to calculate the expected utility of a possible outcome, you need to calculate its *probability-weighted utility*, that is, the probability of the outcome occurring multiplied by the utility of the outcome should it occur. In order to do this, you must assign a numerical probability to the outcome occurring (a number between 0 and 1 representing its likelihood) and you must also assign a numerical utility to the outcome (a number representing the utility/disutility or value/disvalue of the outcome should it occur). So, for the decision whether to bring an umbrella when you leave the house, suppose there is a fifty percent chance that it rains and a fifty percent chance that it does not rain that day (a 0.5 probability assignment to both S1 and S2), and also suppose that you assign the utilities 10, 5, and -10 respectively to outcomes (B2), (A1), and (B1). Then the expected utilities (i.e. probability-weighted utilities) of the outcomes (B2), (A1), and (B1) are 5, 2.5, and -5, respectively. This generates the following *utility matrix*, in which the descriptions of the outcomes are replaced by their expected utilities:

Umbrella Utility Matrix

	S1: it rains	S2: it does not rain
A: bring umbrella	2.5	2.5
B: don't bring umbrella	-5	5

The expected utility of an alternative is the sum of the expected utilities of its possible outcomes. Thus the expected utility of alternative (A) is 5 and the expected utility of alternative (B) is 0.

MEU tells you to choose the alternative with the highest expected utility, and so tells you to choose to (A) bring an umbrella. (Different probability and utility assignments can lead to MEU recommending not bringing an umbrella, of course.)

2.3 Decision-theoretic and economic definitions

In decision theory and economic theory, "risk" is defined as a situation in which the possible outcomes of the alternatives and their probabilities are known (Aven 2003, p. 39; Resnik 2003, p. 332; Steel 2014, p. 95; Hansson 2009, p. 425). For example, the decision whether to bring an umbrella after stipulating probabilities takes place in a situation of risk (also referred to as a "decision under risk"). In that case there are two alternatives and two possible outcomes for each alternative, and it is stipulated that the probability that it rains is .5 and the probability that it does not rain is .5. Thus all of the possible outcomes as well as their probabilities are known. "Risk," in this context, refers to an epistemic situation, or to the quality, nature, or degree of the decision-maker's knowledge. Though the decision-maker (DM) knows the possible outcomes and their probabilities, the DM does not know what the actual outcome of each alternative will be, because there is more than one possible outcome corresponding to each alternative. In a situation of certainty (or a decision under certainty), on the other hand, it is known what outcome will occur upon choosing any of the alternatives, i.e., there is only one outcome corresponding to each alternative.[5] The DM knows less in a situation of risk, and other epistemic situations are defined according to further degradation of the DM's knowledge.

---

[5]For example, assuming the probability that it rains is 1, then the decision-maker knows with certainty what the outcome of taking or not taking an umbrella will be.

"Uncertainty" is defined as a situation in which the possible outcomes of the alternatives are known but not their probabilities (also referred to as a "decision under uncertainty") (Aven 2003, p. 39; Hansson 1996, p. 369; Steel 2014, p. 95). The DM may have some knowledge about the probabilities of the outcomes and could perhaps specify a range (for example, a probability somewhere between .2 and .4), but the DM's knowledge is incomplete in ways that preclude precise probability assignments to the possible outcomes. Alternatively, the DM may have no knowledge at all about the probabilities of the possible outcomes (a situation often referred to as "complete uncertainty").[6]

---

[6] One of the primary sources of the certainty, risk, and uncertainty distinction is Robert Luce and Howard Raiffa (1957):

> As to the certainty-risk-uncertainty classification, let us suppose that a choice must be made between two actions. We shall say that we are in the realm of decision making under:
>
> (a) *Certainty* if each action is known to lead invariably to a specific outcome (the words prospect, stimulus, alternative, etc., are also used.
>
> (b) *Risk* if each action leads to one of a set of possible specific outcomes, each outcome occurring with a known probability. The probabilities are assumed to be known to the decision maker. For example, an action might lead to this risky outcome: a reward of $10 if a "fair" coin comes up heads, and a loss of $5 if it comes up tails. Of course, certainty is a degenerate case of risk where the probabilities are 0 and 1.
>
> (c) *Uncertainty* if either action or both has as its consequence a set of possible specific outcomes, but where the probabilities of these outcomes are completely unknown or are not even meaningful. (p. 13)

Similar notions were expressed even earlier by Frank Knight (1921), where "risk" is defined as a quantitatively measurable uncertainty, while "uncertainty" is defined as a quantitatively unmeasurable one:

> The essential fact is that "risk" means in some cases a quantity susceptible of measurement… It will appear that a measurable uncertainty, or "risk" proper, as we shall use the term, is so far different from an unmeasurable one that it is not in effect an uncertainty at all. We shall accordingly restrict the term "uncertainty" to cases of the non-quantitative type. (pp. 19-20)

"Ignorance" is defined as a situation in which not all of the possible outcomes of the alternatives are known, nor are their probabilities (also referred to as a "decision under ignorance") (Stirling 2008, p. 98; Aven 2011, p. 1520; Betz 2007, p.7; Harremoës et al. 2002, p. 170). In other words, in addition to unknown probabilities, there is an incomplete list of possible outcomes.[7] There are at least two ways a situation of ignorance can occur. One is when there are "known unknowns," situations in which the DM has conceived of an outcome but has not determined whether it is possible or not possible, and thus whether it should be included in the list of possible outcomes. An example of this sort was Oppenheimer's worry that a nuclear explosion could trigger an unstoppable burning of the atmosphere that would propagate globally, an outcome that was later ruled out as a possibility (Hansson 1996, p. 377). Another way is when there are "unknown unknowns," situations in which the DM has not even conceived of all of the possible outcomes of an alternative. An example of this sort was the use of chlorofluorocarbons and the depletion of the ozone layer, an effect that was not even thought of during the early use of CFCs (Stirling 2008, p. 100; Betz 2010, p. 94). In cases of ignorance, there is often the lurking possibility of unanticipated effects (surprises) relative to the DM's knowledge (Ibid.).

Matters are complicated by the fact that the decision-theoretic and economic definitions of risk and uncertainty differ from the risk and uncertainty concepts that are used in every-day speech, as well as in the discipline of quantitative risk analysis. An intuitive notion of the

---

[7] Strictly speaking, the definition of a situation of ignorance given here is a deviation from economic and decision theory, where "ignorance" is sometimes defined as a situation where the possible outcomes are known but there is no information at all about their probabilities (see, for example, Resnik 2003, p. 332). I have instead referred to this kind of situation as "complete uncertainty." Several authors (such as those just cited) have noted that a further situation needs to be recognized where we know even less than in a situation of complete uncertainty, namely where we do not have knowledge of some of the possible outcomes. They reserve the term "ignorance" for this kind of case, and I am doing the same.

referent of the word “risk” is an undesirable, or unwanted, possible event (Hansson 2002, p. 1), while “uncertainty” usually means “lack of knowledge” in every-day speech. In risk analysis, the aim is to provide quantitative descriptions of risks, that is, quantitative descriptions of undesirable possible events. Say for example, that an energy company wants to determine the risk of worker fatality in one of its power plants. Roughly speaking, the risk analyst would use available data, background knowledge, models of the system, and statistical methods to arrive at an estimate of either the probability of a fatality, such as a .002 probability of 1 fatality in one year, or, more commonly, an expected value of a fatality, such as .002 fatalities per year (Aven 2003, pp. 11-14). To avoid confusion, I will use phrases such as “situation of risk,” “situation of uncertainty,” “decision under risk,” “decision under uncertainty,” and so on, to make clear that it is the decision theoretic and economic definition of the term that is intended. Phrases like “risk,” “taking a risk,” “imposing a risk,” “severity of the risk,” “likelihood of the risk,” etc., as well as phrases like “uncertainty,” “degree of uncertainty,” “increase uncertainty”, “reduce uncertainty,” etc., on the other hand, signal that it is the every-day sense of “risk” and “uncertainty” that is being referred to, namely an undesirable possible event, and a lack of knowledge, respectively.

2.4 Interpretations of Probability

What does it mean to say that the probabilities of outcomes are reliable or unreliable, known or unknown? The standard decision-theoretic and economic definitions of the epistemic situations of certainty, risk, and uncertainty assume that known probabilities are those that can be determined objectively (Aven 2003, p. 28, 39; Steel 2014, p. 98). One way that probabilities can determined objectively is if they can be deduced *a priori*, as they can in some gambling situations (for instance, flipping a fair coin, or spinning a roulette wheel, where we can deduce

that each outcome has the same probability of occurring) (Aven 2003, p. 28). Another way that probabilities can be determined objectively is by calculating the relative frequencies of possible outcomes based on data (Ibid.). According to the objective frequency interpretation, probability is a number, *n*, representing the relative frequency of an event in a population (Aven 2011, p. 1518; Steel 2014, p. 98). A population (also called a reference class) is a series of similar activities repeated under similar conditions (or trials). The relative frequency of an event *E* in a population *P* is simply *E*/*P* (also referred to as a "success fraction"). Consider a simple coin toss example. If the event of the coin coming up heads occurs 520 times out of a population of 1000 flips, then the relative frequency, and thus the probability of the coin coming up heads on a single flip, on this view, is 520/1000, or .52. The larger the population, or the reference class, the more confident we can be that the calculated relative frequency closely approximates the *true* relative frequency of the event, where the true relative frequency of an event can be defined as the limit of its relative frequency in an infinite population (Aven 2011, p. 1519). So, after ten flips the coin may come up heads only two times, but we should not conclude that the probability of the coin coming up heads on a single flip is .2, as our sample size is very small. The larger the population, the more confident we can be that we know the probability of an event, or that the probability estimate is reliable, on this view. The notion of "calibration" has been suggested as a standard by which to measure knowledge of probabilities, or their reliability: "a probability assigned to an event is *calibrated* just in case it corresponds to the relative frequency of that event in a large sample of repetitions" (Steel 2014, p. 99; see also Hansson 1993, pp. 20-1).

Thus, in situations of risk (known possible outcomes and probabilities) probabilities can be determined objectively in either of the two ways just described, while in uncertainty (known possible outcomes but partially or completely unknown probabilities) and ignorance (some

possible outcomes unknown and unknown probabilities), they cannot. In the latter situations, probabilities cannot be deduced as they can in gambling situations, and there is typically insufficient data to reliably calculate the objective frequencies of the possible outcomes. Here, there can at most be *subjective* assignments of probabilities – either in the form of estimates (or best guesses) of the true (but unknown) objective probabilities of the possible outcomes, or, alternatively, as representations of the decision makers mental states, specifically their degree of belief or confidence that an outcome will occur (still measured using numbers between 0 and 1) (Aven 2003, p. 28; Steel 2014, p. 98.).

Ch. 3 Assessing Quantitative Approaches

3.1 Introduction

In this section I survey some of the main problems facing quantitative approaches to environmental policy decision-making, such as maximizing expected utility and its applied versions, cost-benefit analysis and risk-benefit analysis. I argue that there is reason to doubt the applicability of quantitative decision rules to many environmental policy problems. It is suggested that the PP, because of its nature as a qualitative decision rule, may fare better in such contexts.

3.2 Overview of quantitative approaches

Maximizing expected utility (MEU) says that the decision-maker should choose the alternative that has the highest expected utility, where the expected utility of an alternative is the sum of the probability-weighted utilities of all of its possible outcomes, and where the probability-weighted utility of an outcome is the probability that the outcome will occur multiplied by the numerical utility of the outcome. The MEU model of rational decision-making is predominant in decision theory, risk analysis, and economics (Hansson 1999, p. 910; Aven 2003, p. 30)). In economic cost-benefit analysis (CBA), as well as in risk-benefit analysis (RBA), utilities are interpreted in terms of money, such that the alternative with the highest expected net monetary benefit should be chosen (Resnik 2003, p. 332; Aven 2003, p. 39; Hansson 2007, p. 163; Ackerman & Heinzerling 2002, p. 1553). (Thus, for a precautionary measure to be justified according to CBA, it must be shown to be the alternative with the highest expected net monetary benefit.) Risk analysis (RA) typically interprets (dis)utility in terms of fatalities (or some other adverse outcome), and calculates the expected value of fatalities for each

alternative in order to compare them (Steel 2014, p. 108; Hansson 1993, p. 17). Quantitative approaches to decision-making such as these can also be described as "reductive-aggregative," since they involve a process of reducing the likelihoods and values of possible outcomes to numbers and then adding them together (Stirling 2008, p. 98). Importantly, the application of the MEU model of rational decision-making requires the assignment of reliable numerical probabilities and utilities to all possible outcomes in order to determine reliable estimations of expected utility.

3.3 The applicability of quantitative approaches

Quantitative approaches are most applicable when probability estimates are calibrated, i.e., when there is sufficient empirical data to reliably determine the objective frequencies of all possible outcomes, i.e. in situations of risk. Such contexts include, but are not limited to: actuarial and insurance settings (Aven 2003, p. 28), transport safety, engineering failure, known epidemics, and local disaster risk-management (such as flood under normal conditions) (Stirling 2008, p. 98), as well as medicine and occupational health (Hansson 1993, p. 20).[8] Quantitative approaches are less applicable when probability estimates are not calibrated, i.e., when there is insufficient empirical data to reliably determine the objective frequencies of all possible outcomes, i.e. in situations of uncertainty and ignorance. This is the case with many emerging technologies that pose risks of rare and extreme events and are likely to entail unknown dangers (Hansson 2009, p. 427; Hansson & Aven 2014, p. 1179). It is very difficult to apply MEU and CBA in such cases because, without reliable probability estimates of possible outcomes, the

[8] For example, "death rates at high exposure to asbestos are known from epidemiological studies" (Hansson 1993, p. 20).

expected utilities of the alternatives cannot be calculated with any degree of accuracy. Since many important cases in environmental and technology policy are like this (i.e. best characterized as decisions under ignorance[9]), this is a highly problematic result for MEU (and CBA).

3.4 Principle of indifference

One possible solution is for the proponent of MEU is to appeal to the principle indifference (principle of insufficient reason), which says that when nothing is known about the probabilities of the possible outcomes, each should be assigned equal probability (Resnik 2003, p. 335; Peterson 2017b, p. 54; Bognar 2011, p. 332; Steel 2014, pp. 23-4). For example, if there are two possible outcomes, then each is assigned .5, and if there are four possible outcomes, then each is assigned .25, etc. In other words, the assigned probability of each outcome is $1/n$, where n is the total number of possible outcomes. The principle of indifference is a transformative decision rule - it (magically) transforms a situation of complete uncertainty into a situation of risk so that MEU can be applied.[10] There are at least three problems with this solution. One problem is that the principle of indifference seems arbitrary – why not ascribe unequal probabilities to the possible outcomes when nothing is known about their likelihoods, or whatever probabilities the decision-maker wants (Peterson 2017b, p. 56)? This would seem just

---

[9] I argue for this claim in detail in section 5.5.

[10] The maximin rule can also be viewed as a transformative decision rule, as it basically says to assume that, whatever alternative you choose, the worst-case result of that alternative is certain to occur. If that were the case, you would ignore the better possible outcomes of each alternative and choose the alternative with the best worst outcome. So maximin basically converts the decision problem to a situation of certainty, where there is only one outcome that will occur when you choose an alternative, that is, only one outcome corresponding to each alternative.

as justified as ascribing equal probabilities, given that there is no evidence suggesting that one probability distribution is more accurate than another, as the principle of indifference assumes. And if it is arbitrary to assign equal probabilities to all of the possible outcomes in a situation of complete uncertainty, which it appears to be, then the expected utility calculations based on those assignments would also be arbitrary. A second problem is that, even if the principle of indifference were not arbitrary, it would only work for transforming a situation of uncertainty into a situation of risk, not for transforming a situation of ignorance into a situation of risk. In a situation of ignorance, since some of the possible outcomes are unknown, the number of total possible outcomes, n, is also unknown, and thus it would not be possible to know the value of $1/n$, which is the probability of each possible outcome. Thus, even if the principle of indifference weren't arbitrary, it would still leave MEU inapplicable in situations of ignorance, and hence inapplicable to many environmental policy problems. A third problem is that often *something* is known about the relative likelihoods of possible outcomes in environmental policy decisions, just not enough is known to ascribe accurate or reliable numerical probabilities to them (Steel 2014, p. 24). Thus the antecedent of the principle of indifference ("when nothing is known about the probabilities of the possible outcomes") is often not satisfied.

3.5 Probability ranges, distributions, and second-order probabilities

When there is uncertainty (lack of knowledge) about the estimation of parameters in a risk analysis, usually due to insufficient data, the standard approach is to represent this uncertainty with second-order probabilities (Aven 2003, p. 18). For instance, in the setting of risk, where the possible outcomes and their probabilities are known (as well as the values of the possible outcomes), the probability of a parameter X might be determined to be .3. Say that X is

the possible state of the world in which climate sensitivity is 3 degrees Celsius, for example. Climate sensitivity is, roughly, how sensitive global climate temperature is to a doubling of the current atmospheric carbon dioxide concentration (Sprenger 2012, p. 886). So to say that climate sensitivity is 3 degrees Celsius is to say that doubling the current carbon dioxide concentration would lead to a 3 degree Celsius increase in global average climate temperature. Suppose a simplified example – that there are four different possible states of the world, climate sensitivity is 1 degree (very low), 2 degrees (low), 3 degrees (high), and 4 degrees (very high). Under a situation of risk, we might determine the probabilities of these four states to be .2, .4, .3, and .1 respectively. We could then use these to determine the expected values of outcomes of different alternatives in different states of the world (combining them with utility assignments). But in a situation of uncertainty, say, there would be different probability distributions that are compatible with the available evidence (comprised of data, models, knowledge of the system), because of the lack of ability to determine relative frequencies (the probability distribution is indeterminate). Suppose that the probability distribution could also be (.1, .3, .4, and .2 respectively), or (.25, .45, .25, and .05 respectively). So we have three possible probability distributions, p1, p2 and p3. The idea would then be to weight the *reliability* of each distribution using a second-order probability (representing the likelihood that the distribution is correct). The analyst could conclude that p1 has a .5 probability of being the correct probability, that p2 has a .3 probability of being true, and that p3 has a .2 probability of being true. The expected utility (or, rather, reliability-weighted expected utility) of each alternative could then be determined by calculating its expected utility under each distribution, weighing the result under each distribution by the reliability of the distribution (e.g. multiplying the expected utility of an alternative under p1 by .6, and doing the same for p2 and p3), and summing them together. MEU

could then be modified to say to choose the alternative with the highest reliability-weighted expected utility. This method allows there to be uncertainty over the correct probability assignments while still resulting in a numerical description of the possible outcomes, which is what is required to apply MEU.

There are at least two problems with reliability-weighted expected utility. One problem is "that we don't know the correct probability distributions that should be applied to various parameters" (Pindyck 2017, p. 103; see also Frisch 2012, p. 120). When there is a severe lack of data, such as in situations of ignorance, there will be significant uncertainty about what probability distributions to include. There could be nearly indefinite interpretations consistent with the available data (or lack thereof) – basically any combination that adds up to 1. Another problem is that there will be significant uncertainty about the second-order probabilities, or about which quantitative reliability weights should be assigned to each probability distribution (Steel 2014, p. 24). The severe lack of data that makes it difficult to calculate reliable first-order probabilities will also make it very difficult to calculate reliable second-order probabilities. Thus reliability-weighted expected utility calculations are going to be very inaccurate in such circumstances as well. Reliability-weighted probability distributions does not solve the problem.

A similar proposal would be to use probability ranges or confidence intervals ("chunks" of the probability space instead of point estimates). There might not be enough data to determine the exact probability of a parameter, but there might be enough data to place it within a range of probabilities with high confidence, say, that there is a 95 percent chance that the probability that climate sensitivity is 2 degrees falls within the range of probability values .2 - .4. We could then use some method to calculate the expected utility of an alternative, say by just using the lowest and highest values in the range, or using the average value of the range. Similar criticisms apply

to this proposal, though. The probability range for some parameter will often be highly uncertain in contexts with few data. There are likely to be different interpretations of cut-off points, or the lowest and highest values the probability of the parameter could be that seem equally justified based on the evidence available.

To clarify, I am not claiming that these methods could not work – a lot would depend on how much data is available. In some cases, there might not be enough data to reliably determine precise probabilities, but perhaps enough to reliably determine a range, for example. The point, rather, is that in many problems in environmental, public health, and technology policy, specifically, there is such a scarcity of data regarding relevant parameters and possible outcomes that MEU and its applied versions (RA and CBA) *would* often not be able to reliably determine probability ranges, distributions, and reliability weights, in addition to precise probabilities, and would thus be rendered indeterminate in many cases. The Precautionary Principle, on the other hand, is widely taken to be a qualitative rule for decision-making, and qualitative rules typically require less information in order to be successfully applied. For instance, less evidence is needed to be justified in describing parameters and outcomes as, say, "possible," "plausible," "realistic," and "scientifically credible" (qualitative descriptions), than is required to be justified in describing them numerically, say, with precise probabilities, probability ranges, and distributions. And precautionary measures (and action generally) can be appropriate in cases where only qualitative descriptions of possible outcomes are justified by the available evidence, especially when the outcomes are potentially very bad. In such cases, though, quantitative approaches to decision-making would be inapplicable. Thus the PP, as a qualitative rule, may apply and yield recommendations in a much wider range of cases in environmental policy than

quantitative approaches such as MEU and CBA. Several authors note this promising feature of the PP and the use of qualitative descriptions:

> The general idea behind PP is that we should use it in a situation where we lack the type of empirical evidence that one would need for risk management, which uses objective probabilities obtained from risk assessment… The PP requires us to respond only to credible threats… A threat that is merely credible is not objectively probable. (Resnik 2003, p. 338).
>
> PP is designed to be capable of guiding decisions even when no quantitative risk analysis exists… PP can utilize qualitative rankings of knowledge conditions that may be justifiable even when scientific uncertainty is unquantifiable… The ability to function as a decision rule given only qualitative information is essential for meeting the demand…that scientific uncertainty not be a reason for inaction in the face of serious threats. (Steel 2014, p. 111)
>
> Acknowledging that probabilistic assessment is inappropriate in some case is by no means then to give up on the assessment. Assigning non-probabilistic qualities can commit us to less than assigning probabilities or probability ranges and thus can better represent uncertainty. Judging that it is a real possibility that climate sensitivity is 2 degrees Celsius does not require taking a position on the full range of climate sensitivity. Nor need rankings of climate sensitivities according to plausibility do so. (Katzav et al. 2012, p. 259)

Ch. 4 Incoherence Arguments

4.1 Introduction

Much of the debate in the philosophy of the Precautionary Principle concerns whether it is inconsistent. In this chapter, I consider two "incoherence" arguments against the PP, one from Martin Peterson (2007) and another from Cass Sunstein (2005). I argue that Peterson's argument can be defended against but that Sunstein's poses a more serious problem for proponents of the PP. I discuss several arguments, tools, and strategies that can be used to cast doubt on Sunstein's incoherence argument, and which pave the way for reasonable reconstructions of the PP.

4.2 Too strong - Peterson's incoherence argument

A strong Precautionary Principle that makes precautionary action mandatory is likely what is needed to address the problem of costly delays. However, critics have accused strong versions of the PP of being an irrational and incoherent guide to decision-making in environmental and public health policy. Martin Peterson (2007), for example, purports to show that the PP is incoherent through the use of examples that stress the great benefits of technology to society and public health that will be forgone by using the PP as a basis for decision-making.[11] According to Peterson's interpretation, the PP "holds that when decision-makers lack sufficient knowledge about the effects of a potentially dangerous activity, one should not proceed" (2007, p. 305). Peterson's main proposed counterexample to the PP is the case of clinical drug trials and his claim that the PP would "prohibit the majority of clinical trials" (Ibid.):

> It is impossible to foresee all possible adverse reactions that a new drug might cause without actually testing it on humans. Even if a trial ends in tragedy, we must

---

[11] This is a specific version of the general argument against the PP that it would stifle technological and societal progress.

> nevertheless continue with clinical trials simply because the benefits outweigh the risks; knowledge about safety and efficacy cannot be obtained in any other way, and it is important that tragic accidents should be accepted, depending on the expected benefits of a new drug. (Ibid.)
>
> Arguably, there is no fundamental difference between using the precautionary principle to abolish clinical trials, mobile phones or genetically modified food. In all three examples, there is scientific uncertainty about possible long-term effects on public or environmental health. (Ibid., p. 306)

Decision-makers lack sufficient knowledge about the potentially harmful effects of drugs being used in clinical trials – the whole point of clinical trials is to determine the possible harmful and beneficial effects of drugs, i.e., to acquire such knowledge – and thus the PP recommends that clinical trials not proceed (by Peterson's formulation of the PP). But clearly we must continue with clinical trials. If we followed the PP as a guide to decision-making, many of the most beneficial drugs to public health would never have been allowed. In fact, any activity or technology about which there is scientific uncertainty regarding possible harmful effects is not permitted by (Peterson's formulation of) the PP, no matter its potential benefits. Peterson concludes that "this shows that the Precautionary Principle is incoherent" (Ibid.). Peterson does not elaborate on how the purported incoherence of the PP is generated.[12] Nevertheless, if his conclusion that the PP would ban clinical drug trials, cell phones, and every other crucially

---

[12] Peterson also has a more advanced incoherence argument against the PP (see Peterson 2006, 2007, 2017a), which I do not assess in this paper. Peterson proves that the Precautionary Principle, defined as the principle that, "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 desirable" (2007, p. 306) is logically inconsistent with three other principles of decision theory, namely, Dominance, Total Order, and Archimedes, and thus that the PP should be jettisoned. This version of the incoherence argument can be challenged as well, either by arguing against Peterson's formulation of the PP again, or by choosing to reject Archimedes instead (see Boyer-Kassem 2017a and 2017b).

beneficial technology to society because of possible though uncertain risks is true, then using the PP as a basis for decision-making is surely irrational.

The main problem with Peterson's argument against the Precautionary Principle is that he seems to think that the PP requires one and only one response to a "potentially dangerous activity," namely a complete ban or stoppage of the activity (recall that the consequent of Peterson's formulation of the PP is "one should not proceed"). But this is a misinterpretation of the PP, for none of the versions so far discussed make this claim. They instead call for "precautionary measures" (Wingspread PP, UNFCC), to "take action" (North Sea PP), "to prevent harm" (CELA), and "the need for a decision" (ECCC). An outright ban on a potentially harmful activity is one kind of precautionary measure (or action, prevention, decision), but it is certainly not the only option available to the decision-maker.[13] The following is a list of just some of the other precautionary measures that are available: a temporary ban, or moratorium, on the activity; postponement of the activity; research on the activity, such as performing an environmental impact assessment to further understand potential harms; substitution by less hazardous alternatives used in the activity, if possible; economic incentives to reduce the prevalence of the activity, such as taxes or subsidies like a carbon tax or subsidizing renewable energy sources; adaptive measures to diminish the harmful effects of the activity, such as building sea walls to protect against rising sea levels due to climate change; and even the adoption of democratic/participatory decision-making procedures so as to identify all potentially affected parties before proceeding (Manson 2002, p. 267; Trouwborst 2006, pp. 165-77; see also

[13] Though it is the most aggressive option: "evidently, where an activity, technology, or substance poses an unacceptable environmental risk, the most cautious approach is simply to abstain from performing or allowing them" (Trouwborst 2006, p. 166).

The Wingspread Consensus Statement of the Precautionary Principle 1998).[14] Thus there are a wide variety of precautionary measures available to choose from, not just a ban. Strictly speaking, none of the versions of the PP so far considered claim that a ban is required, or any other specific measure for that matter; they at most require that *some* precautionary action be taken in response to potential harms but do not specify which action to take.

A decision-maker applying the Precautionary Principle to a clinical drug trial, for example, need not ban the trial because of the potential dangerous effects of the drug on the volunteers. Instead, they can and indeed should choose a precautionary measure that is more appropriate in the given circumstances. In Peterson's own example of a phase 1 clinical trail of an anti-inflammatory drug called TGN1412, in which "six healthy volunteers became critically ill," he notes that "several preventive measures could have been taken" (2007, p. 305). For instance, Peterson explains that "it was a mistake to administer the drug to all six volunteers at almost the same time; it would have been more vigilant to ensure that each volunteer remained healthy before proceeding onto the next" (Ibid.). Administering the drug sequentially and noting any adverse effects before proceeding is a risk reducing measure and thus counts as a precautionary measure.[15] Once it is recognized that there are precautionary measures other than a ban that can be chosen when implementing the PP, which Peterson himself deems rational, the apparent "incoherence" of the PP in his main proposed counterexample disappears.[16] Also, in the

---

[14] Many of these measures are not exclusive and can be combined to achieve the desired result.
[15] I take it as uncontroversial that there is a very close connection between risk reduction and precaution, in that precautionary measures *are* risk reducing measures, or at least they are meant to be.
[16] Another medical example is "replacing the preservative thimerosal with a less toxic alternative," instead of banning the Measles, Mumps, and Rubella vaccine, which is the precautionary measure of substitution by less hazardous alternative (Steel 2014, p. 29).

cases of cell phones and genetically modified food, the PP does not necessarily require a ban, but could call for less aggressive measures such as further research, or, in the case of cell phones, limits on the quantity of electromagnetic fields emitted (Steel 2014). The claim that the PP is committed to banning all technologies and activities that pose uncertain threats of harm regardless of their benefits is thus false, and so cannot be used to prove that it is incoherent or irrational.

4.3 Too strong II - Sunstein's incoherence argument

Cass Sunstein (2005) also claims to show that strong versions of the Precautionary Principle are incoherent. He formulates the (strong) PP as follows:

> For the moment let us understand the principle in a strong way, to suggest that regulation is required whenever there is a possible risk to health, safety, or the environment, even if the supporting evidence remains speculative and even if the economic costs of regulation are high. (2005, p. 24)

Sunstein's incoherence argument is meant to show that the PP is literally incoherent because "it bans the very steps it requires" (Ibid., p. 26). It is based on several examples which are intended to show how applying the PP to protect against one risk can result in the generation of new, countervailing risks. For instance, regulating genetically modified food may decrease the likelihood of ecological harm, but it may also result in the decreased availability of nutritious food sources, and thus may increase risks of harm to human health; reducing greenhouse gas emissions will likely decrease the risk of harmful global warming, but reducing greenhouse gases is also costly and may impose further economic burdens on the poor, thus imposing risks to human health; regulating nuclear power will likely decrease associated health and environmental risks, but may also increase reliance on greenhouse gas emitting energy sources such as coal-

fired power plants, and may thus increase the risk of harmful global warming (Ibid., pp. 26-7).[17] In such cases, Sunstein maintains that "stringent regulation runs afoul of the precautionary principle" (Ibid., p. 29). The idea is that the PP must also be used to consider the risks associated with enacting regulations, not just the risks associated with the activity or technology in question (see also Manson 2002, p. 274). But when this is done, the PP will often end up recommending against those very regulations. Thus the PP will recommend both for and against the same regulation in many cases, which is contradictory. Sunstein concludes that "the real problem is that the principle offers no guidance - not that it is wrong, but that it forbids all courses of action, including regulation" (Ibid, p. 26), and thus that it is "*paralyzing*" (Ibid., p. 33).

Sunstein's incoherence argument against the Precautionary Principle is more difficult for proponents of the PP to defend against. One reason for this is that his formulation of the (strong) PP is very charitable. It does not require a complete ban on any potentially harmful activity, like Peterson's version, only that "regulation is required." Of course, Sunstein holds that the generation of countervailing risks is due to "stringent regulation," but that shouldn't be a problem for his argument because many proponents of the PP will agree that the risks associated with activities like emitting greenhouse gases, nuclear energy, and the large-scale implementation of genetically modified organisms are serious enough to justify aggressive precautionary measures. Sunstein also interprets "possible risk" as requiring "a certain threshold of scientific plausibility" that must be met in order to justify regulation (Ibid., p.24). Thus his critique is not based on an uncharitable formulation of the PP that requires action in the face of any possible threat no matter how implausible it may be, which would easily lead to irrational

---

[17] Such cases are often referred to in the literature as "risk-risk trade-offs," or simply "risk trade-offs" (see Graham & Wiener 2008 and Hansen et al., 2008).

and incoherent recommendations. Instead, there must be some (unspecified) measure of scientific plausibility regarding the threat in question before preventive action can be taken. Lastly, Sunstein's formulation accurately reflects versions of the strong PP found in the environmental law and policy literature, such as the Wingspread PP, for instance (and is actually more sophisticated in certain respects). Thus there is little room for criticism of Sunstein's incoherence argument based on his characterization of the PP. Furthermore, Sunstein's claim that the risks associated with implementing precautionary measures also need to be taken into account, or subject to the PP, is also sensible, for it is possible that regulations themselves will result in unintended negative consequences.

Sunstein completes his case against the Precautionary Principle by combining his incoherence argument against strong versions of the PP with his complaint that weak versions of the PP are "unobjectionable, even banal" (Ibid., p. 24). He characterizes weak versions as holding "that a lack of decisive evidence of harm should not be a ground for refusing to regulate (Ibid., p. 18), and maintains that they "state a truism - uncontroversial in principle and necessary in practice only to combat public confusion or the self-interested claims of private groups demanding unambiguous evidence of harm, which no rational society requires" (Ibid., p. 24). Thus weak versions of the PP, according to Sunstein, are true, but obviously so - clearly we should not demand decisive evidence of harm before taking precautions.[18] Sunstein's two objections can be gathered to form a dilemma for the PP – it can either be formulated as a weak version of the principle, in which case it is trivial, or it can be formulated as a strong version of

---

[18] In other words, in environmental policy we are almost never in possession of *full* scientific certainty with respect to all the causal relationships between an activity and its effects. If full scientific certainty were a requirement for taking precautionary measures, then none would ever be taken, and so such a requirement is obviously too demanding.

the principle, in which case it is incoherent (Steel 2014, p.3; Gardiner 2006, p. 33; Boyer-Kassem 2017, p. 2). In either case the PP fails to guide environmental policy decision-making. As Daniel Steel sees it, Sunstein's dilemma objection "is a - if not *the* - central objection to PP" (2014, p. 21).

4.4 Challenging the incoherence argument

In section 4.3 I suggested that Sunstein's incoherence argument is more difficult to defend against for two reasons. First, it accurately reflects strong versions of the PP in the policy literature, and so does not depend on an uncharitable formulation. Second, the premise that precautionary measures (regulations) themselves can create countervailing risks is plausible, and so precautionary measures also need to be judged in light of the PP. In other words, the PP needs to consider the risks associated with all actions, including regulation. The contentious part of Sunstein's argument, however, is the claim that implementing a stringent regulation to mitigate one risk will often generate countervailing risks such that the PP will also recommend against the initial regulation, resulting in contradictory recommendations. This is the claim that is at the heart of the incoherence argument, but there are a couple of reasons to be skeptical of it.

One reason is that Sunstein's argument mistakenly assumes that the PP could only recommend against the initial regulation that is responsible for the generation of the countervailing risk. But there are other, better options available, such as implementing further precautionary measures to mitigate the countervailing risk (which would be the new target risk). For instance, if regulating nuclear power resulted in an increased reliance on greenhouse gas (GHG) emitting energy sources, and thus increased the risk of harmful global warming, then further regulations on GHG emissions could be implemented to reduce that risk (such as an

increased carbon tax). And if the costs of further regulations on GHG emissions were burdensome and increased harms to human health, then further measures could be taken to decrease that risk as well (such as a proportionate decrease in income tax), and so on, until overall risk is sufficiently reduced. This is a far more sensible approach than simply banning the initial regulation, which is contradictory and thus achieves nothing.

Another reason is that it is not always irrational to reduce or eliminate a risk while substituting another risk in its place, so long as the new risk is less serious and/or less likely than the original risk (Steel 2014, p. 20; Hartzell-Nichols 2002, p. 166).[19] In such a case, the PP would not be committed to recommending against the initial regulation, as that would be contrary to its goal of risk reduction, and thus would not yield incoherent recommendations. So, even if precautionary measures (regulations) create new risks, there is nothing irrational or incoherent about that so long as net risk is being reduced. (In fact, there are very few actions that are completely risk free, and so we should not expect or require that precautionary measures carry no risks of their own.) What would be irrational, and what would give rise to incoherence, is if the recommended regulation resulted in countervailing risks equal to or greater than the target risk. Then the PP would recommend against the initial regulation, as it would either have no effect on net risk or would actually increase overall risk, both of which are contrary to the PP's goal of risk reduction. It is fairly uncontentious that implementing precautionary measures will often create new risks, but it is far from certain that they will often create new risks equal to or greater in seriousness and/or likelihood than the original risk, which would be highly

---

[19] The original risk, the risk that we are initially trying to reduce, can also be referred to as the "target risk," while the new risk that arises can also be referred to as the "countervailing risk." I get this terminology from Graham & Wiener 2008 and Hansen et al. 2008.

dependent on the context and which precautionary measure is chosen. This is the contentious claim that Sunstein's incoherence argument relies upon. In any case, several scholars of the PP have reconstructed the principle such that it is even less likely to result in inconsistent recommendations. Before looking at several examples, I introduce an analytical tool that can be used to clarify and assess different versions of the PP and to formulate new ones.

4.5 The Precautionary Principle Schema

Versions of the Precautionary Principle vary depending on how they characterize the nature and severity of the harm, the likelihood or quality of knowledge of the threat, and the kinds of precautionary actions that should be taken. These elements, which can be referred to as the "harm condition," the "knowledge condition," and the "action condition," respectively, compose what Neil Manson calls the "general, three-part structure shared by every version" of the PP (2002, p. 265). This three-part structure takes the form of a conditional statement, which I will refer to as the "Precautionary Principle Schema" (PPS):

> For some activity A and effect E, if harm condition H and knowledge condition K are satisfied, then precautionary action P is required.[20]

---

[20] Many scholars of the PP have identified this structure. I am largely drawing on Steel's (2014) terminology of "knowledge condition", "harm condition", and "recommended precaution". Trouwborst (2006) characterizes the PP as a "tripod", the three legs being "threat of environmental harm", "uncertainty", and "action". Sandin (1999) uses "threat dimension", "uncertainty dimension", "action dimension", and, additionally, "command dimension" (which specifies the stringency of the recommended precaution, i.e., required or permissible). Manson (2002) uses "damage condition", "knowledge condition", and "e-remedy" (environmental remedy). Raffensberger and Tickner (1999) offer a "dual-trigger" characterization, where, if there is a "potential of harm" and if there is "uncertainty", then "precautionary action" should be taken. Though these authors use differing terminology, it is safe to say that there is general agreement among them about the basic elements that compose the PP, which can be used to generate different versions of the PP depending on how each element is specified.

The weaker, or less stringent, the harm and knowledge conditions, the more stringent and demanding the Precautionary Principle (Sandin, 1999). This is because the weaker H and K are, the less likely E needs to be given A (or the less evidence needed relating E and A), and the less harmful E must be in order for precautionary action P to be required. Also, the stronger the recommended precaution, the more stringent and demanding the PP, as the stronger P is the more activity A will be curbed (an outright ban on A being the strongest precautionary measure). In general, the weaker the knowledge and harm conditions, and the stronger the action condition, the more the PP is liable to run afoul of the charge of irrationality and incoherence. For instance, a PP that requires an absolute ban on any activity that poses any possibility of harm of any kind to the environment or human health would be maximally extreme and irrational - basically every activity satisfies these conditions, and so no activity would ever be allowed, including regulation itself. Once the knowledge and harm conditions are strengthened, and perhaps once the action condition is weakened, though, it is far more difficult to say that the PP is irrational or incoherent. I now discuss several examples from the philosophical literature on the PP that use the "limiting strategy" to do just this.[21]

4.6 Limiting strategies

---

[21] It should also be noted that the PPS is helpful in responding to the objection that the PP is too vague to guide environmental policy (Bodansky 1991). Many formulations of the PP do indeed suffer from vagueness, in the sense that there is a lack of specification of the harm, knowledge, and action conditions, which gives rise to ambiguity in interpretation and application. For instance, the Wingspread PP simply refers to "threats of harm to human health or the environment," and "precautionary measures," but does not specify how serious or likely the harms need to be in order to justify precautionary measures, nor does it indicate which measures should be taken (Turner & Hartzell 2004, pp. 451-52). Vagueness can be reduced, however, by using the PPS to provide more detailed specifications of the component parts of the PP.

Several scholars of the PP have reconstructed the principle such that it is even less likely to result in inconsistent recommendations, usually by limiting application of the PP to disastrous or catastrophic risks (thus not just any harm), by requiring that it only respond to scientifically credible or plausible threats (thus not merely possible, or conceivable, ones), and by putting conditions on which precautionary measures can be chosen. For example, Lauren Hartzell-Nichols (2009, 2012) defends the "Catastrophic Precautionary Principle," which states:

> *Appropriate precautionary measures should be taken against threats of catastrophe*
> Where:
>
> [1] Threats of catastrophe are those in which many millions of people could suffer severely harmful outcomes (defined as severely detrimental to human health, livelihood, or existence).
>
> [2] A precise probability of a threat of harm is not needed to warrant taking precautionary measures so long as the mechanism by which the threat would be realized is well understood and the conditions for the function of the mechanism are accumulating (this draws directly on Shue, 2010).
>
> [3] Appropriate precautionary measures must not create further threats of catastrophe and must aim to prevent the potential catastrophe in question.
>
> [4] Imminent threats of catastrophe require immediate precautionary action.
>
> [5] Threats of catastrophe that involve an imminent threshold or point of-no-return for effective precautionary action (beyond which precautionary measures are limited or unavailable) also require immediate precautionary action aimed at preventing this threshold from being crossed.
>
> [6] Non-imminent threats of catastrophe might warrant further study before further precautionary measures are implemented, provided a delay in taking precautionary measures will not prevent such measures from effectively preventing the catastrophic outcome in question.
>
> [7] The likelihood (or probability) of a catastrophically harmful outcome may affect what precautionary measures are taken, but not that precautionary measures should be taken. That is, a low probability outcome/event might warrant more minimal precautionary measures than a similar high probability outcome/event (which might require aggressive mitigation measures). (2012, pp. 160-61)

Hartzell-Nichols reconstructs the PP by limiting its scope to catastrophic risks, as defined in condition [1], thus making the harm condition very strong. She also strengthens the knowledge condition significantly, interpreting it as "the mechanism by which the threat would be realized is well understood and the conditions for the function of the mechanism are accumulating," thus requiring a high level of evidence regarding the threat. Daniel Steel also formulates a version of the PP that characterizes the harm and knowledge conditions in a similar way: "If a scientifically plausible mechanism exists whereby an activity can lead to a catastrophe, then that activity should be phased out or significantly restricted," (2014, p. 28) where a "scientifically plausible mechanism" is defined as "a causally related sequence of events (a) that is grounded in scientific knowledge, such as physical laws, and (b) for which scientific evidence exists of its actual occurrence" (Ibid., p. 31). Interpreted in this way, in order for the PP to recommend both for and against a regulation, and thus to be incoherent, the countervailing risks created by the regulation must rise to the level of catastrophic harm and there must be scientific evidence of the mechanism by which this would occur. Certainly, some regulations would satisfy these conditions, such as a complete and immediate ban on greenhouse gas emissions, which would probably cause an economic catastrophe affecting millions of people, but there are likely to be plenty of regulations that would not, such as a moderate carbon tax or cap and trade scheme, for example (Ibid., p. 34). Furthermore, Sunstein would tend to agree with these claims, as he is also on board with the project of reconstructing and limiting the application of the PP, and even argues for the "Anti-Catastrophe Principle," which is largely based on Gardiner's interpretation of the PP as a constrained version of the maximin rule.[22]

---

[22] "I do suggest that the Anti-Catastrophe Principle has a definite place in both life and law,

A possible problem with limiting the application of the Precautionary Principle to threats of catastrophe in order to defend against the charge of incoherence is that the PP may no longer be applicable in cases where activities raise threats of harm that are not catastrophic, but are nevertheless very severe. For instance, many of the examples of costly delays found in the EEA report are cases where resulting harms to human health and the environment were relatively local, affecting "only" thousands of people. The PP will not even be applicable to cases where activities raise severe threats of harm to hundreds of thousands of people, if we are using Hartzell-Nichols definition that "threats of catastrophe are those in which many millions of people could suffer severely harmful outcomes" (2012, p. 160). Also, we might think that precautions are justified in cases where knowledge of the mechanism that would lead to harmful effects isn't well understood, but where our understanding of causes and effects is more limited and speculative. The worry is that by limiting the scope of the applicability of the PP to defend against incoherence, we are also greatly reducing the role that it will play in environmental policy. Consider, for instance, another reconstruction of the PP by Taleb et al. that has extremely stringent harm and action conditions (though a much weaker knowledge condition):

> The precautionary principle (PP) states that if an action or policy has a suspected risk of causing severe harm to the public domain (affecting general health or the environment globally), the action should not be taken in the absence of near-certainty about its safety.
>
> We believe that the PP should be evoked only in extreme situations: when the potential harm is systemic (rather than localized) and the consequences can involve total irreversible ruin, such as the extinction of human beings or all life on the planet. (2014, p. 1)

---

applying to uncertain dangers of catastrophe, at least when the costs of reducing those dangers are not huge and when incurring those costs does not divert resources from more pressing problems" (2005, p. 114).

Taleb et al.'s formulation of the PP thus restricts its applicability to cases in which possible harms are maximally extreme, i.e., to existential risks, but this would make the PP applicable to only a handful of cases.

To be clear, though, none of these authors are committed to the claim that precaution is irrelevant outside of the conditions identified in their respective formulations of the PP, nor are they claiming that there are no other plausible reconstructions of the PP that are less limited in scope. For example, Steel formulates another version of the PP which state that, "If there is some scientific evidence that an activity leads to a significant and irreversible harm, then an alternative should be substituted for that activity if feasible" (2014, p. 28). This version, in which the harm, knowledge, and action conditions are each far less stringent, will apply to more cases than his catastrophe version of the PP. Hartzell-Nichols also suggests "that there might be other, independently justifiable Precautionary Principles" than the Catastrophic Precautionary Principle (2009, p. 106). Lastly, Taleb et al. clarify that their "point is that the PP is a decisive consideration for ruin problems, while in a broader context precaution is not decisive and can be balanced against other considerations" (2014, p. 2). This point suggests that other versions of the PP could be formulated for different contexts, since, presumably, if precaution is relevant in other contexts, then it is also appropriate to take precautionary measures in other contexts. We could thus identify the conditions that make precautionary measures appropriate in those contexts and formulate versions of the PP based upon them. Such versions, however, would likely be less justified in requiring a ban on the activity, which I believe is what Taleb et al. mean by precaution being "not decisive."

4.7 Conditioning strategies

In general, reconstructions of the PP that limit its scope to the most serious harms, such as a catastrophes or existential risks, are aiming to identify conditions where the most stringent precautionary measures are justified, such as significant restrictions, phase outs, and bans. In cases where risks are less serious, less stringent precautionary measures may be more appropriate. This is essentially the notion that the action condition, i.e., the required precautionary measure, should be proportional to both the seriousness and likelihood (or level of evidence) of the risk identified in the harm and knowledge conditions (Trouwborst 2006, p. 150; Resnik 2003, p. 341; Hartzell-Nichols 2012, p. 161; Steel 2014, p. 27; Sunstein 2005, p. 114; Manson 2002, p. 274). The "proportionality condition" is commonly included in reconstructions of the PP in order to help avoid the charge of incoherence, and is meant to reflect the idea that "the cure should not be worse than the disease." In the context of regulatory policy, this means that precautionary measures should not create risks equal to or greater than the target risk. Conditions [3] and [7] of Hartzell-Nichols' Catastrophic Precautionary Principle include versions of the proportionality condition: "appropriate precautionary measures must not create further threats of catastrophe," and "a low probability outcome/event might warrant more minimal precautionary measures than a similar high probability outcome/event (which might require aggressive mitigation measures)," respectively (2012, pp. 160-61). The proportionality condition is clearly a sensible requirement for formulating the PP, and for deciding which precautionary measures to take, though there are two caveats worth mentioning. One is that including the proportionality condition does not by itself dissolve the incoherence objection. If it is true that precautionary measures will often create risks greater than or equal to the target risk, as the incoherence argument maintains, then including the proportionality condition simply makes the PP inapplicable in those cases, instead of incoherent, which isn't a much better result for the PP.

Another caveat is that the proportionality condition will prevent precautionary action from being taken when it cannot be proven that a precautionary measure will not create risks equal to or greater than the target risk. But in certain cases, such an action might be our only chance at preventing an impending harm. In chapter 6, I argue that the proportionality condition should therefore be recognized as defeasible.

Another popular condition on precautionary measures is cost-effectiveness, or efficiency. According to Sunstein, "use of the principle should be closely attentive to the idea of cost-effectiveness, which requires regulators to choose the least costly means of achieving their ends" (2005, p. 115; see also Resnik 2003, p. 342). Similarly, Steel explains that "efficiency demands that if more than one precaution can be consistently recommended by the version of PP being used, then those with lower costs should be preferred" (2014, p. 30). Cost-effectiveness, or efficiency, can thus be understood to further condition precautionary measures that satisfy the proportionality condition. If there are multiple proportional measures available to mitigate the target risk, then the least burdensome measure should be chosen, thus further reducing overall risk. It is important to note that the cost-effectiveness (efficiency) condition is meant to be different from cost-benefit analysis' requirement that regulations must be economically efficient in order to be justified in the first place, in the sense that they need to be shown to have greater expected net present monetary benefits than costs. The cost-effectiveness condition is meant to apply to precautionary measures *after* the decision has been made to implement them, based on the harm and knowledge conditions of the PP. It is also important to note that the cost-effectiveness condition is not universally accepted, for it may be difficult to determine which of the available proportional precautionary measures is the most cost-effective (least burdensome) given the presence of uncertainty in many environmental policy decisions (Hartzell 2009, pp. 53-

4; see also Turner 2013). An inability to decide which precautionary measure is most cost-effective could stall the decision process and delay precautions from being implemented.

To conclude, I have argued that Sunstein's incoherence argument relies on the mistaken claim that PP could only recommend against the initial regulation that is responsible for the generation of the countervailing risk, that it depends on the controversial claim that precautionary measures often create new risks equal to or greater in seriousness and/or likelihood than the original risk, and that there are various limiting and conditioning strategies that can and have been employed to reconstruct versions of the PP that are unlikely to fall prey to the charge of incoherence. These arguments provide strong reason to doubt the success of incoherence argument against the PP.

Chapter 5 Gardiner's Proposal

5.1 Introduction

In the previous chapter, I discussed several arguments, tools, and strategies that can be used to cast doubt on Sunstein's incoherence argument, and which pave the way for reasonable reconstructions of the Precautionary Principle that are neither irrational nor vacuous. I now turn to an examination of Stephen Gardener's highly influential proposal for interpreting the PP as a constrained version of the maximin decision rule (thus limiting the scope of the PP). I first examine two general motivations for interpreting the PP as the maximin rule, followed by an explanation of Gardiner's view and the benefits of his proposal. In section 5.5, I assess and critique the uncertainty condition of the RCPP.

5.2 The intuitive argument for maximin

Several authors propose interpretations of the precautionary principle as a version of the maximin rule (Hansson 1997, Gardiner 2006, Aldred 2012, 2013, Sunstein 2005). One motivation for interpreting the PP as the maximin rule is the view that there is a strong intuitive connection between the maximin rule and the concept of precaution. As Sven Hansson explains, "the maximin rule gives priority to avoiding the worst that can happen. Therefore, it corresponds closely to the intuitive notion of being cautious, or exercising precaution" (1997, p. 295). Similarly, Jonathan Aldred writes that, "intuitively, the maximin rule reflects an extremely 'precautionary' outlook, one in which the desire to avoid the 'worse-case' outcome is given overwhelming priority" (2013, p. 134). These remarks involve two claims: first, that the maximin rule gives (overwhelming) priority to avoiding the worst outcome, and, second, that giving priority to avoiding the worst outcome is closely related to – or reflects – the notion of

precaution. Thus, the argument concludes, the maximin rule is closely related to – or reflects – the notion of precaution. Hence there is good reason, at least on the face of it, for interpreting the PP as the maximin rule. Call this the "intuitive argument".

The first claim of the intuitive argument is incomplete; the maximin rule does not just give (overwhelming) priority to avoiding the worst outcome: it gives (absolute) priority to avoiding every outcome that is worse than the least bad worst possible outcome (i.e. to choosing the alternative with the least bad worst possible outcome). But, as Hartzell rightly points out, "… eliminating a worst-case scenario is different from choosing the least bad worst-case scenario" (2009, p. 60). This can be easily seen via the decision-problem represented by the following matrix:

Wager 1

	S1: heads	S2: tails
A	A1: win $10	A2: lose $10
B	B1: win $5	B2: lose $5
C	C1: win $1	C2: lose $1

In this example, you are offered a choice between three gambles (alternatives), (A), (B), and (C). On each gamble a fair coin is tossed, the result of which (either heads or tails) determines your payout or loss, as listed in the matrix. If you are required to eliminate the alternative that has the worst possible outcome, then you must not choose (A), but are permitted to choose either (B) or (C). If you are required to choose the alternative that has the least bad worst possible outcome, though, then you can only choose (C). Thus, the first claim is incomplete – the maximin rule is committed to more than it states. And it is less clear whether giving priority to avoiding every

outcome that is worse than the least bad one is closely related to the notion of precaution. This is because there are cases in which it is appropriate to take precautions to avoid the worst possible outcome, but in which you do not have to choose the alternative with the least bad worst possible outcome in order to do so. Consider the following gambles:

Wager 2

	S1: heads	S2: tails
A	A1: win $10	A2: lose $1000
B	B1: win $500	B2: lose $2
C	C1: win $1	C2: lose $1

In this case, you can avoid the worst possible outcome of (A2) losing $1000 without having to choose the alternative with the least bad worst possible outcome, i.e., by choosing gamble B and not C. So, though taking precautions will sometimes require choosing the alternative with the least bad worst possible outcome – what the maximin rule requires – other times it will not. This provides reason to doubt the alleged close connection between the maximin rule and the notion of precaution.

The second claim of the intuitive argument – that giving priority to avoiding the worst outcome is closely related to the notion of precaution – can also be challenged. Though it will often be the case that it is appropriate to take precautions to avoid the worst possible outcome, especially when it comes to decision-problems in environmental policy, there will also be cases in which it is not. Consider the following gamble:

Wager 3

	S1: heads	S2: tails

A	A1: win $100	A2: win $9
B	B1: win 10$	B2: win $10

In this case, it would be inappropriate to exercise precaution in order to avoid the worst possible outcome because there is nothing to lose – it is a win-win situation. In other words, though there is a worst possible outcome in this case, there is nothing about it that would justify taking precautions to prevent it. While it is easy to construct cases in which the existence of a worst possible outcome does not justify exercising precaution, such as the one just examined, it is much more difficult to imagine a situation in which the existence of a bad possible outcome does not. This suggests that it is the *badness* of an outcome, not its "worseness," which justifies exercising precaution in any given situation. A more accurate version of the second claim, then, is that giving priority to avoiding bad outcomes is closely related – if not essential – to the notion of precaution. This provides further reason to doubt the purported closeness of the relation between the maximin rule and the notion of precaution, since the maximin rule does not mention the badness of possible outcomes *per se*, just their relative badness.

The intuitive argument also depends upon an implicit premise - the assumption that it is reasonable to base formulations of the PP on intuitions about the general concept of precaution. But whether and to what extent the PP should capture intuitive notions of precaution is an open question. A reasonable answer is that the PP should be at least somewhat related to the concept of precaution, otherwise the label of "precautionary" principle would be a misnomer, but that the PP need not reflect every idea that is connected to the concept of precaution, as there may be multiple such ideas not all of which are complimentary or consistent with each other (Steel 2014, p. 47; see also Thalos 2009). But since giving priority to avoiding bad outcomes is likely an

essential component of the concept of precaution, any proposed interpretation of PP should at least reflect this.

5.3 A formal grounding

Another motivation for interpreting the precautionary principle as the maximin decision rule (or as any formal decision rule for that matter) is that such an interpretation allows for the possibility of a formal/mathematical grounding of the PP. There has been a significant amount of formal work done on the maximin rule. For example, Kenneth Arrow and Leonid Hurwicz (1972), assuming some very basic rationality axioms, prove that under conditions of complete uncertainty (where the possible outcomes of a decision problem are known but there is no information whatsoever about their probabilities), it is only the best and worst possible outcomes that are relevant to the optimal decision – all other possible outcomes can be ignored (Aldred 2013, p. 133; Steel 2014, p. 53). This does not prove the maximin rule under such conditions, as the maximin rule only takes into account the worst possible outcomes of a decision problem, but others have built upon Arrow and Hurwicz's work. As Aldred explains, "Gilboa and Schmeidler (1989) provide an influential decision-theoretic extension of Arrow–Hurwicz" arriving at a decision rule that "… under complete uncertainty…yields the standard maximin rule" (2012, p, 1063). An assessment of purported formal proofs of the maximin rule is beyond the scope of the present paper, but the relevant point is that an interpretation that grounds the PP in a formal decision principle such as the maximin rule has the appeal of lending analytical rigor and thus credibility to the PP. The downside of such an interpretation, however, is that whatever criticisms pose problems for the chosen decision rule will then also attach to the PP. Having noted and assessed two preliminary motivations for interpreting the PP as the maximin rule

(namely the intuitive argument and the desirability of a formal grounding of the PP), let us now return to Gardiner's proposal.

5.4 The Rawlsian Core Precautionary Principle

Gardiner interprets the precautionary principle as a *constrained* version of the maximin decision rule. Typically, proponents of the maximin rule do not defend it as a rational principle to be followed in general, but rather defend it as appropriate to use only when certain conditions obtain (thus constraining the scope of its application). This is because it is fairly easy to construct decision-problems in which following the maximin rule clearly seems irrational. Consider, again, the decision whether to bring an umbrella when you leave the house. Suppose that you can determine that the chance of rain that day is very low; it looks sunny outside and the weather forecast does not call for rain. But you cannot, strictly speaking, rule out the chance of rain; weather forecasts are prone to error and you know from having lived in Calgary that the weather can take a turn for the worse rather unexpectedly. The maximin rule still tells you to bring an umbrella in this case, as that is still the alternative that has the least bad worst possible outcome. Furthermore, the maximin rule tells you to bring an umbrella on every day like this. So, if you follow the maximin rule, you will almost always be carrying an umbrella with you, even on the nicest of days. But such behavior is irrational (or, at the very least, quite odd).

John Harsanyi (1975) provides the following often cited counter-example to the maximin rule:

> Suppose you live in New York City and are offered two jobs at the same time. One is a tedious and badly paid job in New York City itself, while the other is a very interesting and well paid job in Chicago. But the catch is that, if you wanted the Chicago job, you would have to take a plane from New York to Chicago (e.g., because this job would have to be taken up the very next day). Therefore there would be a very small but positive probability that you might be killed in a plane accident. (p. 595)

This choice is represented by the following decision-matrix:

Harsanyi's Job Matrix

	S1: plane accident	S2: no plane accident
A: choose N.Y. job	A1: poor job/alive	A2 = A1: poor job/alive
B: choose Chicago job	B1: dead	B2: excellent job/alive

Since (A) choosing the New York job has the least bad worst possible outcome ((A1) having a poor job and being alive is better than (B1) being dead), the maximin rule tells you to choose the New York job. Harsanyi's point in this case is that it is irrational to forgo the highly probable outcome of having an excellent job in Chicago for the sake of preventing the highly unlikely outcome of dying in a plane crash on your way there. The rational choice, in this case, is to (B) choose the Chicago job.

One way that the maximin decision rule can go wrong, as we have just seen, is when the probability of the of the worst possible outcome is very low. Another way the maximin rule can be irrational is when the worst possible outcome, though perhaps quite likely, is not very bad (or not bad at all). Suppose, for example, that you are offered the choice between two wagers, (A) and (B). If you choose wager (A), there are two possible outcomes: you will either win $100 or you will lose $2. If you choose wager (B), there are also two possible outcomes: you will either win $1 or lose $1. Further suppose that the outcomes of each alternative are equally likely. The maximin rule tells you to choose wager (B), as that is the alternative with the least bad worst possible outcome (i.e. lose $1). But the rational choice, at least for most people, would be to choose wager (A) as it affords the opportunity to win a lot while risking very little.

We have seen several cases in which the maximin decision rule yields an irrational recommendation, and no doubt many more could be constructed. One reason why the maximin rule can go wrong, and often would go wrong if you were to use it as a general decision rule, is that it is insensitive to the probabilities of the possible outcomes; it always tells you to choose the alternative with the least bad worst possible outcome even when other alternatives have worse possible outcomes that are highly unlikely and better possible outcomes that are very likely (as in Harsanyi's jobs case). Another reason why the maximin rule can go wrong is that it is insensitive to the goodness and badness of possible outcomes; it always tells you to choose the alternative with the least bad worst possible outcome even when other alternatives have worse possible outcomes that are not very bad or not bad at all and better possible outcomes that are very good (as in the wagers discussed in the above paragraph). For these reasons, proponents of the maximin rule usually constrain its application to situations in which the probabilities of the possible outcomes are unknown (situations of uncertainty) and where the worst possible outcomes are actually quite bad. For example, suppose again that you are offered the choice between two wagers, (A) and (B). If you choose wager (A), there are two possible outcomes: you will either win $100 or you will be shot. If you choose wager (B), there are also two possible outcomes: you will either win $50 or will receive a slap on the wrist (Gardiner 2006, pp. 45-6). Suppose, further, that you have no information regarding the probabilities of the possible outcomes. In such a case, it is rational to use the maximin principle and choose wager (B). (You have very much to lose and very little to gain by choosing (A).)

Gardiner adopts similar constraints to the maximin rule (i.e. uncertainty and bad outcomes) by drawing heavily on the work of John Rawls (1999, 1974). Rawls, in his argument

for the superiority of a contractarian conception of justice over utilitarianism, argues that the following conditions, taken together, make it appropriate to apply the maximin decision rule:

> Thus it must be, for example, that the situation is one in which a knowledge of likelihoods is impossible, or at best extremely insecure. In this case it is unreasonable not to be skeptical of probabilistic calculations unless there is no other way out, particularly if the decision is a fundamental one that needs to be justified to others.
>
> The person choosing has a conception of the good such that he cares very little if anything for what he might gain above the minimum stipend that he can, in fact, be sure of by following the maximin rule. It is not worthwhile for him to take a chance for the sake of further advantage, especially when it may turn out that he loses much that is important to him.
>
> … The rejected alternatives have outcomes that one can hardly accept. The situation involves grave risks. (1999, p. 134)

Gardiner refers to these three conditions as the "Rawlsian Conditions" and asks us to suppose that they "… constitute criteria for the application of the precautionary principle such that if those criteria are met, the precautionary principle demands action" (2006, p. 47). Gardiner calls his version of the PP the "Rawlsian Core Precautionary Principle" (RCPP), which can be stated as follows using Gardiner's slightly altered descriptions of the three Rawlsian conditions: the RCPP demands that decision-makers follow the maximin decision rule – that they choose the alternative with the least bad worst possible outcome – when the following conditions obtain:

> [1] The first is that decision-makers either lack, or have reason to sharply discount, information about the probabilities of the possible outcomes of their actions. ["absence of reliable probabilities condition," or "uncertainty condition"]
>
> [2] Second, the decision-makers care relatively little for potential gains that might be made above the minimum that can be guaranteed by the maximin approach. ["care little for gains condition"]
>
> [3] Third, the decision-makers face unacceptable alternatives. ["unacceptable outcomes," or "disaster condition"] (2006, p. 47)

> [4] … The range of outcomes considered are in some appropriate sense "realistic," so that, for example, only credible threats are considered. ["realistic outcomes condition"] (Ibid., p. 51)

In terms of the Precautionary Principle Schema, the RCPP states that if there is a lack of probabilistic information regarding possible outcomes, and if there are alternatives that have possible outcomes that are unacceptable (i.e. disastrous or catastrophic possible outcomes), then decision-makers are required to follow the maximin rule so long as they care relatively little for the potential gains that might be made above the minimum outcome that maximin guarantees. (For example, when the foregone benefits, or opportunity costs, of following maximin are relatively small. This would be the case when the best-case possible outcomes of alternatives maximin does not recommend are not much better than the worst-case possible outcome of the alternative maximin recommends.)

It is important to clarify that Gardiner does not claim to identify the necessary conditions for the application of the PP, only sufficient conditions – conditions [1] through [4] are jointly sufficient for the application of the PP (Ibid., p. 48). His view is that the RCPP is a plausible version of the PP, but not that it is the only one possible. Gardiner does believe that he has captured the "core" conditions for the application of the PP – the conditions that are characteristic of paradigm cases where it is appropriate to apply the PP. For example, Gardiner maintains that the Rawlsian conditions are met in the problem of global climate change: there is significant uncertainty – i.e. lack of probability information – about climate trajectory and effects; there are possible disastrous or catastrophic outcomes, such as a sea level rise of ten meters; and the costs of climate change mitigation are, according to Gardiner, relatively small compared to the costs of catastrophic climate change – about one to two percent of Global GDP, which would set economic growth back a mere one or two years (Ibid., p. 55). Furthermore,

Gardiner holds that the RCPP does not run afoul of the incoherence objection, and that it would actually be very irrational *not* to follow the maximin rule when the Rawlsian conditions obtain:

> If one really were faced with a genuine possibility of disaster, cared relatively little for the potential gains to be made by avoiding disaster, and had no reliable information about how likely the disaster was to occur, then, other things being equal, choosing to run the risk might well seem like a foolhardy and thereby extreme option. (2006, p. 49)

In other words, it would be foolish for decision-makers not to ensure prevention of the possible disaster or catastrophe in an attempt to obtain potential gains that they care relatively little about. In sum, the RCPP purports to be a sensible, moderate reconstruction of the PP that is applicable to the paradigm cases in which proponents of the PP think it is appropriate. In the next section, I assess the knowledge condition of the RCPP, and in chapter 6 I assess the harm condition of the RCPP, which consists of the disaster condition and the care little for gains condition.

5.5 Gardiner's uncertainty condition

In the last section, Gardiner's Rawlsian Core Precautionary Principle (RCPP) was outlined, which is an interpretation of the Precautionary Principle as a constrained version of the maximin rule. Gardiner's general claim, it is important to note, is that the three Rawlsian conditions plus the realistic outcomes condition together constitute sufficient conditions for the application of the PP. And his interpretation of "applying the PP" is following the maximin decision rule. Gardiner's overall strategy is to determine the relevant features of cases in which it is clear that precautionary measures should be taken, and to use these features to formulate a moderate version of the PP that avoids the standard objections of vacuity, irrationality, and incoherence. In this section, Gardiner's "uncertainty condition" is assessed – the condition "…that decision-makers either lack, or have reason to sharply discount, information about the probabilities of the possible outcomes of their actions" (2006, p. 47). This in effect amounts to an

examination of the epistemology of the PP, that is, of the (lack of) knowledge and evidence that justifies, makes appropriate, is sufficient for, or is necessary for applying the PP in a given case.[23]

Gardiner's uncertainty condition, "…that decision-makers either lack, or have reason to sharply discount, information about the probabilities of the possible outcomes of their actions" (2006, p. 47), clearly rules out situations of risk, in which the probabilities of outcomes are known. However, it does not, by itself, rule out situations of ignorance, as it does not say anything about knowledge, or lack thereof, of the possible outcomes. It is actually because of Gardiner's adoption of maximin that the RCPP does not apply in situations of ignorance:

> The first Rawlsian condition explicitly rules out circumstances of risk, since it stipulates that probability information is known. More importantly, strict ignorance of some possibilities is also ruled out, because the RCPP assumes that in order even to apply a maximin principle one must be able to compare the possible outcomes. And, although this does not mean that one needs to be able to individuate all possible outcomes in a fine-grained way, it at least requires that the range of possible outcomes has been identified, and so excludes the possibility of outcomes outside the range. (Gardiner 2006, p. 50)

In order to apply the maximin rule, all of the possible outcomes need to be known – "the range of possible outcomes" needs to be identified – as that is the only way to determine which alternative has the least bad worst outcome. Thus, Gardiner's uncertainty condition, combined with maximin, *is* the decision-theoretic and economic definition of a situation of uncertainty, i.e. known possible outcomes but unknown probabilities. His claim is that paradigmatic cases in environmental policy in which it is appropriate to apply the PP ("core cases") are decisions under uncertainty, not risk or ignorance.

---

[23]I interpret "justifies" and "makes appropriate" in terms of rationality - applying the PP is justified, or appropriate, if and only if it is rational to apply the PP.

Greg Bognar (2011) and Miriam Thalos (2009, 2012) have objected (independently) to Gardiner's proposal on the basis that the PP is also applicable in situations of risk:

> But suppose for a moment that we had all the relevant probabilities. For instance, suppose we knew that the probability of a 10-meter sea level rise was exactly 0.5 given that no action is taken. Why should we think a precautionary approach is not warranted in this case? Why wouldn't the core precautionary principle apply? (Bognar 2011, p. 335)

> Suppose we were to learn, contrary to what we thought previously, and presumably also contrary to fact, that the "absence of reliable probabilities" condition is no longer met – that, in other words, there is now reliable information about likelihood of climate change. For example, we might learn (and for this hypothetical we do not need to assert this as fact) that levels of CO2 have a proportionate contribution to overall rise in average temperature, and that from this information we can predict tolerably well a host of important climate variables that impact on ecosystem health. While I am not claiming that this is now true, I am claiming that if this were true, and if we were to learn it, it would constitute reliable information about the prospects of climate disaster.
> What if it were indeed true? … My contention is this: if it were true, a precautionary attitude would still be advisable, even though the Gardineriean conditions fail to apply. (Thalos 2009, p. 43)

These hypothetical scenarios illustrate the point that if precautionary measures are justified in cases that are best described as situations of uncertainty, then they can be equally, if not more, justified when those cases are transformed into situations of risk. If all we knew was that a 10-meter sea level rise, which would be disastrous, was one of the possible outcomes of doing nothing in response to climate change, then that information alone would be enough to justify precautionary measures. If our knowledge of the climate system and our predictive abilities then became sufficiently advanced, so as to be able to determine that the probability of a 10-meter sea level rise is .5, then precautionary measures would be no less justified. In fact, since .5 is such a high probability for a disastrous outcome, this additional information would likely justify even more aggressive precautionary measures than in the original case. But because of the uncertainty condition, the RCPP would no longer apply (assuming that the probabilities of the other possible outcomes were also determined).

The following response can be made on behalf of Gardiner's proposal. Given the fact that the Rawlsian conditions pick out sufficient conditions for applying the PP, Gardiner is not claiming that the PP does not apply in situations of risk. He is claiming, rather, that paradigmatic cases in which applying the PP is appropriate – core precautionary cases – are often situations of uncertainty. So, as long as most of the core precautionary cases are decisions under uncertainty, not risk, the hypothetical counterexamples where PP applies in a situation of risk does not undermine the RCPP. Indeed, in the *actual* case of climate change (a presumed core precautionary case), there is an absence of reliable probabilistic information for many of the possible outcomes, a point that Bognar and Thalos would readily agree with:

> …there is considerable uncertainty about the likelihoods of different climate change projections, about the local impacts of climate change, and how soon the changes will occur. (Bognar 2011, p. 334)

> In the case of climate change, the "absence of reliable probabilities" condition is satisfied because the inherent complexity of the climate system produces uncertainty about the size, distribution and timing of change in climate variables. (Thalos 2009, p. 42)

Similar points could be made about other core-precautionary cases, such as geoengineering, genetically modified organisms, nanotechnology, and artificial intelligence, for example. Due to the novelties and complexities of emerging technologies and the systems they interact with (such as the atmospheric system, ecosystems, and the global economic system), it is very difficult to make accurate probabilistic predictions regarding their possible outcomes. Furthermore, one of the initial motivations behind the PP was that it should apply even when there is a lack of full scientific certainty regarding possible outcomes, a condition that is satisfied when there is a lack of reliable probabilistic information regarding such outcomes.

It could be objected that the hypothetical examples reveal a deeper problem for Gardiner's proposal, namely that a situation of uncertainty, by itself, does not actually provide any justification for taking precautions: "In justifying precautions, uncertainty does no real work – unless, of course, you simply insist that the precautionary principle is only relevant whenever there is complete uncertainty. But there is no reason for insisting on that" (Bognar 2011, pp. 335-36). Extending this point, if a decision problem can be transformed from a situation of uncertainty to a situation of risk without precautions becoming any less justified, as in the 10-metere seal level rise example, then the fact that the situation is one of uncertainty isn't playing any role in the justification for the precaution – it is, rather, the fact that there is a possibility of a disastrous outcome that is doing all the justificatory work. Thus a situation of uncertainty is not one of the distinctive features of cases that makes precautionary measures appropriate and, therefore, it is merely coincidental that core precautionary cases are often decisions under uncertainty.

The objection can be defended against by showing that uncertainty, in the sense of a lack of probabilistic information, can, and indeed does play a justificatory role in the implementation of precautions. Consider an altered version of the 10-meter sea level rise example, where this time the case is transformed from a decision under risk to a decision under uncertainty. Suppose, in one scenario, that the probability of a 10-meter sea level rise was determined to be .01. Now suppose, that in another scenario, a 10-meter sea level rise was determined to be one of the possible outcomes but that no probabilistic information for this outcome was available. I submit that the second case calls for a more risk-averse strategy, and thus for more aggressive precautionary measures. This is because, in the second case, the probability of a 10-meter sea level rise could be .01, but it could also be .1, .2, .3, .4, .5, etc., and so the strategy that the

decision-maker chooses must also take into account the possibility that the probability of a 10-meter sea level rise is relatively high, which would call for a more risk-averse strategy than if the DM knew the probability was .01, or relatively low. The relevant point is that a situation of uncertainty, by itself, can intensify risk and thus justify more aggressive precautionary measures.[24]

The main problem with Gardiner's uncertainty condition is not that it excludes situations of risk, but that it excludes situations of *ignorance*. Many core precautionary cases, such as the examples mentioned above – geoengineering, genetically modified organisms, nanotechnology, and artificial intelligence – are best described as decisions under ignorance, where not only is there an absence of reliable probabilistic information for possible outcomes, but also a lack of knowledge of (some) possible outcomes. There are at least two reasons why we cannot confidently say that we have identified all of the possible outcomes of emerging technologies like these. One reason is that there are past examples where we have failed to conceive of all the possible outcomes of emerging technologies, such as the example of ozone depletion from CFCs cited above. Another reason, which explains the first, is the inherent difficulty in forecasting outcomes especially when technologies involve large scale interventions on complex natural and social systems (Hansson 1996, p. 379). Even climate change due to greenhouse gas emissions is probably better characterized as a situation of ignorance, as there is persistent uncertainty and debate "about what possible states the climate system might exhibit" (Katzav et al. 2012, p. 261). And as one author sees it, "ignorance is arguably a pretty good description of the quality of

[24] I believe Rawls uses similar reasoning to support his claim that decision-makers in the original position, which is a situation of complete uncertainty, will be very risk-averse.

our knowledge about the consequences of human activity on the state of the environment in the long run" (Neumayer 2010, p. 102).

While Gardiner's position is that the Rawlsian conditions constitute sufficient conditions, and thus not necessary conditions, for the application of the PP, strictly speaking he is not claiming that the PP is not applicable in situations of ignorance (while the RCPP would not apply to a decision under ignorance, perhaps there is some other version of PP that would). However, Gardiner *is* claiming to have identified the "core" features of paradigm cases in which it is appropriate to apply the PP, namely the uncertainty condition, care little for gains condition, and the disaster condition. Thus, if many of the core precautionary cases are better described as situations of ignorance, which there is good reason to believe they are, rather than situations of uncertainty, then that is a significant problem for Gardiner's proposal – he has overlooked one of the main features of core precautionary cases. Furthermore, since the RCPP, as a maximin rule, is not applicable to situations of ignorance, it would thus not be applicable to many core precautionary cases. Maximin, in order to determine the least bad worst possible outcome of each alternative, "at least requires that the range of possible outcomes has been identified" (Gardiner 2006, p. 50). In situations of ignorance, however, there is significant uncertainty ("lack of knowledge" sense of "uncertainty") about whether the full range of possible outcomes has been identified, which would often render the maximin rule indeterminate. Any plausible interpretation of the Precautionary Principle, though, needs to be applicable and able to yield recommendations in such cases.

This point against maximin suggests a general problem for formal decision theory. In so far as formal decision principles require that the list of possible outcomes be known in order to be applied, as all of them seem to do, they will be inapplicable in situations of ignorance. I

cannot offer a full defense of this claim here, but I will note that this reveals a significant limitation for the use of formal decision theory in environmental and technology policy, given the importance and prevalence of situations of ignorance therein.

Chapter 6 Assessing the Care Little for Gains Condition

6.1 Introduction

In this chapter I will assess the second condition of Gardiner's Rawlsian Core Precautionary Principle (RCPP), the condition that "… the decision-makers care relatively little for potential gains that might be made above the minimum that can be guaranteed by the maximin approach" (2006, p. 47). Gardiner holds that the "care little for gains" condition, as he refers to it, is a characteristic feature of core cases "where it seems clear the precautionary principle applies" (Ibid., p. 48). (For brevity, I will refer to such cases as "core precautionary cases." These are paradigm cases for the application of precautionary principle.) For example, with respect to the problem of global climate change, which is an assumed core precautionary case, Gardiner maintains that "… the "care little for gains" condition is met because the costs of stabilizing emissions, though large in an absolute sense, are said to be manageable within the global economic system, especially in relation to the potential costs of climate change (Ibid., p. 55). Gardiner has in mind projections that peg the cost of stabilizing emissions at roughly 2% of global product per year, which he suggests "… constitutes only a delay of a year or two in projected growth" (though he acknowledges that 2% of global product per year is a very large sum of money in absolute terms) (Ibid.). Arguably, a year or two delay in global economic growth is a tolerable cost in relation to the potentially catastrophic possible outcomes of not responding to a rapidly warming climate. Such a delay would not likely translate into a significant decrease in peoples' quality of life in general, while unmitigated climate change may very well cause extreme reductions in the quality of life of many people around the globe, as well as deaths (for example, due to an increase in incidents of extreme weather events, such as hurricanes, floods, and forest fires).

Gardiner also identifies one of the most optimistic possible outcomes of climate change mitigation in support of his position: "Other models claim that such projections underestimate the indirect benefits of reduced fossil fuel consumption, and that, once this is factored in, mitigation might actually be beneficial even if there were no climate change" (Ibid.). It may be that transitioning to renewable energy based economies will lead to a long term relative *increase* in global economic growth in comparison with a business-as-usual strategy, rather than a reduction, which would be a win-win outcome (i.e. mitigating climate change *and* increasing overall global economic development). (The scenario just described is probably overly simplistic because there are likely to be winners and losers no matter which climate change strategy is chosen – some people, businesses, cities, and countries will fare better if an aggressive mitigation strategy is chosen, while others may be better off with different strategies, such as less aggressive mitigation, adaptation, and even business-as-usual.)

It is helpful to identify Gardiner's explicit as well as implicit premises in his argument for the conclusion that the "care little for gains" condition is satisfied in the case of global climate change. The following, I hold, is an accurate reconstruction of the argument. Premise 1: In the case of climate change, a mitigation strategy to stabilize emissions would cost relatively little (or would even be economically beneficial) compared to the potential catastrophic costs of unmitigated climate change (as per the reasoning in the first paragraph of this section). Premise 2: The maximin approach would recommend a mitigation strategy to stabilize emissions in the case of climate change. Conclusion 1: Thus, in the case of climate change, using the maximin approach would cost relatively little (or would even be economically beneficial) compared to the potential catastrophic costs of not using maximin (from Premise 1 and Premise 2). Premise 3: If using the maximin approach would cost relatively little (or would even be economically

beneficial) compared to the potential catastrophic costs of not using maximin, then decision-makers would care relatively little for potential gains that might be made above the minimum that can be guaranteed by the maximin approach (because any such gains would be small or non-existent). Conclusion 2: Therefore, In the case of climate change, decision-makers care relatively little for potential gains that might be made above the minimum that can be guaranteed by the maximin approach (from Conclusion 1 and Premise 3). Thus the "care little for gains" condition is satisfied in the case of climate change. And since global climate change is a core precautionary case – a paradigm case for the precautionary principle – this supports Gardiner's interpretation of the precautionary principle – the Rawlsian Core Precautionary Principle (assuming that one desideratum for formulating the precautionary principle is that it be applicable to core precautionary cases, which is a reasonable assumption).

6.2 Assessing the premises of Gardiner's argument

6.2.1 Argument against Premise 1

Many of Gardiner's premises can be challenged, as well as the implications of his conclusion. I have already explained Gardiner's justification for Premise 1 in the first two paragraphs. The problem, though, is that Gardiner identifies only two possible outcomes in support of his claim, which also happen to be two of the most optimistic possible outcomes of climate change mitigation, namely relatively little cost to mitigation and mitigation turning out to be economically beneficial. It is also possible that climate change mitigation will be a significant burden on many economies, especially those that are heavily reliant of fossil fuels. As Sunstein notes, "If we followed maximin for global warming, we would spend a great deal to reduce greenhouse gas emissions, and the result would almost certainly be higher prices for gasoline and

energy, probably producing increases in unemployment and poverty" (2005, p. 112). In other words, the regulations (such as carbon taxes or a cap and trade scheme) that are required to successfully stabilize emissions may cause significant economic downturn in some regions, resulting in job losses and increased rates of mental illness and suicide. So, the costs of mitigation may turn out to be high for many people, in which case decision makers would not care relatively little about the potential benefits of not pursing mitigation, such as preventing harmful economic downturns. The relevant issue here is not about which mitigation scenario will occur or which is most likely to occur, but that it is at least possible that climate change mitigation will not be relatively low cost. In such a case, though, Gardiner's Rawlsian Core Precautionary Principle will be inapplicable to the problem of global climate change.

6.2.2 Argument against Premise 3

Premise 3 is also problematic, for it is not always the case that decision makers care relatively little about the foregone benefits of choosing regulation, even if such benefits are relatively small. As Steel points out, "citizens and the politicians they elect generally do seem to care a great deal about small changes in GDP" (2014, pp. 54-55). Elected decision-makers may be overly focused on shorter term, economic benefits because they feel pressure to produce tangible results for the citizens whose interests they represent. So even if it were true that stabilizing emissions would be relatively low cost (say, 2% of GDP per year) compared to the potential catastrophic costs of unmitigated climate change, decision-makers may still care very much about those costs. In such a case, the Rawlsian Core Precautionary Principle would again be inapplicable to the problem of global climate change. (This critique of Premise 3 raises the question of why Gardiner chose to make the RCPP dependent on whether decision makers

happen to care about the costs of precautionary measures. Surely he could have just made the RCPP applicable when the costs of precaution are relatively low, regardless of whether decision makers care much about them, and would have avoided the objection just made.)

6.2.3 Argument against Premise 2

Premise 2, which states that the maximin approach would recommend a mitigation strategy to stabilize emissions in the case of climate change, appears to be the least controversial premise in Gardiner's argument. For suppose a simplistic framing of the climate change decision problem, as follows:

Catastrophe Matrix 1

	S1	S2
X: Mitigation	No catastrophe, but costs of mitigation	No catastrophe but costs of mitigation
Y: No Mitigation	Catastrophe, but no costs of mitigation	No catastrophe and no costs of mitigation

S1 is a possible state of the world where increasing global temperatures due to greenhouse gas emissions will lead to some catastrophic climate event if alternative Y (No Mitigation) is chosen, but will be prevented if alternative X (Mitigation) is chosen (this could be the case if climate sensitivity – roughly, the degree to which climate temperature is affected by an increase in greenhouse gas concentrations – is high). S2 is a possible state of the world where increasing global temperatures due to greenhouse gas emissions does not lead to a catastrophic climate event regardless of whether Mitigation or No Mitigation is chosen (this could be the case if climate sensitivity turns out to be low). Under this characterization of the climate change decision problem, the maximin rule does indeed recommend Mitigation, as that is the alternative with the least bad worst possible outcome (avoiding catastrophe but incurring the costs of

mitigation is better than the catastrophe occurring but not incurring the costs of mitigation, given the reasonable assumption that the climate catastrophe is worse than incurring the costs of mitigation).

However, as soon as we allow for the possibility that (S3) increasing global temperatures will lead to a climate catastrophe regardless of whether Mitigation or No Mitigation is chosen (for example, because of the possibility that the amount of greenhouse gases already emitted may be sufficient to cause a climate catastrophe, or that certain irreversible thresholds have been crossed that will ensure rapid warming regardless of whether we stop emitting now), the maximin rule recommends *against* Mitigation and thus for No Mitigation. Once we include S3, No Mitigation is the alternative that has the least bad worst possible outcome (incurring catastrophe but not incurring the costs of mitigation is better than incurring catastrophe as well as the costs of mitigation). Thus, maximin recommends No Mitigation in the case of climate change if there is a chance that mitigation will not prevent catastrophe.[25]

Catastrophe Matrix 2

	S1	S2	S3
X: Mitigation	No catastrophe, but costs of mitigation	No catastrophe but costs of mitigation	Catastrophe and costs of mitigation
Y: No Mitigation	Catastrophe, but no costs of mitigation	No catastrophe and no costs of mitigation	Catastrophe but no costs of mitigation

This is an intuitively problematic result for the Rawlsian Core Precautionary Principle. Mitigation is generally thought to be *the* precautionary strategy when it comes to dealing with

[25] This counterexample to the maximin rule originally appears in Chisholm and Clarke 1993 in their argument for the minimax regret rule as the correct interpretation of the precautionary principle. The example is also discussed in Clarke 2008, pp. 488-490; Hansson 1997, p. 299, and Steel 2014, pp. 57-60.

the problem of climate change, and that we should try to prevent catastrophic climate events from happening even if there is a chance that we cannot. What are some responses that can be made on behalf of Gardiner's view? It could be that the decision problem, as presented in the above decision matrix, is unfairly simplistic. It could be that there is *some* catastrophe that we may not be able to prevent from occurring given the amount greenhouse gases already emitted, but that there are other, additional possible catastrophes, which may be even worse, that will occur if emissions continue unrestrained. In such a case maximin may recommend Mitigation. But as long as it is possible that those additional catastrophes are unpreventable, then maximin will recommend against trying to prevent them as well, and the problem still remains.

Another option is to bite the bullet – to insist that the maximin rule is correct that Mitigation is an irrational response to climate change. The only way that this seems like a feasible response to me are possible scenarios where the costs of climate change mitigation are extremely burdensome and we have good reason to believe that mitigation will likely not prevent climate catastrophes. In such cases it starts to seem more plausible that mitigation would be an irrational response to climate change and that some other alternative should be chosen, such as adaptation, or even geoengineering, which I discuss below. The most problematic scenarios for the maximin rule, on the other hand, are when Mitigation is very low cost and we have good reason to believe that it is likely to prevent climate catastrophes, but cannot rule out the possibility that it will not. In such a case, the "care little for gains" condition is met, and so the RCPP is applicable, but it will recommend against Mitigation. But in such cases Mitigation is clearly the rational choice, and so we simply cannot bite the bullet and insist that Maximin rule is correct in this case. These are precisely the kinds of cases that Gardiner wants his proposal to choose the precautionary measure, but it can only do so when Mitigation is certain to prevent the

possible catastrophe. Even if all real-life cases were such that precautionary measures were guaranteed to prevent catastrophes (which they are not), and thus that maximin would always yield the rational choice, the theoretical possibility of such cases would be enough to sink Gardiner's view.

There may be a solution to this problem. In cases like the one just mentioned - where the precaution is low cost, there is the possibility of a catastrophic outcome, and there is a chance that the precaution will not prevent the catastrophe - there may be no meaningful difference in value between an outcome where a catastrophe occurs as well as incurring a small cost of implementing a failed precaution, and an outcome where a catastrophe alone occurs. This is due to the fact that when using qualitative descriptions of the value of possible outcomes, which we must do when accurate quantitative ascriptions are unavailable, there will be outcomes differing in minor respects that will be classified as falling into the same qualitative category of harm. As Steel puts it, "An unavoidable consequence of representing harm conditions with qualitative categories is that outcomes differing in severity of harm in some minor respect will often be lumped into the same category" (2014, p. 60). (Steel also points out that this needs to be the case to avoid "paralysis by analysis". If we worked with very fine grained qualitative categories where small differences were enough to justify classifying outcomes as qualitatively different, then we would never be sure which category possible outcomes would fall into, as we would have to know their exact effects, which is information that is not available to us (Ibid.)). Consider, for instance, a case in which a cancer patient can choose between two alternatives, that of undergoing an experimental treatment or not, say in the form of taking a pill. Suppose that there is a chance that the pill will cure the patient's cancer, and a chance that it will not, and that a known side effect of taking the pill is experiencing a mild headache for one hour. Suppose also

that the there is a chance that the patient will die if nothing is done and also a chance that the patient will survive if nothing is done. Suppose also that there are no other treatment options.

Fatal Illness Matrix

	S1	S2	S3
X: Take Pill	Survive, headache	Survive, headache	Die of cancer, headache
Y: Not take Pill	Die of cancer, no headache	Survive, no headache	Die of cancer, no headache

The relevant question for our purposes is the whether there is a difference in value between the two possible outcomes of S3. Is experiencing a mild headache for one hour and then dying of cancer worse than dying of cancer alone? There is a sense in which it does seem to be worse - experiencing some additional discomfort before dying of cancer is *slightly* worse than not experiencing it, and it wouldn't be irrational to prefer the latter to the former. But if the answer is yes, then the maximin rule says the rational choice is to not take the pill, as not taking the pill would have the least bad worst possible outcome, namely dying of cancer but not experiencing the headache. But taking the pill is clearly the rational choice in this decision problem - it could save your life at very little cost (notice that maximin would recommend not taking the pill even if there were a .99 probability of it saving your life). But, picking up on Steel's point, the real question is whether the two outcomes of S3 should be classified under different qualitative categories of harm, and the answer to this question may be more favorable to the maximin rule. If Steel is right, qualitative harm categories must be significantly different, not just differing in minor respects. Is the additional harm of experiencing a mild headache for one hour to the outcome of dying of cancer significant enough to warrant classifying the new outcome as falling into a worse qualitative harm category? Intuitively, it is not – it is so close in value to the original outcome that it should fall under the same harm category, which could be

labelled “personal catastrophe,” for instance. What *would* be worse than dying of cancer? Well, suppose that instead of the pill causing a mild headache for one hour, it caused utter torment - an excruciating migraine lasting one year, say. Then it is more intuitive to say that the that the worst possible outcome of taking the pill - a year of torture following by death - is indeed worse (falls into a worse qualitative harm category) than the worst possible outcome of not taking the pill - death alone.

Coming back to the climate change example, the analogous case to the cancer patient taking the pill that causes only a temporary and mild headache, but offers the chance of curing the disease, is the case where mitigation is very low cost and offers the chance of preventing a climate catastrophe from occurring (assume that mitigation would decrease global product by a small fraction of one percent, and that this translated to only some people having slightly less disposable income). In such a case, it is reasonable to say that there is no real qualitative difference between the possible outcome of a catastrophe occurring as well as incurring the very small costs of mitigation, and the possible outcome of the catastrophe alone occurring - the addition of the relatively small monetary costs of precaution is not enough push the already catastrophic outcome into an even worse qualitative harm category. And also, if the low costs of mitigation are not enough to make a difference to the value of possible outcomes in the worst case, consistency suggests that the same is true in better possible cases, such as in S2 - adding the very small costs of mitigation to the possible outcome where no catastrophe occurs is not enough to make the value of that outcome qualitatively worse. In fact, if the costs of mitigation are so small that they cannot make a significant difference in value to any of the outcomes, then they can be considered negligible costs and can reasonably be ignored in the decision problem. The climate change decision problem where mitigation costs are very low thus becomes:

Catastrophe Matrix 3

	S1	S2	S3
X: Mitigation	No catastrophe	No catastrophe	Catastrophe
Y: No Mitigation	Catastrophe	No catastrophe	Catastrophe

The maximin rule, however, does not fare that much better in the new decision problem. Now the worst possible outcome of Mitigation and No mitigation are equal. The standard move in such situations is to appeal to a sub-principle of maximin - the *leximin* decision rule, which says that if the worst possible outcomes of the alternatives are tied in value, then compare the second worst possible outcomes and choose the alternative with the least bad second worst possible outcome, and so forth. But the second worst possible outcomes of Mitigation and No mitigation are also tied (No catastrophe for each), and there are no further possible outcomes to compare. So, the maximin rule (plus leximin) is indeterminate in this case. But the obvious rational choice is Mitigation, and there is a basic principle of decision theory which explains this, namely *Dominance*: If one alternative achieves results that are at least as good as another alternative in every possible state, and which achieves results that are strictly better in at least one possible state, then the former ought (rationally) to be preferred to the latter. Mitigation dominates No mitigation in this case, and thus ought to be preferred. (Maximin would also be indeterminate in the cancer patient case, after ignoring the slight disvalue of a mild headache, even though taking the pill dominates not taking it, and thus ought to be preferred.) The "qualitative harm categories" point seemed to me the best way to try to rescue the maximin rule in these kinds of cases, but it did not help enough to save it. (The proponent of maximin and the RCPP could insist that cases where the precaution is not guaranteed to prevent the catastrophe

are rare, so the RCPP would be a good principle in practice. I think that this overly optimistic is unlikely to be the case in many situations. It would probably be better to come up with a different formulation of the PP.) Furthermore, as long as the costs of precautionary measures are large enough to make catastrophe + costs of precaution qualitatively worse than catastrophe alone (such as when mitigation causes significant economic downturns, say), the maximin rule will always recommend against the precaution. But it may still be rational to take precautions even when this is the case. It is to this point that I now turn.

6.3 Assessing the Conclusion of Gardiner's Argument

Even if Gardiner's argument is sound, i.e. that the premises 1-3 are all true and they logically imply the conclusion that the "care little for gains" condition is satisfied in the case of climate change, there would still remain a serious problem for the Rawlsian Core Precautionary Principle. Cass Sunstein, I believe, correctly diagnoses the main problem:

> [The RCPP] ... risks triviality, above all because of condition (3) [the "care little for gains" condition]. If individuals and societies can eliminate an uncertain danger of catastrophe for essentially no cost, then of course they should eliminate that risk. But the real world rarely presents problems of this form… In real disputes, the elimination of uncertain dangers of catastrophe imposes both costs and risks." (2005, p. 112)

The diagnosis has two parts. The first is that the "care little for gains" condition limits the application of RCPP to uncontroversial cases where the decision to take precautions is easy. This threatens to make the RCPP trivial and unhelpful because it cannot yield answers in controversial cases, only in cases where the rational choice is clear.[26] The second part is that controversial

---

[26] Critics of Gardiner's view have extended Sunstein's claim (Bognar 2011, Steel 2014) by pointing out that the easiness of decisions in cases where the "care little for gains" condition is met means that other decision rules besides maximin will also recommend the obviously rational choice, such as maximizing expected utility and the minimax regret rule (though remember that

cases are the norm - in most environmental and health policy problems, there is much that can be gained *and* much that can be lost by taking precautions to prevent against possible catastrophes. I have already discussed the possibility that climate change mitigation is higher cost than the optimistic scenarios Gardiner identifies, possibly imposing economic burdens which translate into substantial reductions in peoples' quality of life. Thus, there may be significant possible losses associated with a mitigation strategy, as well as significant possible benefits, namely preventing catastrophic climate events. (Of course, critics of climate change mitigation will often argue that the economic costs of mitigation are very high and that the possible benefits are virtually non-existent, as they typically hold that the possibility of catastrophic climate events is so low as to be practically negligible, or on a par with other extremely unlikely events, such as invasion by space aliens.) Other examples are more obvious than climate change mitigation, in which precautionary measures may or may not be low cost. For instance, geoengineering as a response to climate change (which I discuss in more detail in the next section), genetically modified organisms (GMOs), artificial intelligence (AI), and bioengineering and nanotechnology

---

the maximin rule cannot even choose the obviously rational choice when there is just a slight chance that the precautionary action will not prevent the catastrophe, as I have argued above). The conclusion is that Gardiner's argument does not uniquely support an interpretation the PP as the maximin rule, as it could just as well be an argument for an interpretation of the PP as a version of maximizing expected utility, or a version of the minimax regret rule, or really any other decision rule that yields the rational choice in cases where there is little to be gained and much to lose by not taking precautions. "Several alternative decision rules to maximin, therefore, coincide with it in the narrowly restricted Rawlsian conditions. This raises serious problems for regarding maximin a "core" of PP, because Gardiner's proposal does not provide grounds for claiming that PP should be understood in terms of the maximin rule rather than one of these alternatives" (Steel 2014, p. 56). "My main argument is that Rawls's conditions do not establish maximin as the uniquely rational rule even when the conditions do obtain. In such cases, other principles can also lead to reasonable choices. There is no need for maximin at all. The strategy of using maximin to defend the precautionary principle makes the principle dispensable" (Bognar 2011, p. 331).)

all promise great benefits to society, but also carry great risks. Thus the costs of precaution, such as stringent regulation, may be high, say, in the form of forgoing the potentially great benefits to society of these technologies in order to protect against possible catastrophic outcomes. These are cases in which precaution is nevertheless intuitively appropriate – they are core precautionary cases. But, because of the "care little for gains" condition, the RCPP does not apply, for decision makers will indeed care about the costs of precautionary measures when they are high. In light of Gardiner's own desideratum for interpreting the precautionary principle, namely that it be applicable to core precautionary cases, the RCPP is inadequate because it fails to be applicable to a large swath of core precautionary cases.

Sunstein's conclusion, which we should accept, is that the "care little for gains" condition is too stringent and should be abandoned: "even if the costs of following maximin are significant, and even if choosers care a great deal about incurring those costs, it makes sense to follow maximin when they face uncertain dangers of catastrophe" (2005, pp. 112-13). That is, it is often rational to take precautions (assuming that's what maximin recommends) in order to prevent possible catastrophes from occurring, even if the costs of precaution are high. Sunstein explains that, "The hardest question here is: How much cost does it make sense to incur in the service of maximin?" (Ibid.). How can we determine whether precautionary measures are rational when they are high cost? This is one of the toughest questions to answer when coming up with a good formulation of the precautionary principle, and it is a question that needs to be dealt with, for many environmental and technology policy problems are controversial cases of this kind.

6.4 Costly Precautions

I will now look at some possible options for how the PP should deal with such cases. In cases like these, implementing precautions comes with a high cost, either in the form of forgone opportunity benefits, or direct effects of precautionary measures. Sunstein, through his "Anti-Catastrophe" principle, proposes that moderately costly and even significantly costly precautionary measures are justified in situations of uncertainty to prevent a possible catastrophic outcome, as long as they are "not hugely costly" (2005, p. 114). For example, if mitigating climate change would involve a 50 percent reduction in the quality of peoples' lives worldwide, then Sunstein thinks it would be better not to mitigate climate change, but instead to do further research to increase our knowledge of the likelihood of a climate catastrophe occurring before taking such costly action (Ibid., p. 113). Sunstein adds the condition that precautionary measures themselves not create further threats of catastrophe (proportionality condition). Arguably, a 50 percent reduction in the quality of peoples' lives worldwide is itself a catastrophic outcome, and so mitigation would be too costly a precautionary measure to be taken if that were the case - it would violate the proportionality condition. Sunstein claims that we would need a sufficiently high demonstrable probability that a climate catastrophe will occur (and presumably knowledge that the climate catastrophe is worse than a 50 percent reduction in quality of life) to take such costly precautionary measures (thus basically claiming that hugely costly precautions could only be justified in situations of risk, not uncertainty) (Ibid.). He also adds the condition that precautionary measures should be as least costly as possible (cost-effectiveness condition) (Ibid., p. 115).

This kind of response is fairly common in the literature. Hartzell-Nichols argues for the "Catastrophic Precautionary Principle," which says that appropriate precautionary measures should be taken to prevent threats of catastrophe, and that precautionary measures should be

shown not to create further threats of catastrophe (proportionality condition), and includes several other conditions (2012, p. 160-161). So, a popular answer as to how costly precautionary measures can be in order to prevent possible catastrophic outcomes is that they can be significantly costly, but not so costly as to introduce their own possible catastrophic outcomes. And that if there is more than one precautionary measure that meets this criterion, the least costly one be chosen. I think this kind of proposal is on the right track, though there are a couple of remaining problems. One is that there are certain cases in which a potentially highly costly precautionary measure can be justified, even though it cannot be proven not to introduce new possible catastrophic outcomes. I believe this point can be illustrated by geoengineering, particularly Solar Radiation Management (SRM), as a possible solution to climate change.

6.5 Geoengineering overview

In order to meet our duties of justice to future generations, we need to prevent the occurrence of catastrophic climate effects due to global warming. To do this, say the majority of climate scientists, our goal should be to limit the increase of global average temperatures to 2 degrees Celsius from preindustrial levels. Much of the current debate focuses on strategies for mitigating climate change caused by atmospheric carbon dioxide, that is, cutting emissions so as to limit the concentration of CO2 in the atmosphere. Mitigation strategies will likely involve significant social and economic changes in order to shift world economies away from a heavy reliance on greenhouse gas emitting energy sources. In other words, mitigation strategies will likely involve significant societal restructuring. A relatively new strategy for dealing with anthropogenic climate change (which can be proposed as an alternative or a complement to mitigation strategies) is *geoengineering*. According to David Keith, "Geoengineering is the

intentional large-scale manipulation of the environment, particularly manipulation that is intended to reduce undesired anthropogenic climate change" (2000, p. 245). There are various geoengineering technologies and they are typically classified under two categories: Carbon Dioxide Removal (CDR) and Solar Radiation Management (SRM). CDR technologies, as the name suggests, aim to remove CO2 particles from the environment. Examples include Afforestation, Direct Air Capture, and Ocean Fertilization. SRM technologies aim to increase the earth's albedo, or the percentage of incoming sunlight that is reflected by earth back into outer space. Examples include Cloud Whitening, Space Reflectors, and Stratospheric Aerosols (small particles such as Sulphate, which reflect a portion of sunlight in the stratosphere). Geoengineering is one of three possible strategies for responding to the problem of anthropogenic climate change, the other two being Mitigation and Adaptation. (These strategies, it is important to note, are not mutually exclusive.)

Dale Jamieson (1996) is widely credited as the first philosopher to open door to discussion on the ethics of geoengineering. Jamieson puts forth four conditions that must be met in order for the deployment of geoengineering technology to be morally permissible: "(1) the project is technically feasible; (2) its consequences can be predicted reliably; (3) it would produce states that are socio-economically preferable to the alternatives; (4) implementing the project would not seriously and systematically violate any important, well-founded ethical principles or considerations" (1996, p. 326). Condition (2) is not currently met: the science of geoengineering is still its infancy and scientists are not now in a position to say with confidence what the environmental effects of geoengineering the climate will be. (Jamieson raises the issue of whether it is even possible for us to reliably predict the consequences of geoengineering: "If we couple the pervasiveness of unintended effects and the tendency of experts to overestimate

their expertise with the incredible complexity of the climate system, the grounds for skepticism about reliably predicting the effects of ICC [Intentional Climate Change] seem very strong" (1996, p. 327). Because condition (2) is not met, due to what Jamieson calls "cascading uncertainties," condition (3) is not met either. If we do not know the consequences of geoengineering, then we cannot say whether a geoengineering proposal would produce states that are socio-economically preferable to the alternatives. (This also makes condition (4) far more difficult to assess, because if we do not know the socio-economic outcome of geoengineering, we cannot assess the outcome from the point of view of, say, distributive justice.) The main point here though is that geoengineering (of any sort) should not at present be deployed, and this is exactly what proponents of geoengineering readily admit (Crutzen 2006, Keith et al. 2010). What, then, is the issue? The issue is whether there should be an intensive research and development program on geoengineering technology. The debate, then, is not about deployment now, but about research into possible future deployment.

One of the main arguments for geoengineering research and development is known in the literature as the "Lesser-evil Argument." Dale Jamieson (1996) first stated the argument and it was later recast in a very influential paper by atmospheric chemist and Nobel Laureate Paul Crutzen (2006). The basic idea, in Jamieson's words, is this: "we may reach a point at which ICC is the *lesser of two evils*" (1996, p. 333, emphasis added). Two questions are pertinent to the understanding to this idea: Why might geoengineering become the "lesser of two evils"? And why might geoengineering be considered an "evil"? The answer to the first question is that the possibility of high climate sensitivity combined with blatant political inertia on climate change mitigation could very well put us in a situation in which we (or future generations) face catastrophic climate impacts (such as the drowning of coastal cities due to sea level rise).

Crutzen sites "studies by Andreae et al. (2005) and Stainforth et al. (2005) indicat[ing] that global average climate warming during this century may even surpass the highest values in the projected IPCC [Intergovernmental Panel on Climate Change] global warming range of 1.4-5.8 degrees Celsius" (2006, p. 211, citation suppressed). Combine this with the fact that global carbon dioxide emissions have actually increased since the early 1990s, the point in time at which it is generally agreed that the seriousness of anthropogenic climate change became public knowledge. ("For example, from 1990-2005, global emissions rose by almost thirty percent (from 6164 to 7985 million metric tons of carbon), and U.S. emissions rose by almost over twenty percent" (Marland et al., 2008; cited in Gardiner 2010, p. 286). Emissions are also growing more "rapidly in the recent past (from an average of 1.5-2% per annum to around 3% in 2007). Indeed, this growth is so rapid that they are currently at the very high end of projected emissions given back in 1990" (Ibid.). Most advocates for geoengineering research maintain that aggressive carbon dioxide mitigation strategies are "by far the preferred way" (Crutzen 2006, p. 211) for dealing with the problem, but are very skeptical (based on the dismal international track record) that the necessary mitigation will happen. And if mitigation efforts fail we will certainly face severe and perhaps catastrophic climate impacts. Thus, the reasoning goes, we need to research geoengineering as a possible way of avoiding catastrophic climate impacts. Here is Crutzen again: "Given the grossly disappointing international political response to the required greenhouse gas emissions…research on the feasibility and environmental consequences of climate engineering…, which might need to be deployed in the future, should not be tabooed. Actually…*its research should anyhow be intensified*" (2006, p. 214, emphasis added).

The answer to the second question, the question of why geoengineering might be considered an "evil", is that the kind of geoengineering a lesser evil argument of the sort just

described calls for is the most controversial kind of geoengineering technology available, namely Sulphate Aerosol Geoengineering (SAG). The Lesser Evil Argument (LEA) imagines a possible future state in which we need a relatively fast acting solution to an impending climate catastrophe. Using SAG, say by dumping quantities of sulphate particles in the stratosphere via airplane, is projected to begin reducing global temperatures fairly quickly (an estimated time frame of as little six months to two years). No other technology promises such a fast reduction in global temperatures. For example, just as the effects of increased carbon dioxide concentrations are realized many years after the initial addition (50-100 years), so too are the effects of lowering carbon dioxide concentrations via Carbon Dioxide Removal methods, which are generally thought to be safer than SAG methods (though removing carbon dioxide from the atmosphere may also have unforeseen and unintended consequences as well). Thus the Lesser Evil Argument likely calls for research into Solar Radiation Management technologies, and particularly SAG (this is indeed why much of the debate in the current literature focuses on SAG). Just a few of the possible consequences of SAG, however, are as follows: Ozone depletion, warming of the stratosphere, changes in the distribution of rainfall, and rogue implementation (a single company or country deciding to implement SAG without political consent), as well as unknown and unforeseeable consequences. So deciding to geoengineer may turn out to be less bad than certain climate catastrophes, but this decision might have its own bad consequences and possible violations of principle of ethics and justice. It is in this sense that geoengineering is thought to be an "evil."

6.6 Geoengineering and the Precautionary Principle

Taken in isolation, the PP would recommend against implementing SRM technology, due to the possible known negative consequences as well unknown possible consequences. But if we come to a point where we have failed to implement aggressive mitigation and adaptation strategies, and there is strong evidence for an impending catastrophe due to runaway climate change, SRM could actually be viewed as a precautionary measure against the impending catastrophe. However, it would be very difficult to prove that any SRM technology, at least within the timeframe available for effective action, would not create a climate catastrophe of its own, or would not exacerbate the situation. That's because, with any large-scale intervention on complex systems, it is nearly impossible to foresee all possible consequences. However, if in this situation we don't implement SRM, we aren't protecting ourselves at all from the catastrophic outcomes of climate change. Arguably, the precautionary action would be to implement SRM in order to prevent the highly likely catastrophic outcomes of climate change, while knowing that we might be introducing further catastrophic outcomes due to SRM. This argument lends support to the conclusion that we should be researching SRM technologies now, so that *if* we find ourselves in such dire circumstances (which is a real possibility given the current mitigation effort), we at least have the option to implement SRM to prevent catastrophic climate change.

Matters are complicated by the fact that researching SRM might further weaken the mitigation effort, thereby making the worst-case scenario a self-fulfilling prophecy. There is also the worry that the PP will be incoherent in the SRM worst case I am describing, recommending both for and against the implementation of SRM. But, as I argued in section 4.4, it is not necessarily irrational to prevent a catastrophe from occurring, while at the same time introducing new risks. Much would depend on our knowledge of the situation. If we knew the impending climate catastrophe was huge and highly likely, but only knew that a SRM induced catastrophe

was not impossible, but had no idea of its nature, then perhaps it would be rational to implement SRM. I am not claiming, as Sunstein does, that we would need a demonstrable probability of the impending catastrophe, but if there is enough scientific evidence and understanding of the situation that can justify qualitative descriptions such as "highly likely," then that might suffice to justify a risky precautionary measure in the form of SRM.

The SRM case is admittedly highly controversial, and a thorough discussion of debate cannot be given here, but the point that there can be worst case scenarios where our only chance to prevent an impending catastrophe is to implement a highly risky precautionary measure. But if we adopt the condition that a precautionary measure itself cannot introduce new threats of catastrophe (proportionality condition) we will almost never be allowed to implement such measures, for it will be very hard to prove that they will not introduce new threats of catastrophe. The discussion suggests that we should recognize situations in which the proportionality condition can be defeated. Or, alternatively, it could be stated that in such situations the PP simply does not apply, because there is no "precautionary" action available; if all actions involve threats of catastrophe, we are beyond the scope of precautionary action and thus the precautionary principle (Hartzell-Nichols 2012, pp. 168-69). I am inclined to disagree with this position, however, because a precautionary approach involves recognizing all possible situations, especially the very worst cases, and giving ourselves the best chance to avoid catastrophes in worst-case scenarios. In some cases, the only way to do that would be to use risky technologies or alternatives as a last-ditch effort. And so, if such cases are plausible or likely to occur, then a precautionary approach likely calls for the development of the option to implement those alternatives.

6.7 Typology of catastrophe cases

I will now introduce a typology of catastrophe cases in order to help understand when highly costly or risky actions are justified and when they are not. This will be a useful conceptual tool for recognizing which features of cases justify which precautionary actions, and can also be used to identify the motivation behind different arguments for or against different activities and technologies. In "waterfall" cases, there is an impending catastrophic outcome. In such cases, implementing the PP does not automatically call for any action that might prevent the catastrophe from occurring, but once less risky options are no longer available (such as mitigating or adapting to climate change), riskier / potentially costly measures can be recommended by the PP. As I have suggested, in the limiting case, the PP can even recommend measures that themselves are not proven not to introduce new risks of catastrophe assuming no other options are available. The "dying patient" case is a waterfall case, as it is rational for a dying patient to try risky and unproven drugs or treatments to try to save themselves even if the treatments might kill them, whereas this would not be justified for a mildly sick patient.

In "Frankenstein" cases, there is no impending catastrophe that we are trying to prevent. Instead, cases of this type involve the introduction of risky activities and technologies in order to try to achieve great benefits. For example, a healthy person trying to decide whether to take a "super pill" that either turns them into superman or kills them. The real-life case that most accurately fits under this category, as I see it, is the creation of superhuman artificial intelligence. Creating the AI could lead to a development of an information revolution that would make the industrial revolution (and its benefits) pale in comparison, but it may also end up destroying us because of our inability to control it or, say, if we programmed it improperly. Alternatives that

carry risks of catastrophe are far less rational in Frankenstein cases and the PP is more likely to recommend against them, or for their stringent regulation. That's because the possible downsides they introduce are far too costly and unnecessary.[27]

Accurately identifying decision-problems using this typology will help to determine what kinds of measures the precautionary principle will recommend. Often, proponents of risky activities and technologies frame the decision-problem as a waterfall case. For instance, proponents of genetically modified foods often argue for them as a precautionary measure against the global starvation crisis (thus framing GMOs as a waterfall case) (Taleb et al., p. 10). Opponents of GMOs point out that starvation can be solved through government policy changes in distributions and uses of the current food stock, and, furthermore, that starvation, while terrible, is a relatively local and contained harm compared to the global systemic ecological and health risk introduced by GMOs (Thus framing GMOs as a Frankenstein case) (Ibid.). Similarly, geoengineering can be framed as a Frankenstein case. If we had a relatively healthy though suboptimal climate and we wanted to engineer it to be even more favourable for human goals, geoengineering seems completely unjustified, as we would be introducing systemic risk in order to improve something that's already pretty good. Correctly identifying the decision problem as a waterfall or Frankenstein case can help to determine when costly/risky precautionary measures are justified and when they are not and to better understand the arguments of proponents and opponents of emerging technologies.

---

[27] I thank Allen Habib for recommending the Waterfall/Frankenstein terminology in conversation.

Conclusion

To conclude, I have argued that the main objection to the Precautionary Principle – the incoherence argument – can be defended against, and that there are significant limitations to quantitative approaches to decision-making in the context of environmental policy. I take this to show that there is a need for the Precautionary Principle, as a qualitative decision rule, in environmental policy decision making. I also examined Gardiner's maximin interpretation of the PP, one of the main interpretations of the PP in the literature, and considered three arguments against it. The first was that maximin requires that the full range of the possible outcomes of our actions be known in order to be applied, but many environmental policy decision problems are such that there is significant uncertainty about the full range of possible outcomes, i.e., are decisions under ignorance. This point also suggests that there are significant limitations to formal decision theory in environmental policy in so far as formal decision principles require a complete list of possible outcomes. The second was that maximin requires that precautionary measures be guaranteed to prevent threats of harm in order to be enacted, but this is often unknown as well. The third was that Gardiner's maximin principle only applies in cases where the costs of implementing precautionary measures are relatively low, and where decision-makers do not care much about the foregone benefits of regulation, thus leaving out many cases where precautions are likely to be moderate or high cost. I argued that precautionary measures can be justified in each of these three cases, but that Gardiner's maximin principle will be unable to recommend them. I thus conclude that the PP should not be interpreted as a version of the maximin decision rule, despite the initial plausibility of that project. Any plausible formulation of the PP needs to be applicable in situations in which there is uncertainty about the range of possible outcomes, about the effectiveness of precautionary measures, and where precautionary measures are

potentially high cost. Additionally, I argued that a common condition on how costly precautionary measures can be in order to prevent threats of harm – the proportionality condition – should be recognized as defeasible in certain kinds of cases involving catastrophic possible outcomes. Finally, I provided a typology of catastrophe cases to help identify situations where risky precautionary measures are more likely to be justified and when they are not.

References

Ackerman, F. & L. Heinzerling. (2002). "Pricing the Priceless: Cost-Benefit Analysis of Environmental Protection," *University of Pennsylvania Law Review* 150: 1553-1584.

Aldred, J. (2013). "Justifying Precautionary Policies: Incommensurability and Uncertainty," *Ecological Economics* 96: 132-140.

Aldred, J. (2012). "Climate Change Uncertainty, Irreversibility and the Precautionary Principle," *Cambridge Journal of Economics* 36: 1051-1072.

Andreae, M., C. Jones, and P. Cox. (2005). "Strong Present-Day Aerosol Cooling Implies a Hot Future," *Nature* 435: 11187-1190.

Arrow, K. & L. Hurwicz. (1972). "An Optimality Criterion for Decision Making Under Ignorance," in Carter, C. and Ford, J. (eds.), *Uncertainty and Expectations in Economics: Essays in Honour of G. L. S. Shackle*, Oxford, Blackwell.

Aven, T. (2011). "On Different Types of Uncertainties in the Context of the Precautionary Principle," *Risk Analysis* 31: 1515-1525.

Aven, T. (2003). *Foundations of Risk Analysis: A Knowledge and Decision-Oriented Perspective*, Chichester, John Wiley & Sons Ltd.

Betz, G. (2010). "What's the Worst Case? The Methodology of Possibilistic Prediction," *Analyse & Kritik*: 87-106.

Betz, G. (2007). "Probabilities in Climate Policy Advice: A Critical Comment," *Climatic Change* 85: 1-9.

Bodansky, D. (1991). "Scientific Uncertainty and the Precautionary Principle," *Law, Environment: Science and Policy for Sustainable Development* 33: 4-44.

Bognar, G. (2011). “Can the Maximin Principle Serve as a Basis for Climate Change Policy?” *The Monist*: 94: 329-348.

Boyer-Kassem, T. (2017a). “The Precautionary Principle Has Not Been Shown to Be Incoherent: A Reply to Peterson,” *Risk Analysis* 37: 2039-2040.

Boyer-Kassem, T. (2017b). “Is the Precautionary Principle Really Incoherent?” *Risk Analysis* 37: 2026-2034.

Cameron, J. & J. Abouchar. (1991). “The Precautionary Principle: A Fundamental Principle of Law and Policy for the Protection of the Global Environment,” *Boston College International and Comparative Law Review* 14: 1-27.

Canadian Environmental Law Association. (2017). *The Precautionary Principle*. Retrieved from http://www.cela.ca/collections/pollution/precautionary-principle.

Clarke, H. (2008). “Classical Decision Rules and Adaptation to Climate Change,” *The Australian Journal of Agricultural and Resource Economics* 52: 487-504.

Crutzen, P. (2006). “Albedo Enhancement by Stratospheric Sulphur Injections: A Contribution to Resolve a Policy Dilemma?” *Climatic Change* 77: 211-219.

Environment and Climate Change Canada. (2017). *Precaution in Science-Based Decision Making*. Retrieved from https://www.ec.gc.ca/lcpe-cepa/default.asp

Farber, D. (2015). “Coping with Uncertainty: Cost-Benefit Analysis, the Precautionary Principle, and Climate Change,” *Washington Law Review* 90: 1659-1725.

Gardiner, S. (2010). “Is “Arming the Future” with Geoengineering Really the Lesser Evil?” In S. Gardiner, S. Caney, D. Jamieson, and H. Shue (eds.), 284-312, *Climate Ethics: Essential Readings*, Oxford, Oxford University Press.

Gardiner, S. (2006). "A Core Precautionary Principle," *The Journal of Political Philosophy* 14: 33-60.

Gilboa, I. & D. Schmeidler. (1989). "Maximin Expected Utility with a Non-Unique Prior," *Journal of Mathematical Economics* 18: 141-153.

Graham, J. & J. Wiener. (2008). "The Precautionary Principle and Risk-Risk Tradeoffs: A Comment," *Journal of Risk Research* 11: 465-474.

Hansen, S., M. Krayer, and J. Tickner. (2008). "The Precautionary Principle and Risk-Risk Tradeoffs," *Journal of Risk Research* 11: 423-464.

Hansson, S. & T. Aven. (2014). "Is Risk Analysis Scientific?" *Risk Analysis* 34: 1173-1183.

Hansson, S. (2009). "From the Casino to the Jungle: Dealing with Uncertainty in Technological Risk Management," *Synthese* 168: 423-432.

Hansson, S. (2006). "Economic (Ir)Rationality in Risk Analysis," *Economics and Philosophy* 22: 231-241.

Hansson, S. (2002). "Philosophical Perspectives on Risk," Keynote address at the conference of *Research in Ethics and Engineering*, Delft University of Technology, April 25, 2002:1-32.

Hansson, S. (1999). "Adjusting Scientific Practices to the Precautionary Principle," *Human and Ecological Risk Assessment: An International Journal* 5: 909-921.

Hansson, S. (1997). "The Limits of Precaution," *Foundations of Science* 2: 293-306.

Hansson, S. (1996). "Decision Making Under Great Uncertainty," *Philosophy of the Social Sciences* 26: 369-386.

Hansson, S. (1994). *Decision Theory: A Brief Introduction*, Department of Philosophy and the History of Technology, Royal Institute of Technology (KTH), Stockholm.

Hansson, S. (1993). "The False Promises of Risk Analysis," *Ratio* 6: 16-26.

Harremoës, P., D. Gee, M. MacGarvin, A. Stirling, J. Keys, B. Wynne, and S. Guedes Vaz. (2002). *The Precautionary Principle in the 20th Century: Late Lessons from Early Warnings*, London, Earthscan Publications.

Harsanyi, J. (1975). "Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls's Theory," *The American Political Science Review* 69: 594-606.

Hartzell-Nichols, L. (2013). "From 'the' Precautionary Principle to Precautionary Principles," *Policy & Environment* 16: 308-320.

Hartzell-Nichols, L. (2012). "Precaution and Solar Radiation Management," *Ethics, Policy & Environment* 15: 158-171.

Hartzell, L. (2009). "Rethinking the Precautionary Principle and its Role in Climate Change Policy," PhD dissertation, Stanford University.

Jamieson, D. (1996). "Ethics and Intentional Climate Change," *Climatic Change* 33: 323-336.

Jordan, A. & T. O'Riordan. (1999). "The Precautionary Principle in Contemporary Environmental Policy and Politics," In C. Raffensberger and J. Tickner (eds.), *Protecting Public Health and the Environment: Implementing the Precautionary Principle*, Washington, DC, Island Press, 15-35.

Jordan, A. & T. O'Riordan. (1995). "The Precautionary Principle in Contemporary Environmental Politics," *Environmental Values* 4: 191-212.

Katzav, J., H. Dijkstra, and A. de Laat. (2012). "Assessing Climate Model Projections: State of the Art and Philosophical Reflections," *Studies in History and Philosophy of Modern Physics* 43: 258-276.

Keith, D., E. Parson, and M. Morgan. (2010). "Research on Global Sun Block Needed Now," *Nature* 463: 426-427.

Keith, D. (2000). "Geoengineering the Climate: History and Prospect," *Annual Review of Energy and the Environment* 25: 245-84.

Knight, F. (1921). *Risk, Uncertainty, and Profit*. Boston, Houghton Mifflin Company.

Luce, R. & H. Raiffa. (1957). *Games and Decisions: Introduction and Critical Survey*, New York, John Wiley & Sons.

Manson, N. (2002). "Formulating the Precautionary Principle," *Environmental Ethics* 24: 263-274.

Marland, G., T. Boden, and R. Andreas. (2008). "Global CO2 Emissions from Fossil-Fuel Burning, Cement Manufacture, and Gas Flaring: 1751-2005," Carbon Dioxide Information Analysis Center, U.S. Department of Energy.

McKinnon, C. (2009). "Runaway Climate Change: A Justice-Based Case for Precautions," *Journal of Social Philosophy* 40: 187-203.

Neumayer, E. (2010). *Weak Versus Strong Sustainability: Exploring the Limits of Two Opposing Paradigms*, Cheltenham: Edward Elgar Publishing Limited.

Peterson, M. (2017a). "Yes, The Precautionary Principle is Incoherent," *Risk Analysis* 37: 2035-2038.

Peterson, M. (2017b). *An Introduction to Decision Theory* (2nd ed.), Cambridge, Cambridge University Press.

Peterson, M. (2007). "The Precautionary Principle Should Not be Used as a Basis for Decision-making," *European Molecular Biology Organization (EMBO Reports)* 8: 305-308.

Peterson, M. (2006). "The Precautionary Principle is Incoherent," *Risk Analysis* 26: 595-601.

Pindyck, R. (2017). "The Use and Misuse of Models for Climate Policy," *Review of Environmental Economics and Policy* 11: 100-114.

Raffensberger, C. & J. Tickner, eds. (1999). *Protecting Public Health and the Environment: Implementing the Precautionary Principle*, Washington, D.C., Island Press.

Rawls, J. (1999). *A Theory of Justice*, Revised ed. Cambridge, Mass., Harvard University Press.

Rawls, J. (1974). "Some Reasons for the Maximin Criterion," *The American Economic Review* 64: 141-146.

Resnik, D. (2003). "Is the Precautionary Principle Unscientific?" *Studies in History and Philosophy of Biological and Biomedical Sciences* 34: 329-344.

Sandin, P. (1999). "Dimensions of the Precautionary Principle," *Human and Ecological Risk Assessment: An International Journal* 5: 889-907.

Science and Environmental Health Network. (2018). *The Wingspread Consensus Statement on the Precautionary Principle*. Retrieved from http://sehn.org/wingspread-conference-on-the-precautionary-principle/.

Scott, K. (2013). "International Law in the Anthropocene: Responding to the Geoengineering Challenge," *Michigan Journal of International Law* 34: 309-358.

Soule, E. (2000). "Assessing the Precautionary Principle," *Public Affairs Quarterly* 14: 309-328.

Sprenger, J. (2012). "Environmental Risk Analysis: Robustness Is Essential for Precaution," *Philosophy of Science* 79: 881-892.

Stainforth, D. et al. (2005). "Uncertainty in Predictions of the Climate Response to Rising Levels of Greenhouse Gases," *Nature* 433: 403-406.

Steel, D. (2014). *Philosophy and the Precautionary Principle: Science, Evidence, and Environmental Policy*, Cambridge, Cambridge University Press.

Stirling, A. (2008). "Science, Precaution, and the Politics of Technological Risk: Converging Implications in Evolutionary and Social Scientific Perspectives," *Annals of the New York Academy of Sciences* 28: 95-110.

Sunstein, C. (2005). *Laws of Fear: Beyond the Precautionary Principle*, Cambridge, Cambridge University Press.

Svobada, T. (2012). "Is Aerosol Geoengineering Ethically Preferable to Other Climate Change Strategies?" *Ethics and the Environment* 17: 111-135.

Taleb, N., R. Read, R. Douady, J. Norman, B. Yaneer. (2014). "The Precautionary Principle (with Application to the Genetic Modification of Organisms)," *Extreme Risk Initiative – NYU School of Engineering Working Paper Series*: 1-24.

Thalos, M. (2012). "Precaution Has Its Reasons," in W. Kabasenche, M. O'Rourke & M. Slater (eds.), *Topics in Contemporary Philosophy 9: The Environment, Philosophy, Science and Ethics*, MIT Press, 171-184.

Thalos, M. (2009). "There is No Core to Precaution," *Review Journal of Political Philosophy* 7: 41-49.

Trouwborst, A. (2006). *Precautionary Rights and Duties of States*, Leiden: Martinus Nijhoff.

Turner, D. & L. Hartzell. (2004). "The Lack of Clarity in the Precautionary Principle," *Environmental Values* 13: 449-460.