

CAN THE MAXIMIN PRINCIPLE SERVE AS A BASIS FOR CLIMATE CHANGE POLICY?

1. Introduction

Climate change can potentially cause massive harm and perhaps even a global catastrophe. It is estimated that if no action is taken, CO₂^e (carbon dioxide equivalent) concentrations may increase from 430 parts per million (ppm) today to 750ppm by the end of the century. Even if they are stabilized there, there is a 50% chance of a temperature increase in excess of 5°C. In this case, enough of the planet's ice may melt for sea levels to eventually rise by 10 meters or more.¹ But even if we manage to avoid catastrophe on a global scale, smaller temperature increases are still going to cause local disasters. For instance, under most scenarios, the proportion of arid and semi-arid land in Africa is forecast to increase by 5–8% by 2080 (IPCC 2007, 448). This will directly harm agriculture, but its indirect effects may be no less bad. A large part of the rural population will have to abandon agriculture and move to cities where they are more likely to engage in unsafe sex or become sex workers. This will increase the incidence of HIV infection (McMichael *et al.* 2008).

It is sensible to try to avoid the worst scenarios associated with climate change, even if we do not know precisely how likely they are. It is warranted to take precautions. In environmental policy, this idea is reflected in the *precautionary principle*. Although the principle has no widely accepted formulation, its different versions express the thought that in cases in which serious harms are possible, policies and regulations should be designed with the aim of avoiding the worst outcomes even if we lack precise information about the magnitude of the harms, the probability with which they might obtain, or the causal relations involved in bringing them about. For instance, the United Nations Framework Convention on Climate Change (1992) declares:

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.

The UN Framework Convention urges the adoption of a precautionary approach to climate change, but it does not attempt to formulate a precautionary *principle* to capture that approach. As I said, the precautionary principle has no widely accepted formulation. It can be given weaker and stronger forms depending on how bad the adverse effects have to be to trigger a precautionary response, what degree of scientific uncertainty makes precautionary measures necessary or reasonable, what sort of measures are called for, and so on. Each formulation sets a level of certainty and a threshold of harm that justify forgoing precautionary responses and allow us to use the most cost-effective means to aim for the best outcomes.²

In recent years, however, the precautionary principle has come under attack. Critics have complained that its weaker formulations are vacuous, while its stronger, more robust formulations lack adequate justification and lead to inconsistent and irrational policy choices. In response, defenders of the precautionary approach looked for a justification that allows for a reasonably robust formulation but restricts the application of the principle to the range of cases where a precautionary approach seems the most warranted. The strategy is to justify and reduce the scope of the precautionary principle at the same time. And the defenders claim that they have found a way to do that by using Rawls' defense of the maximin decision-making rule.

Their argument, in rough outline, runs as follows. The maximin rule tells you to evaluate each alternative action by looking at the worst outcome that might obtain if you choose that action, and choose the action that has the best worst outcome. That is, you maximize the value of the minimum outcome that might obtain. Now as a general rule, maximin is irrational. But Rawls did not propose it as a general criterion for decision making; rather, he proposed it for the choice in the original position—his interpretation of the hypothetical social contract in which free and equal

individuals select a set of principles for regulating fundamental social institutions. The original position is characterized by certain highly unusual conditions. Given these conditions, following maximin is not irrational.

Moreover, these conditions turn out to be not so unusual after all. They also characterize the situation in which we find ourselves with respect to certain issues in environmental policy. Hence Rawls's defense of the maximin rule in his hypothetical original position is applicable to some real-life policy issues. Precautionary measures are warranted whenever the conditions that Rawls described are approximated. The maximin rule can serve as the basis for a narrow version of the precautionary principle that applies only in a limited range of important cases. Climate change is one of them.

Evidently, the Achilles' heel of this strategy is the application of Rawls's defense of the maximin rule. So in the first part of this paper I exploit that vulnerability. I begin by raising some doubts about the claim that the Rawlsian conditions characterize our situation when it comes to climate change. This is a problem for the strategy of formulating the precautionary principle in terms of the maximin rule, since climate change seems a paradigmatic case when a precautionary approach is reasonable. But this is not my main argument. 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.

Since there are cases when the precautionary approach is widely considered reasonable, this result is hard to accept. But I argue that there is another principle that can capture the precautionary approach—a principle familiar from the moral view known as *prioritarianism*. I argue that a prioritarian principle can be used to make more precise what the precautionary approach involves and serve a useful role in climate change policy. I develop this proposal in the second part of the paper.

2. *The Maximin Rule*

Rawls (1999) famously argued that his principles of justice are the “maximin solution” to the choice problem presented in the original position. A crucial feature of the original position is the veil of ignorance:

the parties do not know (among other things) what social and economic position they themselves will occupy once the veil is lifted. They have to make their choice in conditions of uncertainty.

One of Rawls's most important aims in *A Theory of Justice* is to provide an alternative to utilitarianism. Some utilitarians argue that the rational choice in the original position is to use the *principle of insufficient reason* and maximize expected utility. The principle of insufficient reason tells you that in complete uncertainty, you should treat each outcome as equally probable. That is, if there are n possible outcomes, the probability that you should assign to each outcome is $1/n$. Using the principle together with expected utility maximization is equivalent to maximizing average utility. At the price of some imprecision, I shall call this procedure the *utilitarian principle*.

This is imprecise because utilitarians are not required to use the principle of insufficient reason for handling uncertainty. They are free to combine their moral view with some other principle—even one that assigns a probability of 1 to the worst outcome. Nevertheless, Rawls and the defenders of maximin versions of the precautionary principle take their opponents to hold the combination of the principle of insufficient reason and expected utility maximization. Since that's quite a mouthful, perhaps in this context it's not too misleading to shorten it to "utilitarian principle."

A well-known objection to Rawls's view is that following the maximin rule is irrational. John Harsanyi illustrates this objection through a famous example:

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. (Harsanyi 1975, 595)

It seems clear that making this choice on the basis of the maximin rule is indeed irrational. The rule is insensitive to the probability of the worst outcome and the value of the other outcomes. It tells you to avoid the worst outcome even if it's extremely unlikely and the other outcomes are good. Plainly, in this case the rational choice is to maximize expected utility. You can do that by getting on the plane.

But it's not altogether clear what this example is supposed to show. The decision makers in the original position are behind the veil of ignorance. They do not know the probabilities of different outcomes; they have to make their choice in conditions of uncertainty. Harsanyi's example, however, is *not* a choice under uncertainty: staying in New York City is irrational because you know the probability of being killed in a plane accident is minimal (Harsanyi himself tells you this). So the example cannot show that the adoption of the maximin rule is not rational in the original position. In fact, even though maximin has been proposed as the appropriate rule in conditions of uncertainty, Rawls is not even committed to this proposal. As he says, "the maximin rule is not, in general, a suitable guide for choices under uncertainty" (1999, 133). He argues that following maximin is rational *only* in circumstances which are marked by some unusual conditions, and that these conditions characterize the original position.

Rawls enumerates three conditions that make the adoption of maximin rational in the original position. The first condition is uncertainty, which is ensured by his use of the veil of ignorance. The parties have no basis for assigning probabilities to possible outcomes. In particular, they do not know their own position in society. If they knew with high probability that they would not end up in a bad social position, it would be rational to use a decision-making rule that selects an alternative with a worse minimum but overall better possible outcomes.

The second condition is that there is no special reason for trying to obtain more than the minimum. As Rawls formulates the condition, "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" (1999, 134). Again, if the parties wanted to secure more than the minimum, it might be rational to use a decision-making rule that selects an alternative with a worse minimum but overall better possible outcomes. But this is not what they care about in the original position.

Finally, to make maximin rational, it must be the case that the alternatives that maximin rejects have very bad possible outcomes. If the parties were to choose on the basis of some other principle, they might end up with one of these outcomes. Maximin guarantees they avoid them.

Rawls emphasizes that the use of the maximin rule is rational only if all three conditions characterize the choice situation "to the highest

degree.” What is unique about the original position is that all of the conditions do.

3. *The Core Precautionary Principle*

Stephen Gardiner (2006) argues that Rawls’s conditions provide a way to develop the precautionary approach. The conditions can be treated as criteria for the application of the precautionary principle. His version of the principle is formulated on the basis of the maximin rule. It says that whenever the Rawlsian conditions are approximated, it is rational to choose the alternative that has the best worst outcome regardless of its expected utility. But since the conditions seldom obtain, maximin will be used only in rare, “core” cases. Gardiner calls this formulation the *core precautionary principle*.

Among the core cases are some environmental and public health regulatory choices—and, in particular, climate change policy. With respect to climate change, our situation is characterized by the Rawlsian conditions. First, 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.³ Second, stabilizing CO₂^e concentrations at an acceptable level (usually taken to be no more than 550ppm) is estimated to require around 1% of world GDP per year.⁴ Thus, it is appropriate to care very little for what we have to forgo, given the relatively low necessary expenditure and the risks and harms that can be avoided. And third, the impact of climate change will be severe and potentially catastrophic. If concentrations end up around 750ppm by the end of the century, the resulting temperature increase may be in excess of 5°C, leading to a rise in sea levels of 10 meters or more. No one denies that would be a global catastrophe.

The defense of the core precautionary principle rests on the Rawlsian conditions. These conditions simultaneously reduce its scope. In the case of climate change, the conditions are approximated. It is a core case. Hence, it is rational to follow the maximin rule in climate change policy.

There are at least two ways to attack this position. You might argue that the Rawlsian conditions do not actually characterize our situation with respect to climate change. So the core precautionary principle does not apply. This would be a significant problem, since climate change seems to be a paradigmatic case when a precautionary approach is warranted.⁵

But you might also argue that there is a deeper problem. It is this: Rawls's defense of the maximin rule in terms of his conditions does not succeed. There are other decision-making rules that are just as rational to use when the conditions apply. So the core precautionary principle is not uniquely rational in such cases. We are better off using some other rule, especially if it is rational to follow whether or not the conditions apply. Once again, this would be a significant problem if a precautionary approach were warranted in climate change policy. It would make the precautionary principle dispensable.

I will make only one point in connection with the first line of attack. It has to do with the first condition—uncertainty. As Gardiner himself points out, we are not really in a position of complete uncertainty with regard to climate change. Rather, we are in a combination of risk and uncertainty. After all, we can assign probability estimates—or at least ranges of estimates—to possible outcomes under alternative scenarios. The Intergovernmental Panel on Climate Change provides such probability estimates in its assessment reports. Technically, this makes our situation more like one of risk, as opposed to uncertainty, when such estimates are not available.⁶

Perhaps these estimates are not sufficiently reliable to take us from uncertainty to risk. Moreover, many possibilities remain for which there are no probability estimates at all. Thus, even though we are not under complete uncertainty, it's not unreasonable to hold that our situation sufficiently approximates the Rawlsian uncertainty condition. So the core precautionary principle applies.

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?

As time passes, scientists learn more about the effects of climate change. They can build better models. Presumably, their probability estimates become more reliable. Unless we have widely overestimated the probabilities of bad outcomes, it seems to get things the wrong way around to hold that the precautionary approach becomes *less* rational as we learn more. 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.

For illustration, consider the example Gardiner uses to introduce the maximin rule:

Suppose that in a given situation you have two actions, A and B, available to you. If you choose A, then there are two possible outcomes: either (A1) you will receive \$100, or (A2) you will be shot. If you choose B, there are also two possible outcomes: either (B1) you will receive \$50, or (B2) you will receive a slap on the wrist. (2006, 45–46)

Maximin tells you to choose B, which has the best worst outcome. So it gives the right answer in the example. As Gardiner puts it, “we seem both to behave in accordance with maximin thinking in important areas of decision-making in real life and also to reflectively endorse such behavior as paradigmatically rational” (2006, 46). If, by “maximin thinking,” he means focusing on avoiding the worst outcome without taking account of its probability, then our behavior does not become less “paradigmatically rational” as we continue to choose B even as we learn more about the probabilities of being shot and receiving a slap on the wrist. I suspect the probability of being shot has to be extraordinarily low for anyone to be tempted by A.

So much for the first line of attack. The second line of attack, targeting maximin directly, presents a more formidable problem for the core precautionary principle. Gardiner’s example can be used as a starting point for developing it. Recall that a rival decision-making rule for uncertain choices is what I called the utilitarian principle: assigning equal probabilities to all of the possible outcomes and maximizing expected utility. Utilitarians like Harsanyi argue that this is the rational strategy to follow in uncertainty. Rawls disagrees. His conditions are intended to guarantee that following maximin is the only rational strategy.

But if you look at Gardiner’s example, it is noticeable that there is no disagreement about the rational choice between the two sides. The utilitarian principle tells you to choose B as well. When you follow this rule, you assign an equal chance to receiving \$100 and being shot, and an equal chance to receiving \$50 and a slap on the wrist. No one would have any doubts what to do if they made their choice this way. Hence the example does not show that “maximin thinking” is paradigmatically rational. It could equally well be employed to show that it is the utilitarian principle that is paradigmatically rational.

Putting the point more generally, the reason for the agreement has to do with Rawls's third condition. According to that condition, there are some alternatives that have very bad outcomes. This condition applies to the example: being shot is a very bad outcome. The condition also translates into our position with respect to climate change, since there are scenarios that could lead to catastrophe. But because of the presence of these very bad outcomes, maximin is not unique in selecting alternatives that avoid them. If they are truly awful, a number of other principles—including the utilitarian principle—will do that too.⁷

It is worth spelling out in more detail why the maximin rule and the utilitarian principle lead to similar choices in circumstances characterized by the Rawlsian conditions. One way they might lead to different choices is if the worst outcomes of the alternatives that maximin rejects are only slightly worse than the worst outcome of the alternative that maximin selects. That is, the best worst outcome is only slightly better than the other worst outcomes. In this case, the rejected alternatives might have better outcomes that "compensate" for their worst outcome such that the utilitarian principle selects one of them. However, the third condition seems to exclude this possibility. As Rawls puts it, "the rejected alternatives have outcomes that one can hardly accept" (1999, 134). But if the worst outcomes of the rejected alternatives are only slightly worse, either they are not hardly acceptable or the worst outcome of the alternative that maximin selects is also hardly acceptable.

Another way the maximin rule and the utilitarian principle may lead to different choices is if there is an alternative that maximin rejects that has some exceptionally good outcomes compared to the alternative that maximin selects. Again, these very good outcomes might compensate for having hardly acceptable outcomes. But this is inconsistent with the second condition: that the decision maker cares very little for what she can gain above the minimum. She would therefore sharply discount these better outcomes.⁸

Recall the way Rawls formulates the second condition: "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" (1999, 134). This condition is strange, even in the context of Rawls's own theory. For one thing, he elsewhere assumes that the parties in the original position do not

know their own conception of the good. But then why would they think that they have a conception of the good so that they are satisfied with the minimum? Rawls also assumes that the parties in the original position aim to maximize their bundle of social primary goods. But then why would they not care about what they can gain above the minimum?⁹

In any event, Gardiner says very little about this condition. He claims that in the case of climate change it is met because the costs of reducing emissions to acceptable levels are relatively low. This suggests that the *best* outcomes of the alternatives of responding to climate change and not doing anything are roughly equally good. But if this is so, it only seems to provide further support for the utilitarian principle. For if there is little difference with respect to the best outcomes of responding to climate change and not doing anything, but there are large differences with respect to the worst outcomes of these policies, then the utilitarian principle would equally select the alternative that has the best worst outcome.

Peculiarly, Gardiner thinks the congruence between the maximin rule and the utilitarian principle when the Rawlsian conditions apply provides further support for maximin. He suggests that the fact that utilitarians would make the same choices in these cases shows that following maximin can be reasonable and useful. In order to defend the core precautionary principle, it is sufficient to show that following maximin is rational when the Rawlsian conditions apply. There is no need to show that following it is *uniquely* rational. If other principles lead to the same choices, the case for maximin is strengthened, not undermined.

This is a puzzling argument. For suppose you have two decision-making rules, U and P. You know that whenever using P leads to a reasonable choice, using U leads to the same choice. But you also know that sometimes using P leads to an unreasonable choice, whereas using U always leads to a reasonable choice. Clearly, ever using P is irrational. For you know that you will always make a reasonable choice by using U, but when you want to use P, you first have to figure out whether you are in the sort of situation when it leads to a reasonable choice. That's surely a waste of time. You could just use U instead.

By the same token, it is not only that following maximin is not uniquely rational when the Rawlsian conditions apply. It is actually irrational. Since it is irrational when the conditions don't apply, you are better off not using it at all, since then you don't have to worry about finding out

whether the conditions apply. Whenever the Rawlsian conditions are approximated, the core precautionary principle offers no advantage over utilitarianism.

This conclusion is surprising. It means that if it is defended in terms of the Rawlsian conditions, the precautionary principle is dispensable.

4. *The Anti-Catastrophe Principle*

There are many formulations of the precautionary principle that would apply only when we are potentially faced with a catastrophe. At least one of these—Cass Sunstein’s *anti-catastrophe principle*—has been defended in terms of the Rawlsian conditions. More precisely, Sunstein accepts two of the three conditions: uncertainty and the possibility of catastrophic outcomes. Whenever there is a realistic but uncertain possibility of a catastrophic outcome, policy making should follow the maximin rule (Sunstein 2005).

To begin, there is a worry about anti-catastrophe principles. What should be considered a catastrophe? After all, “catastrophe” is not a value-neutral term. No one would deny that a rise in sea levels by 10 meters would be a catastrophe. The anti-catastrophe principle evidently applies to this case. But what about a one-meter rise? Would that be a catastrophe? For the inhabitants of small island states, certainly—after all, their countries would be wiped out. But relatively few people would be directly affected, avoiding a one-meter rise might be exceedingly costly, and for the rest of the world it might be more cost-effective to adapt. The notion of catastrophe suggests that it is clear when the principle comes into play. But this is not determinate at all.

Perhaps the Rawlsian conditions can help make the scope of the anti-catastrophe principle more determinate. But Sunstein rejects the second condition. Recall that on Gardiner’s interpretation, this condition is met whenever the gains that have to be forgone by following maximin are relatively small. As Sunstein points out, this makes the precautionary approach trivial. After all, if the opportunity costs of doing so are low, there is no reason not to take precautions. Hence, he suggests, this condition should be abandoned.

I agree with Sunstein that Gardiner’s interpretation of the second Rawlsian condition is problematic. But the problem is not merely that it makes the precautionary approach trivial. Rather, it makes it dispensable. Whenever there are alternatives whose worst outcomes are catastrophic

and their best outcomes are only slightly better than the best outcomes of alternatives with acceptable worst outcomes, it is rational to ensure that you avoid the catastrophic outcomes. But you don't need maximin for that. The utilitarian principle also selects alternatives with acceptable worst outcomes.

Moreover, there is another problem with dropping the second condition. It leaves unclear how we should respond to cases when the best outcomes of the alternatives whose worst outcomes are catastrophic are much better than the best outcomes of the alternatives whose worst outcomes are acceptable.

Sunstein gives the following example. Suppose the possibility of catastrophic climate change can be eliminated only if living standards all over the world are reduced by 50 percent. Apparently, he thinks that this is such a great sacrifice that it can make it worth taking the risk of not responding, hoping that no catastrophic climate change will occur. He argues that following maximin may not be reasonable in this case. So the anti-catastrophe principle does not apply, even though there remains the possibility of a catastrophe. His solution for cases like this is to implicitly abandon the *first* Rawlsian condition: "to incur costs of this magnitude, we might want to insist that the danger of catastrophe rise above a minimal threshold—that there be a demonstrable probability, and a not-so-low one, that the catastrophic risk will occur" (2005, 113).

Once again, abandoning the first condition is not unreasonable. But that's because it restricts the application of the precautionary principle to uncertainty. It should not be replaced by a condition that restricts its application to risk. The precautionary principle should apply to both. Sunstein's maneuver makes the anti-catastrophe principle inapplicable in uncertainty. But precautionary measures appear especially warranted in such circumstances. If there is a plausible but uncertain possibility of a catastrophe, it makes sense to take precautions to avoid it.

Furthermore, dropping both the first and the second conditions helps in no way whatsoever justify the precautionary principle on the basis of the maximin rule. It leaves it wide open to the sort of objection that I quoted Harsanyi make in section 2. Even Rawls would agree that the presence of very bad outcomes in itself cannot make following the maximin rule rational. If that is all that Sunstein offers as its justification, the anti-catastrophe principle remains irrational.

I have been arguing that the precautionary principle cannot be defended on the basis of the maximin rule with the help of the Rawlsian

conditions. A precautionary approach certainly seems reasonable when you are faced with uncertainty, but uncertainty is not necessary for precautionary measures to be warranted. The presence of very bad outcomes certainly seems necessary for precautionary measures to be warranted, but the precautionary principle is either not the only principle that leads to reasonable choices in such cases, or it is irrational. So the first condition is not necessary and the third condition is not sufficient for defending the precautionary principle. As for the second condition: if it is accepted, it makes the precautionary principle dispensable; if it is rejected, it makes it irrational.

5. Salvaging the Precautionary Approach

No doubt many people would find worrisome the conclusion that the precautionary principle is dispensable. After all, adopting a precautionary approach in certain circumstances seems eminently reasonable. Some people may even turn my argument around: they might agree that when the Rawlsian conditions are approximated to a high degree, the utilitarian principle delivers a reasonably “precautionary” result. It is safe to use in these circumstances. But they might worry that if we adopted the utilitarian principle, we would make choices in other circumstances that are unsatisfactorily precautionary. In this regard, they reject the core precautionary principle and the anti-catastrophe principle as well. They consider them too narrow.

In particular, many people might believe that when it comes to climate change, there are risks that we should take great care to avoid even if the alternatives that contain these risks also have outcomes that are sufficiently good to compensate for them. The precautionary approach, in their view, guarantees a reasonable degree of risk aversion. The utilitarian principle, they argue, is inadequately sensitive to the badness of these risks.

One way to interpret this worry is in terms of the second Rawlsian condition. This condition is met when you care very little for what you might gain above the minimum. Gardiner believes the condition is satisfied when it comes to climate change because the costs of responding are relatively low. But an alternative interpretation—which, incidentally, seems to me closer to what Rawls might have intended—is that the condition is normative: you *ought to* care little for what you can gain above the minimum, even if the sacrifice that needs to be made is considerable. You should give more weight to avoiding the worst outcome.

Interpreted this way, the second condition can be regarded as an expression of part of the moral view known as *prioritarianism*. Prioritarianism

is the view that benefiting a person matters more the worse off that person is. Utilitarians hold that a benefit has the same moral weight no matter how well off the beneficiary is; prioritarrians hold that the moral importance of a benefit depends, in addition to its size, also on how well off the beneficiary is. If two people can derive the same net benefit from your aid, but one of them is worse off than the other, then the benefit you can bestow on the worse off person has greater moral importance—and the worse off she is, the greater its importance.¹⁰

Although prioritarianism is usually formulated as a view about the distribution of benefits among different people, it has a formal connection to risk aversion.¹¹ Whenever risk is present, giving more weight to the worse off is equivalent to giving more weight to avoiding the worst outcome. So giving more weight to avoiding the worst outcome is part of the prioritarian view. As I am interpreting it, the second Rawlsian condition reflects this idea. For giving less weight to what you can gain “above the minimum” is equivalent to giving more weight to avoiding what you can lose below the minimum.

For illustration, I give an example. Although greatly simplified, it’s not completely unrealistic: the numbers I use roughly correspond to those which can be found in various scenarios outlined in reports on the impact of climate change. Moreover, I’ve already introduced half of the example. Suppose we don’t do anything to reduce emissions. In this case, CO₂^e concentrations may stabilize at around 750ppm by the end of the century, increasing from the current level of approximately 430ppm. Under this “business as usual” scenario, there is a 50% chance that the climate will become warmer by more than 5°C. This would precipitate the melting of enough of the planet’s ice for at least a 10-meter rise in sea levels. This would be a global catastrophe. But there is also 50% chance that warming will not exceed 5°C. We could get lucky. Even though there would be harms caused by climate change, these harms would not be catastrophic. Under this scenario, we would also save the costs of mitigation.¹²

Suppose the alternative is pursuing aggressive mitigation strategies. At considerable cost, we could stabilize CO₂^e concentrations at around 500ppm by 2100. In this case, there is still a roughly 50% chance that the resulting increase in temperature would be in excess of 3°C. Climate change would still lead to substantial harms, regardless of our efforts. But with some luck, the increase in temperature would stay below 3°C.

Climate change would cause disruption, but our policies will have been largely successful in averting the worst harms, albeit at considerable costs.

In this example, we can choose between two alternatives (business as usual and mitigation), each of which may lead to two equally probable outcomes. The four outcomes, ranked from worst to best, are as follows:

- (1) warming by more than 5°C, but no costs;
- (2) warming by more than 3°C, and paying the costs;
- (3) warming not exceeding 3°C, but paying the costs;
- (4) warming not exceeding 5°C, and not paying the costs.

This ranking of the outcomes is not unrealistic. Warming by more than 5°C could lead to a global catastrophe. Even if we respond, warming may be in excess of 3°C with a substantial negative impact on human well-being. We would avoid the worst, but our costly efforts would be insufficient to avoid serious harm. It would be better if our efforts paid off—which means warming remains less than 3°C. Since climate change is already under way, this would still be harmful, but the harms would be less bad. Finally, warming may stay below 5°C even if we do very little to reduce emissions. Even though there would be substantial negative impacts on human well-being, we would be spared the costly efforts of mitigation and avoid catastrophe. We would not need to make sacrifices.

I will now, however, make an assumption that is completely unrealistic. I shall assume that the magnitude of the difference between the harms of (1) and (2) equals the magnitude of the difference between the harms of (3) and (4). That is, the difference between the worst outcomes of business as usual and mitigation is equal to the difference between the best outcomes of business as usual and mitigation. This assumption allows me to illustrate the difference between the utilitarian and the prioritarian views.

Given this assumption, utilitarians would be indifferent between business as usual and responding to climate change. The expected benefits of the two alternatives are equal. Business as usual has equally probable outcomes that are the best we can achieve and the worst that can happen. Mitigation has equally probable outcomes that are the second best we can achieve and the second worst that can happen.

Prioritarians, however, would respond differently. Since they give more weight to avoiding the worst outcome, they would prefer mitigation. The harms that can occur if we do nothing are given more weight than the possible benefits of business as usual. The harm between (1) and (2) is morally more important than the harm between (3) and (4). It is more important to avoid the loss from (2) to (1) than to secure the gain from (3) to (4). It is a risk that should not be taken. For this reason, many people may find prioritarianism more reasonable than utilitarianism when it comes to climate change. It seems able to capture the precautionary approach.¹³

I should explain why the assumption I made—that the difference between the harms of (1) and (2) equals the difference between the harms of (3) and (4)—is completely unrealistic. If concentrations stabilize at around 750ppm and warming exceeds 5°C, the harms of climate change are going to be *much worse* than the harms associated with unsuccessful mitigation. Whatever the details of the different scenarios, in practical terms there is not much of a choice between utilitarianism and prioritarianism when it comes to climate change. Utilitarians and prioritarians would both recommend mitigation. Nevertheless, the example is useful to show how the precautionary approach can be associated with the prioritarian view.

Prioritarianism can be used to make more precise what the precautionary approach involves. On the view I am exploring here, the precautionary principle is interpreted as a prioritarian principle. This view fits well with the idea that the precautionary principle can have many different versions. It allows for weaker and stronger formulations of the principle. For I have said nothing about *how much* priority should be given to avoiding worse outcomes. The more weight given, the more robust the precautionary response. The maximin version of the precautionary principle would give absolute weight to avoiding the worst outcome. But that would often lead to irrational choices. The prioritarian version is more flexible. It also accounts for the idea that a precautionary approach is reasonable in many cases.

Since maximin is irrational as a general decision-making rule, defenders of maximin versions of the precautionary principle have to work out the conditions under which the principle applies. Hence the importance of the Rawlsian conditions in this discussion. But the prioritarian version has no need for such conditions. The first Rawlsian condition is the presence of uncertainty. But the application of the prioritarian principle does not depend on whether uncertainty is present. In the example I used, for instance, the

alternatives had outcomes with equal probabilities. The probabilities were taken to express risks. But the recommendation of the prioritarian view would not change if we were in complete uncertainty and used the principle of insufficient reason to assign equal probabilities to the outcomes. Of course, just as utilitarians, prioritarians may favor some other account for dealing with uncertainty. Whatever that account may be, the application of their view does not depend on the presence of uncertainty.

The second Rawlsian condition is interpreted as an expression of part of the prioritarian view, rather than a criterion for its application. And the third condition is unnecessary: a prioritarian principle applies whether or not there are possible outcomes that are hardly acceptable or catastrophic.

6. Conclusion

The precautionary approach has been highly influential in discussions of environmental and climate change policy. It has proved difficult, however, to formulate the precautionary principle in a satisfactory way. The most promising strategy is to base the principle on the maximin decision-making rule and use the conditions Rawls offers for the justification of maximin as its criteria of application.

In these conditions, however, maximin is not the only rational decision-making rule. Whenever the Rawlsian conditions are approximated, at least one version of a utilitarian decision-making rule fares no worse than maximin. This makes the maximin versions of the precautionary principle dispensable. In practice, when it comes to issues like climate change, utilitarianism gives us the guidance we need. But those who worry that utilitarianism takes too little care to avoid grave risks can build on the connection between risk aversion and prioritarianism. They can offer a prioritarian theory that is able to capture the precautionary approach.¹⁴

Greg Bognar

New York University

NOTES

1. This scenario is discussed in Stern (2010, 42–45).
2. For an analysis of the structure of the precautionary principle, see Manson (2002).

3. For a brief discussion, see Gardiner (2010, 7–9). In addition, we also know very little about certain threshold effects—for instance, the point in warming that would cause the collapse of the West Antarctic Ice Sheet. Thus, we know little about the possibility of abrupt, as opposed to gradual, climate changes. On abrupt climate change, see National Research Council (2002).

4. I borrow this estimate from Stern (2010, 45–49). Gardiner (2010, 10) uses the figure of 2% of GNP per year.

5. Gardiner suggests that debates about the precautionary principle and climate change policy are ultimately about the question whether the conditions apply (2006, 55–56).

6. See Gardiner (2010, 7–8) and IPCC (2005) for an explanation of the language it uses to express risk and uncertainty. More accurately, there are different aspects of the problem of climate change, some of which involve uncertainty and some of which involve risk. Still, we are hardly in the circumstances of radical uncertainty that Rawls envisaged in the original position. (I thank Dale Jamieson for discussion on this point.)

7. It has long been recognized that Rawls’s own illustration of the maximin rule suffers from the same problem. Rawls uses the following example (1999, 133, n. 19):

	c_1	c_2	c_3
d_1	-7	8	12
d_2	-8	7	14
d_3	5	6	8

d_1 – d_3 are the alternatives you can choose from; c_1 – c_3 are the circumstances that determine the outcomes, represented by the numbers in the cells of the table. Because of the first condition, you don’t know the probabilities of these circumstances. The third condition obtains because d_1 and d_2 have very bad outcomes (-7 and -8). Using the maximin rule, you would choose d_3 . But a quick calculation shows that assigning 1/3 probability to each outcome and calculating the expected payoffs of the three alternatives give the same result. What’s more, a number of other proposed rules for choice under uncertainty do that too. (For details, see Ihara 1982 and Kaye 1980.)

8. This paragraph and the last follow Ihara (1982, 63–64). He also discusses a third possibility: there is an alternative that maximin rejects but which has very few hardly acceptable outcomes and a great many slightly better outcomes than the worst outcome of the alternative that maximin selects. Then the large number of these slightly better outcomes could together compensate for having hardly acceptable outcomes. As he argues, this requires that the decision maker is able to finely individuate possible outcomes, which is inconsistent with what Rawls says elsewhere about the choice in the original position. Intriguingly, Gardiner makes a similar point: “it seems likely that many actual cases where the Rawlsian conditions are met would not allow for such a fine-grained individuation of the alternatives” (2006, 49, n. 53).

9. Some commentators argue that perhaps Rawls makes some implicit assumption about the value of primary goods to the decision makers behind the veil. (See, for example, Hubin 1980.) One thing Rawls suggests is that the parties do not want to jeopardize equal basic liberties (1999, 135). But it is unclear to me why he thinks that the other two conditions are insufficient to secure them.

10. See Parfit (2001).

11. A detailed demonstration, which is beyond the scope of this paper, can be found in Broome (1991, 209–13).

12. By the harms of climate change, I mean the decrease in human well-being caused by its impact. I don't mean to deny that there are harms which are not harms to humans. Animals can be harmed. In addition, perhaps nature, species, and ecosystems also have value independently of their value to us. However, I will ignore these complications here.

13. There is, however, an important complication here. I have defined the harms of climate change as the loss of human well-being (including the loss of well-being due to the costs of mitigation). I also said that utilitarians would be indifferent in the example, since they want to maximize expected utility, and they interpret utility as the mathematical representation of well-being. So they want to maximize expected well-being. Although I believe this is in line with the standard interpretation of utilitarianism, it implicitly assumes that utilitarianism implies neutrality about risk to well-being. But in fact utilitarians can take into account the badness of risk to well-being by incorporating it into the utility representation. Their theory is consistent with risk-aversion about well-being. In this sense, perhaps utilitarianism and prioritarianism are not separate theories at all. (I thank John Broome for discussion on this point.) This complication raises further issues that I can't attempt to sort out here. For now, I will simply assume that prioritarianism can be developed so that it is either a distinct moral theory or at least a special form of utilitarianism. With this caveat, I will treat them as separate views, as they usually are.

14. For detailed comments on an earlier version of this paper, I would like to thank John Broome, Chlump Chatkupt, Dale Jamieson, Ben Sachs, and an anonymous referee. I would also like to thank audiences at the University of Cape Town, Penn State University, Northeastern University, the University of York, and participants at the International Society for Environmental Ethics session at the 2011 meeting of the American Philosophical Association in Minneapolis.

REFERENCES

- Broome, John 1991. *Weighing Goods: Equality, Uncertainty and Time*, Oxford: Basil Blackwell.
- Gardiner, Stephen M. 2006. "A Core Precautionary Principle," *Journal of Political Philosophy*, 14: 33–60.
- . 2010. "Ethics and Global Climate Change," in Gardiner et al. 2010, 3–35.
- Gardiner, Stephen M., Simon Caney, Dale Jamieson, and Henry Shue, eds. 2010, *Climate Ethics: Essential Readings*, New York: Oxford University Press.
- Harsanyi, John C. 1975. "Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls' Theory," *American Political Science Review*, 59: 594–606.
- Hubin, D. Clayton 1980. "Minimizing Maximin," *Philosophical Studies*, 37: 363–72.
- Ihara, Craig K. 1982. "Maximin and Other Decision Principles," *Philosophical Topics*, 12: 59–72.
- IPCC 2005. Guidance Notes for Lead Authors of the IPCC Fourth Assessment Report on Addressing Uncertainties, <http://www.ipcc.ch/pdf/assessment-report/ar4/wg1/ar4-uncertaintyguidancenote.pdf>.
- . 2007. *Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge: Cambridge University Press, http://www.ipcc.ch/publications_and_data/ar4/wg2/en/contents.html.
- Kaye, David H. 1980. "Playing Games with Justice: Rawls and the Maximin Rule," *Social Theory & Practice*, 6: 33–51.

- Manson, Neil A. 2002. "Formulating the Precautionary Principle," *Environmental Ethics*, 24: 263–74.
- McMichael, A.J., A. Nyong, and C. Corvalan 2008. "Global Environmental Change and Health: Impacts, Inequalities, and the Health Sector," *British Medical Journal*, 336: 191–94.
- National Research Council 2002. *Abrupt Climate Change: Inevitable Surprises*, Washington, DC: National Academy Press.
- Parfit, Derek 2001. "Equality or Priority?" in John Harris, ed., *Bioethics*, Oxford: Oxford University Press, 347–86.
- Rawls, John 1999. *A Theory of Justice*, revised edition, Cambridge, MA: Harvard University Press.
- Stern, Nicholas 2010. "The Economics of Climate Change," in Gardiner et al. 2010, 39–76.
- Sunstein, Cass R. 2005. *Laws of Fear: Beyond the Precautionary Principle*, Cambridge: Cambridge University Press.
- United Nations 1992. United Nations Framework Convention on Climate Change, <http://unfccc.int/resource/docs/convkp/conveng.pdf>.