

MYOPIA, THRESHOLDS AND PRECAUTION

CHRISTOPHER KNAPP

Binghamton University

In order to solve environmental problems, governments, businesses and citizens must often pay costs now so as to avoid or mitigate the costs of environmental damage in the future. In order to mitigate the future effects of global climate change, for instance, we must now pay the costs involved in dramatically reducing the amount of greenhouse gasses we emit; in order to prevent the future loss of valuable species, we must now pay the costs of keeping their habitat free from development. The resistance environmentalists meet when they encourage us to pay such costs is notorious, and there is no shortage of possible diagnoses for the myopia that infects many institutions and individuals and that often frustrates responsible environmental initiatives. We are reluctant to make sacrifices now in order to avoid or mitigate environmental problems later because we lack temperance, imagination, understanding, direct contact with nature, ... the list of familiar explanations continues. The familiar list is, however, incomplete. Consequently, even when taken together, these common diagnoses of our myopia do not do justice to the depth of the problem that effective environmental activism must somehow solve.

What the familiar list leaves out is that our myopia is sometimes a consequence of perfectly reasonable intuitions about tradeoffs between different kinds of value. In certain cases, putting off addressing environmental problems that will arise in the future in order to avoid paying economic or social costs in the present can seem rational, even though we can see that such procrastination will ultimately prove detrimental to the things we care about. Enacting effective environmental policies will thus involve persuading people and institutions to make choices that will seem to some of them to be irrational. This need not be an insuperable obstacle, as our occasional success in solving collective action problems shows. Indeed, I will offer a strategy for correcting our myopia below. But before considering solutions, it pays to understand the problem we face.

1. Torture and Catastrophe

At the root of this practical problem lies a theoretical quandary that was first pointed out by Larry Temkin and Stuart Rachels. They ask us to consider the following three claims:

- (1) For any unpleasant or “negative” experience, no matter what the intensity or duration of that experience, it would be better to have that experience than to have one that was only a little less intense but twice as long.
- (2) There is a continuum of unpleasant or “negative” experiences ranging in intensity, for example, from extreme forms of torture to the mild discomfort of a hangnail.
- (3) A mild discomfort for the duration of one’s life would be preferable to two years of excruciating torture, no matter the length of one’s life.¹

Each of these claims is intuitively compelling. Taken together, however, they are puzzling. For consider the following way of representing them graphically:




Fig. 5

Claim (2) insures that the slope of the line representing the range of pain intensity from excruciating torture down to a hangnail is smooth. And, if true, claim (1) would insure that for each pair of adjacent pains in the series from A to

¹ This formulation is from Temkin (1996), p. 179. Rachels’ formulation can be found in Rachels (1998).

Y, the pain that is slightly more intense but half as long is the better pain to experience. For example, claim (1) tells us that pain A (i.e., two years of excruciating torture) would be a better pain to experience than pain B (i.e., four years of slightly less excruciating torture); and it tells us that pain B would be better to experience than pain C (i.e., 8 years of torture that is slightly less excruciating than pain B); and so on down the series of pains until we reach pain Y. But pain Y is the pain of a hangnail that is experienced for all of a 32 million-year long life, and claim (3) tells us that 2 years of excruciating torture must be worse than any such hangnail pain.

For present purposes, the lesson to draw from this example is that a single person can make a series of choices, each of which seems rational, that nonetheless leads her to an outcome that is worse than one she could have otherwise achieved.² Imagine, for instance, that you were given the choice between the final two pains represented in the chain above as X and Y. You are likely to find X the rational choice – that is, you would think it rational to choose 16 million years of hangnail pain over 32 million years of hangnail pain that was only slightly less intense. Now imagine that, having made this choice, you are presented with the option of trading the pain you have just chosen for the pain represented above as W – that is, a pain that is only slightly more intense than the hangnail pain you just choose, but that lasts only half as long. Again, your choosing W over X seems eminently rational. If this sequence of choices continues, you will keep choosing the slightly more intense pain that lasts only half as long until you wind up choosing A over B. And insofar as each choice seems to be for the less bad pain, each choice will seem rational. But at the end of the sequence you wind up choosing to experience two years of excruciating torture, whereas had you chosen differently, you could have experienced only the annoyance of a hangnail, albeit for a very long time. And, according to most of us, this is deeply unfortunate, as it would be have been much better for you to live with a hangnail for a very, very long time than to undergo two years of excruciating torture.

² This is a significantly weaker conclusion than Temkin and Rachels argue for. They take the example to show that the ‘better than’ relation is intransitive. This is, however, a controversial interpretation. Indeed, I have argued for a less revolutionary interpretation in Knapp (forthcoming). Others who have challenged this as an argument for intransitivity include Norcross (1997) and Binmore and Voorhoeve (2003).

This is puzzling. We are confident that you should not make this series of choices. But it is far from clear where in the series you should stop. No matter where you stop, you will be choosing to suffer a pain that is only slightly less intense but lasts twice as long as a pain that you could have chosen instead. And since to make such a choice seems to be to choose the greater of two pains, keeping yourself from two years of torture will involve making what seems, at the time, to be an irrational choice.

No one has actually faced this particular series of choices; and, let us hope, no one ever will. But this series is just a dramatic instance of a kind of situation that is all too common. Many environmental problems present us with a structurally similar series of choices. And just as in our hypothetical series, governments, businesses and individuals will have to make what seem to them to be irrational choices in order to avoid producing very bad outcomes.

To see that this is so, it will help to understand what generates the puzzle in the Temkin/Rachels example. The badness of a pain depends on both its intensity and duration. Consequently, when we compare the badness of two pains, we normally must compare both their pain and intensity. Sometimes a pain that is more intense may nonetheless be less bad if it does not last as long; and sometimes a pain that lasts longer may nonetheless be less bad if it is less intense. This, of course is just what makes each individual choice in the hypothetical example seem so rational. Since the slightly more intensely painful option lasts only half as long, it seems less bad, and hence it seems rational to choose it. But when we compare two years of torture to the pain of a hangnail, the duration of the hangnail no longer seems to matter. This is because these two pains are qualitatively different: torture and hangnails produce very different *kinds* of pain. It is this qualitative difference that makes the duration of the hangnail irrelevant to the comparison of hangnail pain to the pain of torture. So with respect to comparing pains, our evaluations are sensitive to whether or not there is a qualitative difference in intensity at stake. If not – if we are comparing only quantitative differences in intensity and duration – then trading off lesser intensity to gain shorter duration can be rational. But if one pain is qualitatively worse than the other, then it can be that no such tradeoff makes sense.

The problem arises in the Temkin/Rachels example because the qualitative difference in intensity between the pain of torture and the pain of a hangnail seems to arise gradually rather than at a specific point. There seems to be no specific intensity of pain that is the precise minimum amount necessary to be of the same kind of intensity as torture. As a result, when we are faced with each

particular choice in the series, it never seems that there is a qualitative difference in intensity at stake. Our choices always seem to be between two pains comprised of different quantities of intensity and duration, but the intensities are always of the same kind. This explains why we always think choosing a slightly more intense pain so as to experience pain for half as long makes sense. Over the entire course of the series, however, there is a qualitative difference at stake, and this explains why we think that, taken together, our choices lead to a worse outcome than we could have attained, despite the fact that each individual choice seems to have been for the best.³ Let us call the practical problem that arises in the Temkin/Rachels case the 'evaluative threshold problem.'

We are now in a position to see that some environmental issues present us with series of choices that are analogous to the Temkin/Rachels series, and hence will generate the evaluative threshold problem. For in some environmental cases it does seem that: 1) there is a qualitative difference that is evaluatively relevant in much the same way that qualitative differences in pain intensity are; 2) whether the qualitatively worse outcome is produced depends on how we make a series of choices; and 3) in none of the individual choices does it seem that the qualitative difference is at stake. Let us consider an example.

Suppose there is a country that emits a large quantity of carbon dioxide. The current administration in this country accepts that these emissions contribute to the greenhouse effect, and hence to changes in the global climate. Provided that these climactic changes are relatively small and relatively gradual, they believe, the net amount of harm such changes will cause to the things they care about will be less than the cumulative costs they would have to incur to significantly reduce their CO₂ emissions. However, they also believe that if concentrations of atmospheric CO₂ get sufficiently high, enough radiant heat would be retained to dramatically slow the currents in the Atlantic Ocean that provide warmth to the northern parts of the Western hemisphere.⁴ If the currents slow enough, the effects on the things they care about would be catastrophic. And were they presented with the choice of whether to pay the cumulative costs of significantly

³ This diagnosis is proposed in Temkin (1996), p. 194 and Rachels (1998), p. 71. I discuss it in detail in Knapp (forthcoming). It is not the only possible diagnosis of the case; but it is the diagnosis that best fits the environmental analogues of the case that are our primary concern here.

⁴ The scientific bases for this and other possibly catastrophic effects of climate change can be found in National Research Council (2002).

limiting their emissions or the costs involved with catastrophic cooling, they would gladly choose the former.

The administration must decide whether to now pay the costs associated with significantly reducing CO₂ emissions, or to put off paying such costs and revisit the issue in, say, three months. They might reason thus. If, in order to prevent catastrophic cooling, it were necessary to pay the costs right now, then we would surely do it. But three months of higher emissions could not possibly make the difference between catastrophic and non-catastrophic cooling.⁵ In the very worst case scenario, those three months of higher emissions would mean that the ocean currents slow only ever so slightly more, and hence that we experience only ever so slightly cooler temperatures. That is, at worst, such a small difference might make the difference between a catastrophe and an ever so slightly worse catastrophe, but it could not possibly constitute the difference between a catastrophe and no catastrophe at all. Thus, since there is no chance that putting the decision to institute a greenhouse gas abatement policy off for three months will cause a catastrophe, let us put off paying the costs a bit longer. Three months later, the administration revisits the question, but none of the basis for their prior reasoning has changed. So they – and their successors – continue to put off reducing emissions, satisfied at each step that doing so makes the most sense. Eventually the Atlantic currents slow, and catastrophe ensues.

To put the point in more theoretical terms, this government's evaluation of whether the benefits of reducing emissions outweigh the costs depends on whether the qualitative evaluative difference of a catastrophic harm is at stake. When there is no question of creating a catastrophe, they believe that the cost of producing each little bit of climate stability is greater than the benefit that small amount of stability will bring. The economic costs of three additional months of abatement are greater than the quantitative improvement in climactic stability three months of abatement would produce. And since there seems to be no precise level of cooling that is catastrophic, at each decision point, there seems to be no chance of creating a catastrophe. So at each decision point, it seems better to defer paying the costs of reducing emissions. The government thus always chooses to defer making significant reductions in emissions. But the cumulative

⁵ If epistemic accounts of vagueness are true, this claim is not. See, for instance, Williamson (1994), Chapters 7 and 8. But it is hardly unreasonable to believe that epistemicism is false. Moreover, the kinds of precise thresholds epistemicism postulates are not the kind of thresholds that matter evaluatively. For a discussion of this point in another context, see Alter and Rachels (2004).

effect of all these deferrals is sufficient to cause the catastrophe – the costs of which outweigh what the CO₂-abatement costs would have amounted to had the government started paying such costs at the outset. They have thus made a series of decisions, each of which seemed to be for the better, but which taken together, produce a worse outcome than they could have achieved.⁶

Let me reiterate that I am not claiming that each of the administration's decisions to delay really is for the best, nor that each of their decisions is, in fact, rational. Whether they are for the best and whether they are rational depend on the resolution of difficult questions about the nature of value and its relation to rational choice this sort of example raises. The point, rather, is that people might reasonably believe that these decisions are rational, and hence that persuading these people to pay costs now so as to avoid eventual catastrophe will involve persuading them to do what they have good reason to believe is irrational.⁷

⁶ It requires no great leap of the imagination to take the hypothetical administration of this example to be the current administration in the United States. The Bush administration has acknowledged the existence of anthropogenic climate change. It does not believe, however, that the marginal climate stability that would be achieved by relatively ambitious reductions in emissions are worth the cost. Even the relatively modest reductions required by Kyoto Protocol would have been, in the words of EPA Administrator Mike Leavitt, "bad for the United States." On the other hand, at least some in the administration also seem to be deeply troubled by the possibility of dramatic cooling resulting from weakening ocean currents. For instance, a report commissioned by the Pentagon and released in 2003, predicted that greenhouse effect-induced cooling would thrust the U.S. into a world in which "Europe is struggling internally, large numbers of refugees are washing up on its shores and Asia is in serious crisis over food and water. Disruption and conflict will be endemic features of life." But so far, at least, the preferred response to any such dire eventualities is to defer making significant sacrifices. "As we advance science and develop technology to substantially reduce greenhouse gas emissions in the long term," the Bush administration explains, "we do not want to risk harming the economy in the short term. ...a rapid reduction in emissions would be costly and threaten economic growth."

⁷ Andreou (2006) argues for a similar conclusion by drawing an analogy between environmentally destructive behavior and the example of the self-torturer presented in Quinn (1990). My diagnosis of our environmental myopia is different than the one suggested by Quinn's puzzle, however, in that Quinn's puzzle relies on the claim that imperceptible effects have no evaluative

This example raises a number of interesting questions. First, how close is the analogy between the Temkin/Rachels case and the catastrophic cooling case? Secondly, how far does the evaluative threshold problem extend? And finally, what resources are available to environmentalists to address these problems? In the following sections, I take up these questions in turn.

2. Tipping Points and Foresight

A crucial feature of the Temkin/Rachels example is that in none of the individual tradeoffs of intensity for duration does it seem that one could be trading one kind of pain for another that is qualitatively worse. Thus, if the catastrophic cooling case is to be strictly analogous to the Temkin/Rachels example, it must also be that trading off three months of high emissions in order to pay fewer economic costs never involves trading a level of emissions that has non-catastrophic effects for a level of emissions that has catastrophic effects. But here, the analogy might break down. For instance, there might be a specific level of emissions such that, once it is reached, a chain of events is set into motion that makes a drastic slowing of the Atlantic currents inevitable. If so, then there will be a specific three-month period during which this tipping point is reached. And this can make it seem that there will be at least one point at which it should not seem rational for the government to put off paying the costs of abatement for another three months.

We should be clear about the source of the disanalogy. In the Temkin/Rachels case, the fact that a qualitative difference never seems to be at stake in any individual decision is due to two factors. First, no precise intensity of pain seems to mark the threshold between torturous and non-torturous pain. And second, the intensity of pain can be finely calibrated so that it is always the case that one can be given the choice between one pain and another that is only ever-so slightly more intense. The potential disanalogy in the catastrophic cooling case concerns not the first, but the second factor. That is, the question here is not

significance. This is not the right diagnosis of the examples considered here, however. For in these cases, we readily acknowledge that what we are trading off is significant, it is just that these values seem to be outweighed. This is important, since Arntzenius and McCarthy (1997) have given a compelling argument that, contrary to Quinn, effects that are imperceptible in Quinn's sense are evaluatively significant, and hence that Quinn's self-torturer is demonstrably irrational after all.

whether there might be a precise level of bad effects that marks the difference between a catastrophe and no catastrophe: there is not. Rather, the question is whether the effects of greenhouse gas emissions also occur gradually, so that an extra three months of higher emissions will always only make the world only ever so slightly worse than it would have been had we emitted less during those three months. If there is a tipping point, the answer to this question will be no.

Unfortunately, we have, at present, no way to know whether this potential disanalogy is actual. As is often the case with such complicated environmental systems, scientists cannot now say exactly how the slowdown in the Atlantic currents would proceed, and so cannot say whether there is a tipping point beyond which catastrophic cooling is inevitable. Moreover, even if such a tipping point exists, we cannot know whether our past emissions have already pushed us past it. For even assuming that there is a precise level of emissions that represents a point of no return on the way to catastrophic cooling, the cascade of climactic and oceanographic changes may unfold quite some time after that level has been exceeded. And finally, even if the tipping point exists and is at a level beyond the current concentration of greenhouse gases, we are far from being able to say during which three-month period of high emissions that point will be reached.

The possibility of a tipping point for catastrophic cooling thus complicates our imagined administration's decision-making in the following way. Whereas we previously assumed that they could be absolutely confident that putting off deciding to pay the costs of reducing emissions for three months would not cause a catastrophe, we must now allow that such absolute confidence would be misplaced. That said, with respect to any particular delay, they have virtually no grounds for believing that it will constitute the difference between a catastrophe and no catastrophe at all. They have insufficient reason to believe that a tipping point exists; they have even less reason to believe that such a tipping point has not already been passed; and virtually no reason to believe that it will be passed in the next three months. Thus, if they found it rational to delay taking abatement action before they considered the possibility of a tipping point, it is likely that they will, quite reasonably, continue to believe that it is rational after taking this remote possibility into account.

To put the point in more general terms, it is true that there might be a tipping point which makes it the case that the evaluatively significant qualitative difference between a catastrophic and a non-catastrophic outcome really is at stake in some particular decision about whether to curtail emissions rather than

to put off abatement for another three months. But with respect to any given decision to delay, the administration has virtually no reason to believe that the tipping point is about to be crossed. So when each decision is made, it is reasonable for them to believe that one of the following three things is true: the tipping point does not exist; it has already been passed; or it will only arise at some later date.⁸ And given this disjunctive belief, it remains reasonable for them to conclude that, since they can be confident that they are not making a decision that will create catastrophic cooling, it is better to put off deciding to curtail emissions for another three months. The evaluative threshold problem thus remains.

There is, however, another important disanalogy between the Temkin/Rachels case and the catastrophic cooling case. In presenting the Temkin/Rachels case, it is easy to imagine that the person who is presented with the series of choices between pains does not know that she will continue to receive such options in the future, and thus cannot know that her choices will result in torture. Our administration, on the other hand, knows that the country will continue to face the choice of whether to curtail greenhouse gas emissions into the foreseeable future, and thus should foresee that if they and future administrations continually choose to delay, catastrophe will ensue. Since they can see that their decisions to delay are part of a pattern that, if continued, will lead to catastrophe, they seem to have an additional, and particularly pressing, reason to stop making such choices.

This is an important consideration to keep in mind, as it is likely to figure in any solution to the general problem. But we should be careful to remember that while such considerations do put pressure on the country to stop delaying abatement at some point or other, it is far less clear how it puts pressure on the country to stop at any given point rather than, say, three months in the future.

Here is one way to bring such pressure to bear. The country must acknowledge that, if it is to avoid a catastrophe, it must stop delaying abatement at some point, and at any given point, it will seem better to delay than begin (or continue with) abatement. Any particular time will seem an equally unsuitable time to do what nonetheless needs to be done. Given this, the country should recognize that now

⁸ Or, perhaps more precisely, it is reasonable for the administration to have a much higher degree of confidence that this disjunction is true than that the proposition “The tipping point will occur in the next three months” is true.

is as good a time as any to begin doing what is necessary. And since it must do what is necessary at some point, it might as well start now.

At best, this line of argument establishes that it is reasonable for the administration to begin abatement now, rather than put the decision off any further. It does not, however, show that it is unreasonable for them to continue to delay: to say the country might as well begin paying costs now is not to say that it must do so. Moreover, this argument for the reasonableness of starting to pay costs now is most forceful only when one adopts an impersonal or objective point of view. It becomes far less compelling when we descend from such a God's-eye perspective. First, from the perspective of many decision-makers, costs that must be borne now are worse than equally great costs that are borne in the future: many decision-makers discount the value of future costs (and benefits). Thus, from our administration's perspective, it may not be that any particular time is an equally unsuitable time to begin paying the necessary costs of abatement; indeed, the present may be the worst time, since the costs involved with abatement now would now be greater than any such costs that will be paid in the future. Of course, such thinking may be a mistake. There are good arguments for thinking that future costs and benefits should not be discounted.⁹ But the arguments for this conclusion may be controversial enough that someone who discounts the future need not be unreasonable.

Secondly, and more importantly, the current administration can be confident that it will not be in office indefinitely. As a result, while some administration must decide to begin (and continue) abatement policies, it need not be they. And they are likely to have reasons generated by their own projects and commitments to think it better that future administrations bear the costs associated with abatement. The current administration, for instance, might have other policy initiatives that it thinks are crucially important to implement, and beginning an abatement policy might detract from such efforts. More cynically, the administration might also think that abatement policies are unlikely to be popular, and since one of their projects is to stay in power, they have a reason not to pursue such policies. To be sure, the administration should recognize that similar considerations will arise for future administrations. But it is not unreasonable to believe that the projects and commitments of future administrations do not provide the current administration with reasons that are nearly as strong as those provided by their own projects and commitments. And

⁹ For arguments against discounting future costs and benefits, see Parfit (1984), Appendix F.

if the current administration's projects provide them with stronger reasons than future administrations' projects do, the current administration can conclude that they have more reason leave abatement to future administrations than to start it now: from the current administration's perspective it is better that future administrations sacrifice their own projects for the cause of abatement than that the current administration do so. Now perhaps consequentialists are right that the strength of the reasons provided by one's own projects and commitments is solely a function of the value pursuing them has from the impersonal perspective, but since this is hardly a settled matter in moral theory, it is not unreasonable for the current administration to believe otherwise.¹⁰

There are important differences between the Temkin/Rachels example and the catastrophic cooling case, then. But the differences are not sufficient to undermine the apparent rationality of each individual decision to trade off a three-month delay in instituting abatement policies in exchange for an extra three months free of the economic costs those policies would entail. Even given the possibility of a tipping point, and the administration's ability to foresee the possibility of catastrophe if the pattern of delay continues, they are not being unreasonable if they decide that it now makes most sense to put off deciding to institute abatement policies for another three months. And the reasons for thinking putting off abatement makes most sense now are unlikely to change in the future.

3. Vagueness and Ignorance

Before discussing strategies for overcoming the evaluative threshold problem, it worth pausing to consider how far it extends. It is not possible to give an exhaustive list of all the environmental contexts in which it arises, of course. But it is instructive to consider a few more examples.

Many environmentalists are concerned that we insure the continued existence of wild or natural places. Wildness and naturalness, though, are often in tension with other desirable things. To increase public access to a place by building paths and infrastructure is to ever so slightly decrease its naturalness, for instance. To control the spread of invasive species by removing individual plants is to ever so slightly decrease a place's wildness. So long as there is no threat of turning a once wild or natural environment into one that is controlled or artificial,

¹⁰ For a defense of the reasonableness of giving special weight to one's own projects, see Scheffler (1982), Chapter 3.

particular and small efforts to increase access or preserve native species will seem to make sense. The problem, however, is that while no particular choice to blaze a trail, build a shelter, root out a weed from a pond, or to monitor an animal's reproduction will change a wild or natural environment into one that is controlled or artificial, taken together, a series of such incursions will. And if such a series of incursions is made, we may well feel that the cumulative gains were not worth the sacrifice of a wild or natural place that they required.

Similarly, many of us are concerned that the water people drink, the air people breathe and the soil people's food grows on be safe; or, more generally, that the environment people live in be a healthy one. We recognize, however, that it is folly to try rid our environment of all potentially harmful contaminants. We tolerate impurities and pollutants, that is, so long as their level does not make our water air and soil unsafe or unhealthy. Thus, it can often seem to make sense to allow small increases in the levels of contaminants if this is required to produce more jobs in an economically struggling neighborhood, or to maintain the ability to drive cheaply and conveniently, or to avoid paying the costs of cleanup. For any relatively small addition to the level of toxic substances will not make the difference between a healthy or safe environment and an unhealthy or unsafe one. But again, a series of such additions can. And we may believe that, given the choice between the jobs and convenience and money that were saved by allowing ever more contaminants into our environment and a healthy and safe place for people to live in, we would surely choose the latter.

Both of these problems arise in part because "wild," "natural," "healthy," and "safe" are, like "catastrophic," vague terms. Their vagueness underlies our belief that none of the qualitative differences they mark is at stake in many of the small-scale decisions in which we trade off small decreases in these values in order to achieve other things we care about. We should be particularly on guard for the evaluative threshold problem, then, in any other context in which we use vague terms to mark evaluatively important differences in kind. It would be a mistake, however, to infer that the evaluative threshold problem will only arise in contexts in which our views about permissible trade-offs depend on qualitative boundaries marked by vague terms. For vagueness is not the only thing that can make us believe that a small-scale decision will not be what pushes over an evaluatively significant threshold.

Recall the earlier discussion of the possibility of a greenhouse-effect tipping point: a precise degree of greenhouse gas concentration that, if reached, would set off a cascade of effects that would make catastrophic cooling inevitable. Here

I argued that even if it is the case that such a tipping point exists, and hence that there is a particular point at which a decision to delay instituting abatement policies for three months really does make the difference between catastrophic and non-catastrophic cooling, the evaluative threshold problem still arises. This was because at each decision point, it is not unreasonable for the administration work under the assumption that now is not the crucial time – given their ignorance about the existence or location of the tipping point, it is not unreasonable for them to assume that they a tipping point will not be crossed in the next three months. This opens the possibility that there are precise evaluatively significant qualitative thresholds that also generate the problem.

Consider, for instance, the following scenario. In many cases, housing development reduces the size of wild animal and plant populations by reducing their habitat. Many people think that the benefits such development provides often outweigh the cost of smaller animal and plant populations. But some of the same people also think that some plant and animal species are so valuable that no reasonable number of tract homes could compensate for any of these species' extinction. That is, these people believe that trading off a quantitative decrease in the size of animal or plant populations for the benefits associated with new housing typically makes sense, but building so many new houses that a species goes extinct often does not. These people treat some species extinctions as evaluatively significant qualitative changes.

Now suppose that these people have the power to control the pace of development in areas that a valuable threatened species uses as habitat, and that the development is a slow, piecemeal process, stretching out over the course of many years. As each proposal for development is made, these people must decide whether to approve a small amount of development – and thereby gain the benefits associated with it – or to reject the small amount of development – and thereby preserve a small amount of habitat for a threatened species. Even if we assume that the issue of whether this species goes extinct is a precise one, these decision-makers are deeply ignorant of whether the loss of any particular parcel of the species' habitat will mean the difference between its survival and its extinction.¹¹ And since at every decision point, it is exceptionally unlikely that the

¹¹ Not all matters of extinction are precise, because species terms admit of vagueness. See, for instance, the discussion of transitional "semi-species" in Mayr (1996). In many actual cases of extinction, however, we can explain the process of extinction without referring to any individuals who are borderline cases of the species in question.

development of the parcel under review will be what dooms the species – it will be far more likely that the threshold lies either in parcels that will be proposed in the future or parcels that have already been developed – the decision-makers work under the assumption that the specific parcel under review is not crucial, that they are only choosing between forgoing the benefits of development and a quantitative decrease in the threatened population. On this basis, they approve the proposed development. But their epistemic situation does not change dramatically over the course of all the proposals they consider; to be sure, the likelihood that the crucial parcel lies in the future decreases, but this is offset by the increasing likelihood that it lies in the past. So they continue to make such approvals, believing at each point that it is for the best, thereby damning to extinction the species they care about more than development.

Like the case of catastrophic global warming, this case highlights a moral challenge associated with our ignorance of extraordinarily complex relationships that constitute ecosystems. In particular, it warns us that we should be on our guard for the evaluative threshold problem not just when we are dealing with evaluatively significant qualitative differences marked by truly vague boundaries, but also when we do not know whether an evaluatively significant precise threshold lies before us or whether we have already crossed it some time ago. When such conditions hold, as they frequently do, reasonable people can fall into the trap of making a series of apparently rational tradeoffs that they will later regret.

4. Precaution and Irrationality

How can we avoid falling prey to the evaluative threshold problem? Since the problem arises when we try to choose the best of the options available to us at any given moment, the answer will involve making our choices in a different way. We will need to implement a different, non-maximizing, decision-making strategy.

Fortunately, philosophers and environmentalists have been refining and advocating a strategy that, in broad outline, seems to fit the bill – namely, the Precautionary Principle. The often-cited “Wingspread” formulation of the principle reads:

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.

One of the great advantages of the Precautionary Principle is that – like the idea that we always ought to choose what we believe will be best – it is intuitively appealing. Indeed, many of its defenders present it as little more than an application of the commonsense maxim “Better safe than sorry” to environmental policy decisions. This intuitive appeal gives us reason to hope that if adopting a version of the precautionary principle could help us solve the evaluative threshold problem in theory, it might also be able to do so in practice.

We do not need anything so broad as the principle proposed by the Wingspread Statement to solve this particular problem, however. And since it is doubtful that the principle advocated by the Wingspread Statement can be justified across the broad range of cases in which it purports to apply, for present purposes it is best to consider the following, more restricted principle:

The Threshold Precautionary Principle: When we have good reason to believe that continuing an activity will push us across an evaluatively significant threshold, and crossing that threshold would be worse than foregoing the gains of continuing the activity, then we should take the precautionary measure of stopping the activity immediately, even if we lack sufficient reason to believe the threshold is imminent.

That is, when we believe we are facing an evaluative threshold problem, it is better to be safe (stop the risky activity immediately) than sorry (produce an outcome that by our own lights is worse than we could have achieved had we stopped the activity immediately).

Were we to reasonably believe that we currently faced any of the instances of the evaluative threshold problem above, following the Threshold Precautionary Principle would stop us from making a series of decisions that we will later regret. For if we found ourselves in the examples described above, the principle would instruct us to immediately reduce our greenhouse gas emissions; immediately stop intervening in the natural processes occurring in wilderness areas; immediately stop polluting the air, soil or water; and immediately halt development in critical habitat. And doing these things would be the most effective way we have to avoid causing a climate catastrophe, eliminating valuable wild places, creating unhealthy environments, and driving important species to extinction.

The Threshold Precautionary Principle also has the virtue of avoiding or diffusing the most potent objections raised against other, more common formulations of the Precautionary Principle. For instance, some commentators have complained that principles like the one advocated in the Wingspread Statement unjustifiably single out environmental and health-related values as being the only values deserving the protection of the better-safe-than-sorry strategy. This worry cannot be pressed against the Threshold version, though, as it applies to any evaluative threshold problem, no matter what values are stake.

Another complaint raised against some versions of the Precautionary Principle is that they imply that no amount of other values could outweigh any threat to an environmental or health-related value. This extreme conservatism with respect to the environmental and health status quo is implausible – when the potential gains are great, running small environmental or health risks must sometimes be permissible. The Threshold version of the Precautionary Principle, however, is not nearly so conservative as this. It applies only when we reasonably believe that we are facing a threshold problem, and even then it demands precaution only if we believe the gains of continuing the threatening activity would not outweigh the costs of crossing the threshold. When the threatened environmental or health-related harm is small and the benefits of the threatening activity are great, the Threshold Precautionary Principle simply will not apply.

The restrictions built into the Threshold Precautionary Principle also mean that it will not paralyze policy-makers. There is a concern that unrestricted versions of the Precautionary Principle would recommend that policy-makers both regulate and not regulate all activities that threaten the environment or human health, for there is always a risk that regulation itself will harm the environment or human health.¹² But since it is extraordinarily unlikely that both regulating and failing to regulate a given activity will generate an evaluative threshold problem, this worry does not apply to the Threshold principle. In all, or nearly all, the evaluative threshold problems we face, the Threshold principle will tell us to stop the threatening activity without also telling us to continue it.

There is, finally, one objection to the Precautionary Principle that the Threshold principle may not be able to avoid – namely, that following the Precautionary Principle is irrational.¹³ If rationality requires always choosing what you

¹² Sunstein (2003), p. 1028.

¹³ Lomborg (2001), p. 350.

reasonably believe to maximize the expected value of the outcome, the Threshold Precautionary Principle does indeed counsel irrationality. For it sometimes tells us to immediately stop a productive activity when we lack sufficient reason to believe that immediately stopping that activity is necessary to avoid incurring losses that outweigh the gains that activity promises.¹⁴ But we are now in a position to see that this objection has no bite in the cases in which the Threshold Precautionary Principle applies. For in these cases, choosing what we reasonably believe to maximize expected value is likely to produce an outcome that is worse than we will produce if we follow the Threshold Precautionary Principle. That is, the Threshold Precautionary Principle tells us to be irrational in this sense only when being rational would be self-defeating.

I do not mean to suggest that the Threshold Precautionary Principle is the only defensible version of the Precautionary Principle. There may well be other formulations that manage to avoid or diffuse the objections typically raised against it.¹⁵ If so, we may have reason to embrace them as well. What I do claim, however, is that we now have sufficient reason for accepting the Threshold principle as a warranted departure from the strategy of choosing what we reasonably believe will produce the best outcome. Our impacts on the environment sometimes threaten to push us across evaluatively significant thresholds. When we see that they do, we ought to now pay the costs associated with ceasing those activities.

5. Conclusion

Environmentalists often lament that we are unwilling to protect the environment because we are not persuaded that there is a dire emergency, but that by the time the emergency becomes apparent, it will be too late to do anything about it. Our resistance to making sacrifices for the sake of environmental protection, they say, is dangerously myopic: we put off paying costs now in order to avoid environmental catastrophes in the future at our peril. This paper has provided a diagnosis of some instances of that myopia, according to which it is the product

¹⁴ However, if the analysis of threshold cases I offer in Knapp (forthcoming) is correct, then the reasonable belief that continuing the productive activity is for the best is often mistaken. For, according to that analysis, when we are in the borderline region of a vague evaluatively significant distinction, we can neither affirm nor deny that the tradeoffs we make the outcome better: they are tradeoffs whose consequences are borderline cases of 'better.'

¹⁵ See, for instance, the proposal in Gardiner (2006).

of reasonable people choosing what they take to be best, and hence doing what a certain conception of rationality requires. Appreciating that our myopia can have the weight of reason and rationality behind it is important. It shows us what we are up against, and it suggests strategies for overcoming it. For if we can help others see that doing what they reasonably believe is for the best will, in certain cases, predictably make things worse, we stand a chance of persuading them to choose another way.

REFERENCES

- Alter, Torin and Rachels, Stuart. (2004) "Epistemicism and the combined spectrum argument." Ratio 17: 241-255.
- Andreou, Chrisoula. (2006) "Environmental damage and the puzzle of the self-torturer." Philosophy and Public Affairs 34: 95-108.
- Arntzenius, Frank and McCarthy, David. (1997). "Self torture and group beneficence." Erkenntnis 47: 129-144.
- Binmore, Ken and Voorhoeve, Alex. (2003) "Defending transitivity against Zeno's paradox." Philosophy and Public Affairs 31: 272-279.
- Gardiner, Stephen. "A core precautionary principle." The Journal of Political Philosophy 14: 33-60.
- Knapp, Christopher. (forthcoming) "Trading quality for quantity." Journal of Philosophical Research.
- Lomborg, Bjørn. (2001) The Skeptical Environmentalist: Measuring the Real State of the World. Cambridge: Cambridge University Press.
- Mayr, Ernst. (1996) "What is a species, and what is not?" Philosophy of Science 63: 273-4.
- National Research Council, Committee on Abrupt Climate Change. (2002) Abrupt Climate Change: Inevitable Surprises. The National Academies Press.
- Norcross, Alastair. (1997) "Comparing harms: headaches and human lives." Philosophy and Public Affairs 26:135-167.
- Quinn, Warren. (1990) "The puzzle of the self torturer." Philosophical Studies 59: 79-90.
- Parfit, Derek. (1984) Reasons and Persons. Oxford: Oxford University Press.

- Rachels, Stuart. (1998) "Counterexamples to the transitivity of *better than*." Australasian Journal of Philosophy 76: 71-83.
- Scheffler, Samuel. (1982) The Rejection of Consequentialism. Oxford: Oxford University Press.
- Sunstein, Cass, R. (2003) "Beyond the precautionary principle." University of Pennsylvania Law Review 151: 1003-1058.
- Temkin, Larry, S. (1996) "A continuum argument for intransitivity." Philosophy and Public Affairs 25: 175-210.
- Williamson, Timothy. (1994) Vagueness. New York: Routledge.