

Discounting, Climate Change, and the Ecological Fallacy*

Matthew Rendall

Discounting future costs and benefits is often defended on the ground that our descendants will be richer. Simply to treat the future as better off, however, is to commit an ecological fallacy. Even if our descendants are better off when we average across climate change scenarios, this cannot justify discounting costs and benefits in possible states of the world in which they are not. Giving due weight to catastrophe scenarios requires energetic action against climate change.

In recent years concern has grown that runaway climate change could mean disaster. Martin Weitzman estimates a roughly 5 percent chance of a rise in global temperatures of more than ten degrees Celsius over the next two centuries in the absence of energetic intervention, with a chance of roughly 1 percent that it could go over twenty—changes which would pose “an extreme threat to human civilization and global ecology as we now know it, even if it might not necessarily mean the end of *Homo sapiens* as a species.”¹ Provided that they did not lead to our extinction, these losses would burden many generations of future people. A vast amount of expected utility is at stake, even if the probability of catastrophe is low.

* I thank participants at workshops and conferences in Salamanca, Graz, Frankfurt, Lille, Ithaca, Cardiff, and London for comments and advice—especially Elizabeth Baldwin, Simon Dietz, Anja Karnein, Paul Kelleher, Antony Millner, and Kian Mintz-Woo—two anonymous referees, and the editors of *Ethics*.

1. Martin L. Weitzman, “On Modeling and Interpreting the Economics of Catastrophic Climate Change,” *Review of Economics and Statistics* 91 (2009): 1–19, 1; Martin L. Weitzman, “Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change,” *Review of Environmental Economics and Policy* 5 (2011): 275–92, 282.

Ethics 129 (April 2019): 441–463

© 2019 by The University of Chicago. All rights reserved. 0014-1704/2019/12903-0001\$10.00

One might think that economics, with its roots in the utilitarian tradition, would give this great weight. Yet some prominent economists have recommended only modest near-term measures to mitigate climate change. When the British government's 2006 *Stern Review on the Economics of Climate Change* called for energetic action, its conclusions were attacked as extreme. Perhaps the most common criticism was that the *Stern Review* had employed an inappropriately low discount rate, giving too much weight to costs and benefits in the future. William Nordhaus presented a *reductio ad absurdum*:

Wrinkle: A "wrinkle" in the climate will reduce world consumption of goods and services by a tiny amount starting in 2200, but this loss can be prevented by timely intervention today. How great a sacrifice would be justified in order to do so?

"Using the *Review's* growth projections," Nordhaus charged, "the *Review* would justify reducing per capita consumption for one year today from \$10,000 to \$4,400 in order to prevent a reduction of consumption from \$130,000 to \$129,870 starting two centuries hence and continuing at that rate forever after."²

Nordhaus was describing what Derek Parfit calls a "Rawls-Scanlon case." In these situations, made famous by John Rawls and T. M. Scanlon, "we can either save one person from some great burden, or give much smaller benefits to many other people, who are all much better off."³ Examples include the following:

1. *Chocolates for the Well-off*: By allowing a single person to starve, we can give one chocolate each to a very large number of well-off people. Should we do it?
2. *World Cup*: Jones has had an accident in a TV station that is currently broadcasting the World Cup and is suffering agonizing pain. Should we interrupt the broadcast to rescue him at the price of frustrating hundreds of millions of soccer fans?

2. William D. Nordhaus, "A Review of the *Stern Review on the Economics of Climate Change*," *Journal of Economic Literature* 45 (2007): 686–702, 696.

3. Derek Parfit, *On What Matters*, ed. Samuel Scheffler (Oxford: Oxford University Press, 2011), 2:246. *Chocolates for the Well-off* is based on an example in Roger Crisp, "Equality, Priority, and Compassion," *Ethics* 113 (2003): 745–63, 754. I have borrowed Yitzak Benbaji's term for the case ("Sufficiency or Priority?," *European Journal of Philosophy* 14 [2006]: 327–48, at p. 330). *World Cup* is from T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Belknap Press of Harvard University Press, 1998), 235. *Blue and the Many* is a composite of examples in Alistair Norcross, "Comparing Harms: Headaches and Human Lives," *Philosophy and Public Affairs* 26 (1997): 135–67; John Broome, "All Goods Are Relevant," in *Summary Measures of Population Health: Concepts, Ethics, Measurement, and Applications*, ed. Christopher J. L. Murray et al. (Geneva: World Health Organization, 2002), 727–29, 728; and Parfit, *On What Matters*, 2:246–47. The name is Parfit's.

3. *Blue and the Many*: If the National Health Service (NHS) can either spare Blue a year of severe pain or relieve the day-long migraine headaches of a large number of people, how many of the latter would be needed to make relieving the headaches the better choice?

These cases reflect a familiar worry about utilitarianism: since it allows the aggregation of costs and benefits, it can justify placing great burdens on some in order to bring small benefits to others. Nordhaus's "wrinkle experiment" is a version of this problem: unless future costs and benefits are severely discounted, utilitarianism implies that the prospect of small benefits to a vast number of the rich (our descendants) could justify imposing severe sacrifices on the poor (ourselves). This "argument from excessive sacrifice"⁴—that giving equal weight to future costs and benefits would impose intolerable obligations on the present generation—has been one of the most influential arguments for the economic practice of discounting. "If there were no time preference," writes Kenneth Arrow, "what would the optimal solution be? Each unit sacrificed would yield a finite utility loss to the first generation, but to compensate, there would be a gain, however small, to each of an infinity of generations. . . . Given any investment, short of the entire income, a still greater investment would be preferred." As a justification for a pure rate of time preference—which accords greater importance to costs and benefits because they arrive sooner—Arrow invokes Samuel Scheffler's concept of an "agent-centered prerogative," according to which agents can defensibly give extra weight to their own interests.⁵

Discounting as usually practiced, however, fails to solve either Nordhaus's or Arrow's paradox. Even if increasing the *Stern Review's* discount rate of roughly 1.4 percent to Nordhaus's preferred 4 percent avoids the conclusion that we should sacrifice more than half of next year's income in *Wrinkle*,⁶ to revive the problem we need only stipulate that more people will benefit. Consider the following:

Wrinkle_{DS}: Between now and the time in which Nordhaus's wrinkle appears in 2200, human beings will be shrunk to the height of five inches, as in the movie *Downsizing*, enabling the earth to support a much larger population. How great a sacrifice would be justified to prevent the wrinkle now?

4. Derek Parfit, *Reasons and Persons* (Oxford: Clarendon, 1987), 484.

5. Kenneth J. Arrow, "Discounting, Morality, and Gaming," in *Discounting and Intergenerational Equity*, ed. Paul R. Portney and John Weyant (Washington, DC: Resources for the Future, 1999), 13–21, 14.

6. William D. Nordhaus, *A Question of Balance: Weighing the Options on Global Warming Policies* (New Haven, CT: Yale University Press, 2008), 10; Olivier Godard, "Time Discounting and Long-Run Issues: The Controversy Raised by the Stern Review of the Economics of Climate Change," *OPEC Energy Review* 33 (2009): 1–22, 4.

So long as world population were more than 108,000 times that of today, cost-benefit analysis would recommend the sacrifice, even if the \$130 gained by maintaining consumption at \$130,000 instead of \$129,870 were a one-off.⁷ The Stakhanovite investment rates that Arrow hopes to avert would also then be justified.⁸ Yet surely both authors mean to claim that these sacrifices would not be required however many future people stood to gain. If discounting is to yield this conclusion, it must be understood differently.

This article probes the ethics underpinning Nordhaus's and Arrow's arguments. It argues that an agent-centered prerogative can justify refusing to sacrifice oneself for any number of people who would in any case be much better off. Costs and benefits to everyone else should be discounted to a limited extent, at most. Economists commonly assume that economic growth will leave future generations richer than the present one, in spite of climate change. Even the *Stern Review* assumed a 1.3 percent base rate of growth, leading to accusations that it demanded too much of the present generation. "It is a peculiar feature of the Review," wrote a group of critics, "that while forecasting that people in the future will be vastly richer than today, it also proposes that the present generation should make substantial new sacrifices on behalf of these more prosperous generations. It is as though, looking back two hundred years . . . we claimed that people living in the early days of industrialisation ought to have made sacrifices on behalf of those living today, even though we are rich beyond the dreams of anyone in those distant times."⁹ In fact, as Thomas Schelling noted, some members of future generations—and particularly victims of climate change—will be poorer than the present-day inhabitants of rich countries. To discount costs and benefits to these people on the ground that future people will be much richer on average is to commit an ecological fallacy—that is to say, treating a statistic that describes a group in the aggregate as if it described its individual members. That a group enjoys an average per capita level of consumption does not imply that this is the level that each member, or even the average member, enjoys.¹⁰ Nevertheless, Schelling suggested that given the expected growth of GDP in the

7. At a 4 percent rate of discount, the value of \$130 in two hundred years is $130/(1 + 0.04)^{200} = \$0.05$. Then, $\$0.05 \times 108,000 = \5400 , the per capita loss to present-day people from the sacrifice.

8. See Stephen M. Gardiner, *A Perfect Moral Storm: The Ethical Tragedy of Climate Change* (New York: Oxford University Press, 2011), 294.

9. Ian Byatt et al., "The Stern Review: A Dual Critique. Part II: Economic Aspects," *World Economics* 7 (2006): 199–229, 214; Nordhaus, "Review of the *Stern Review*," 694.

10. Thomas C. Schelling, "Intergenerational Discounting," *Energy Policy* 23 (1995): 395–401, 398; Jerry Ratcliffe, "The Ecological Fallacy," April 27, 2015, <http://www.jratcliffe.net/blog/the-ecological-fallacy/>. Schelling's description of this as a "fallacy of composition" was a misnomer, as I discuss below.

coming decades, resources might be better spent on poverty alleviation in the present—an argument that writers like Bjørn Lomborg have pressed energetically.¹¹

What Schelling overlooked is that we must also distinguish among possible states of the world. In most scenarios envisioned by economists, future generations in the aggregate are richer—and, implicitly, better off. This assumption may seem optimistic—and clearly would not hold for future nonhumans.¹² But even if we grant it, it does not justify discounting future damages away. We can defensibly disregard costs and benefits to people who will be much better off. In at least a few scenarios, however, most future people are worse off. To treat them as if they will be better off because, averaged across a set of scenarios, future people are better off is to commit another ecological fallacy. Assume that many economists' optimistic assessment of the odds is right, and that there is—say—only a 1 percent chance that climate change will be so bad as to impoverish the lives of our descendants. The expected number of people impoverished—understood as the number of future individuals multiplied by their chance of impoverishment—is nonetheless enormous. There is no excuse for ignoring these costs—or for foot-dragging on climate change.

I. IS THE APPEAL TO A PREROGATIVE NEEDED?

The force of Nordhaus's "wrinkle experiment" comes from our feeling that there must be something wrong with any analysis that implies that we could be obliged to submit to deprivation and even death (let us call it "the sacrifice") in order to prevent a large number of rich people from suffering a minor reduction in consumption. Can this problem be solved without appealing to pure time preference?

We might think that the sacrifice would not in fact make the world a better place, that it would be better if we did not make it. On this view, there is something wrong with the value theory underpinning the utilitarian judgment. It might seem that if the *Stern Review's* discount rate recommended the sacrifice in *Wrinkle*, or unreasonable rates of investment, the problem was insufficient concern about how benefits are distributed. Partha Dasgupta suggests that the solution is not pure time preference, but for discounting to give greater weight to equality.¹³ Nevertheless, such weighting would not be absolute. While it might avoid implying that the

11. Schelling, "Intergenerational Discounting," 399–400; Bjørn Lomborg, *Cool It: The Skeptical Environmentalist's Guide to Global Warming* (London: Marshall Cavendish, 2010).

12. James Brander, review of *The Skeptical Environmentalist*, by Bjørn Lomborg, *Journal of Economic Literature* 40 (2002): 972–74, 973; Wayne Hsiung and Cass R. Sunstein, "Climate Change and Animals," *University of Pennsylvania Law Review* 155 (2007): 1695–1740.

13. Partha Dasgupta, "Discounting Climate Change," *Journal of Risk and Uncertainty* 37 (2008): 141–69.

sacrifice would be optimal in *Wrinkle*, the problem would resurface in *Wrinkle_{DS}*. So long as concern for equality—or priority given to the worse off—is merely weighted, even tiny benefits to the well-off will trump it, provided that there are enough of them.¹⁴

The sufficiency view offers an appealing alternative. It holds that what is morally important is that everyone should “have enough.”¹⁵ If maintaining quality of life above a certain level has absolute priority, then the sacrifice would not have the best result. There would be no need to discount the future to avoid excessive obligation.¹⁶ Yet the sufficiency view faces a severe objection: the “black hole problem,” or the “problem of waste.”¹⁷ Many people suffer medical conditions that render their lives seriously flawed or limited. “There can easily exist medical procedures,” Arrow warned many years ago, “which are . . . so expensive as to reduce the rest of the population to poverty.” So long as we can do anything to ameliorate their condition, however minor, the sufficiency view will tell us to accord them priority. Even in the case of incurable diseases, so long as there is a chance that medical research could discover a way of curing or mitigating them, we will have to do so before we do anything for those who are better off.¹⁸ Long before we eliminate all illnesses, disabilities, and risks, sufficientarian policies will drive everyone in society down to the level where life is just barely good enough, and this will be the best feasible outcome.¹⁹ The sufficiency view blocks the sacrifice in *Wrinkle*, but it has counterintuitive implications of its own.

14. Crisp, “Equality, Priority, and Compassion,” 754; T. M. Scanlon, “Replies,” *Ratio* 16 (2003): 424–39, 433.

15. The original formulation of the view, by Harry Frankfurt, stated this in terms of people having enough resources, particularly money. Nevertheless, “some of [Frankfurt’s] remarks seem to hint quite strongly at welfare,” as Robert Huseby observes, and this has been the focus of most sufficientarian analyses. See Harry Frankfurt, “Equality as a Moral Ideal,” *Ethics* 98 (1987): 21–43; and Robert Huseby, “Sufficiency: Restated and Defended,” *Journal of Political Philosophy* 18 (2010): 178–97, 181 n. 11.

16. Matthew Rendall, “Climate Change and the Threat of Disaster: The Moral Case for Taking Out Insurance at Our Grandchildren’s Expense,” *Political Studies* 59 (2011): 884–99. Dominic Roser defends a nonconsequentialist version of this argument: “The Discount Rate: A Small Number with a Big Impact,” in *Applied Ethics: Life, Environment and Society* (Sapporo: Center for Applied Ethics and Philosophy, Hokkaido University, 2009), 10–25.

17. Huseby, “Sufficiency,” 186–87; Clark Wolf, “Intergenerational Justice, Human Needs, and Climate Policy,” in *Intergenerational Justice*, ed. Axel Gosseries and Lukas H. Meyer (Oxford: Oxford University Press, 2009), 347–76, 356–57, 366–67.

18. Kenneth J. Arrow, “Some Ordinalist-Utilitarian Notes on Rawls’s *Theory of Justice*,” *Journal of Philosophy* 70 (1973): 245–63, 251; Mark S. Stein, *Distributive Justice and Disability: Utilitarianism against Egalitarianism* (New Haven, CT: Yale University Press, 2006), 78, 187–88; Karl Widerquist, “How the Sufficiency Minimum Becomes a Social Maximum,” *Utilitas* 22 (2010): 474–80.

19. Norcross, “Comparing Harms,” 59–60; Widerquist, “How the Sufficiency Minimum.”

A final response would be to grant that the sacrifice would make things go best but to deny that we would be obliged, or even permitted, to bring about the best outcome. This is Rawls's and Scanlon's approach. Scanlon maintains that our actions must be consistent with principles "that no one could reasonably reject as a basis for informed, unforced general agreement." Agents can, on this view, reasonably reject principles aggregating costs and benefits when the burden for themselves would be grossly disproportionate. Thus, in the *World Cup* case, Jones could reasonably reject any principle that left him to suffer, whereas no soccer fan could reasonably reject a principle requiring us to rescue Jones.²⁰ Similarly, we might think that we could reasonably reject any principle requiring us to sacrifice 50 percent of next year's income for the sake of preventing miniscule losses from climate change stretching far into the future. Both Rawlsian contractualism and Scanlonian contractualism, however, face the same problem as the sufficiency view. Why, we might ask, could someone not reasonably reject a principle whereby society chose to treat ten million people's headaches instead of seeking a cure for her rare disease—particularly if she had been born with the ailment and had never had even a chance of benefiting from this arrangement? And yet our societies make such decisions, and intuitively it seems permissible to do so.²¹ If these small gains for the many can outweigh large losses for the few, we need some explanation of why this should be true within generations but not, seemingly, between them. Can an appeal to agent-centered prerogatives provide one?

II. AGENT-CENTERED PREROGATIVES

Some have seen discounting as a way of capturing a prerogative to favor our own generation's interests over the general good. Most of us believe that we can defensibly give some priority to ourselves and to those close to us. Samuel Scheffler advocates recognizing an agent's moral right "to assign a certain proportionately greater weight to his own interests than to the interests of other people." Citing Scheffler's "agent-centered prerogative," Arrow argues that the present generation can legitimately discount

20. John Rawls, *A Theory of Justice*, rev. ed. (Cambridge, MA: Belknap Press of Harvard University Press, 1999); Scanlon, *What We Owe to Each Other*, 153, 235–40. Rawls did not believe that his difference principle applied to relations between generations; *Theory of Justice*, 253–54.

21. James Griffin, *Well-Being: Its Meaning, Measurement, and Moral Importance* (Oxford: Clarendon, 1986), 170–74. On Rawlsian contractualism and health care, see John Kekes, "A Question for Egalitarians," *Ethics* 107 (1997): 658–69, 666–67; Stein, *Distributive Justice and Disability*, chap. 6; Allen E. Buchanan, "Health-Care Delivery and Resource Allocation," in *Justice and Health Care: Selected Essays* (Oxford: Oxford University Press, 2009), 37–75, 63–65. Cf. Norman Daniels, "Health Care Needs and Distributive Justice," *Philosophy and Public Affairs* 10 (1981): 146–79.

the future, precluding any obligation to sacrifice everything for the sake of others.²² Understood thus, pure time preference is not really based on time, but rather on personal distance from the beneficiary.²³ Each generation could accord greater weight to its own interests by a factor which Scheffler calls “M.”²⁴

Formulated thus, Scheffler’s argument cannot do the work that Arrow wants it to do. In the first place, if discounting reflects personal distance, then the standard approach to discounting—under which costs and benefits receive less and less weight as they recede into the future—is senseless. While we care more about our children and grandchildren than about strangers, we do not attach many times as much weight to the welfare of people who will live in one century as to the welfare of those living two centuries from now.²⁵ Second, while Scheffler does not state what the value of M should be, or claim that it can be precisely specified, he clearly intends this weighting to be limited. So long as discounting is understood as weighting, however, a large enough population of future people—as in *Wrinkle_{DS}*—will always overwhelm it.²⁶ More people would have to benefit before we would be required to impoverish ourselves, but so long as enough stood to gain, we would still have to make the investment. How could Arrow’s argument be repaired to avoid the problem?

Here it will be helpful to consider two rationales for an agent-centered prerogative. One is rule consequentialist and appeals to overall consequences. Rule consequentialism, in its most plausible versions, requires us to observe principles that collectively yield the best results. Part of the costs and benefits of a moral code are its effects on our own feelings and on our relationships with others. A code that required us to assign equal weight to the welfare of strangers would be psychologically wrenching and would require us to attach far less importance to much of what makes our own lives worth living.²⁷ An across-the-board duty to treat strangers’ interests on a par with our own would also create a moral hazard. Such a strong duty of assistance would discourage people from helping them-

22. Samuel Scheffler, *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions* (Oxford: Clarendon, 1982), 20; Arrow, “Discounting, Morality, and Gaming,” 16.

23. Parfit, *Reasons and Persons*, 485.

24. Samuel Scheffler, “Prerogatives without Restrictions,” *Philosophical Perspectives* 6 (1992): 377–97, 378.

25. Martin L. Weitzman, “Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate,” *Journal of Environmental Economics and Management* 36 (1998): 201–8, 201.

26. Duncan Purves, “The Case for Discounting the Future,” paper presented at the workshop “Ethical Underpinnings of Climate Economics,” University of Helsinki, November 12, 2014, p. 5; see also Liam B. Murphy, *Moral Demands in Nonideal Theory* (Oxford: Oxford University Press, 2000), 64.

27. Parfit, *On What Matters*, 1:404.

selves. Obligations of assistance must be limited to avoid these perverse effects. Rule consequentialism will thus license us to give ourselves and those close to us some degree of preference. Let us call an agent-centered prerogative that is justified in this fashion PR_{RC} .

Scheffler, in contrast, does not claim that his prerogative will promote the best consequences, even in the long term. Rather, he defends it as the natural consequence of the fact that “people do not typically view the world from the impersonal perspective.”²⁸ Agents, in this conception, are not impartial optimizers. As Orwell remarked, “Many people genuinely do not wish to be saints.” While to treat others’ interests on a par with one’s own may be worthy of praise, few people do, and certainly one is not to be blamed if one puts a finger on the scales. But unless an action is blameworthy, in one key sense it is not wrong.²⁹ Let us call a prerogative justified in this fashion PR_{PPV} .

If an agent-centered prerogative can be justified in either fashion, we can plausibly believe that it would be optimal to make the sacrifice in *Wrinkle* while denying that we would be morally obliged to make it. Consider the following variation on *Blue and the Many*:

Mistake: The NHS can either cure Bill’s rare disease, sparing him a year of severe pain, or treat the day-long migraine headaches of ten million people. When Bill receives a letter from the NHS, he realizes that a computer glitch has led it to decide to treat his rare disease.

We might accept that treating the headaches would be best, and that in *Blue and the Many* the NHS ought to make this choice. Yet we may also think that Bill would not be obligated to notify the NHS of its mistake. Even if the NHS would be justified in imposing the best solution on Blue, we could not reasonably expect Bill voluntarily to accept a year of torment just to spare ten million strangers day-long headaches. Why?

On the first justification for an agent-centered prerogative, PR_{RC} , for Bill to be prepared to make the sacrifice in *Mistake* would require self-abnegation—essentially to regard himself as a means for maximizing impersonal goodness. For us all to have such attitudes would render life much less worth living. Bill is thus justified in regarding himself in a way that leads him to reject the sacrifice in *Mistake*. For us to be prepared to impoverish ourselves whenever it brought small benefits to enough rich people would also require extreme self-abnegation. That too would greatly re-

28. Scheffler, *Rejection of Consequentialism*, 62.

29. George Orwell, “Reflections on Gandhi,” in *In Front of Your Nose: 1945–1950*, ed. Sonia Orwell and Ian Angus (New York: Harcourt Brace Jovanovich, 1968), 463–70, 467; Brian McElwee, “The Rights and Wrongs of Consequentialism,” *Philosophical Studies* 151 (2010): 393–412; Stephen Darwall, “Agreement Matters: Critical Notice of Derek Parfit, *On What Matters*,” *Philosophical Review* 123 (2014): 79–105, 97–98.

duce human happiness. No plausible moral theory, so we might argue, will demand it. Rejecting the sacrifice in *Wrinkle*, on this rule consequentialist view, would be a suboptimal act which was nevertheless justified by the optimal moral code.

On the second justification, PR_{PPV} , Bill could justifiably keep mum because the sacrifice would be his. Douglas Portmore claims that impartiality in such a case would be downright irrational: “If . . . my choice is between saving my own life . . . and saving some stranger’s life, and the only thing that speaks in favor of saving the stranger is that this would result in there being one more utile overall, then it seems that I have not just sufficient, but decisive, reason to save my own life.”³⁰ If that is true where the stakes for each party are equal, then surely it is true a fortiori when Bill would lose so much more. But we need not go even this far. We can grant that Bill rationally could make the sacrifice without holding that he is rationally required to do so.³¹ Likewise, we could not reasonably be expected to make the sacrifice in *Wrinkle*. We could argue that we could not be expected to impose harsh sacrifices on ourselves for the sake of people who would in any case be much better off than we are.³²

When would it be unreasonable on PR_{PPV} to expect such a sacrifice? T. M. Scanlon has famously proposed the concept of *relevance*. On this view, aggregating costs and benefits is appropriate if and only if the sacrifices for the respective parties would be, so to speak, in the same ballpark. If it is a matter of saving one person’s life or many people from total paralysis, then we might justifiably save the latter. But if it is a choice such as the one in *Mistake*, then the disadvantaged party can reasonably reject a principle requiring the sacrifice, no matter how many stand to benefit. That is true even if it would be impartially better to treat the headaches. While the disadvantaged party might have greater impartial reason to accept a principle requiring her to make the sacrifice, Scanlon argues, “taking *all* reasons into account”—both impartial and agent-centered—she can reasonably reject it. The numbers simply do not count.³³

In Rawls-Scanlon cases, “we can either save one person from some great burden, or give much smaller benefits to many other people, who are all much better off.” We might think that both the much greater loss the disadvantaged party would suffer and the fact that the others are much better off make it permissible for her to reject it. Parfit shows that the for-

30. Douglas W. Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford: Oxford University Press, 2011), 4.

31. Parfit, *On What Matters*, 1:382–88, 2:258–59.

32. Hilary Greaves, “Discounting for Public Policy: A Survey,” *Economics and Philosophy* 33 (2017): 391–439, 407.

33. Scanlon, *What We Owe to Each Other*, 239–40; T. M. Scanlon, “How I Am Not a Kantian,” in Parfit and Scheffler, *On What Matters*, 2:116–39, 137; emphasis in the original.

mer consideration is irrelevant. Consider the following variation on *Mistake*:³⁴

Shared Pain: The NHS can either cure Bill's rare disease, sparing him a year of severe pain, or reduce the severe pain of Bill and ten million other people from a year to ten months. When Bill receives a letter from the NHS, he realizes that a computer glitch has led it to decide to treat his rare disease.

Here again, Bill has more to lose than the others. Informing the NHS will cost him ten months of pain, while keeping mum will mean only two months more pain for the other ten million people. Nevertheless, he could not reasonably reject a principle requiring him to do so. The reason that Bill could defensibly reject the burden in *Mistake* is not that his sacrifice would be greater, but that it would leave him so much worse off than the others.

III. SCANLONIAN AND KANTIAN CONTRACTUALISM

This raises a puzzle. Intuitively, Bill is under no moral obligation to volunteer information to the NHS. Yet commonsense morality accepts that the state can impose aggregative trade-offs in cases such as *Blue and the Many*. Were Blue to resist such a decision—say, by hacking the NHS computers to change it—the government would also be justified in forcing her to stop. It appears, then, that the state can defensibly compel Blue to accept a decision with which Bill would have no moral obligation to comply on his own.

The justification for this conclusion under PR_{RC} is straightforward. Many policies, when imposed, place less strain on our moral willpower than would the duty of voluntary compliance. Governments, as Thomas Nagel points out, can justifiably collect taxes to help the poor, without taxpayers being morally obliged to contribute voluntarily, because the latter “requires voluntary decisions that are quite difficult to make.”³⁵ To volunteer for the sacrifice in *Mistake* or *Wrinkle*—where the beneficiaries will be better off—would be harder still. Plausibly, it would require a degree of self-abnegation that it would be very bad to have.³⁶ In *Blue and the Many*, in contrast, Blue has no choice and need not screw herself up for a sacrifice or diminish her self-esteem in the same way. On the logic of PR_{RC} , the best

34. For a similar example, see Parfit, *On What Matters*, 2:197–98.

35. Thomas Nagel, “Libertarianism without Foundations,” *Yale Law Journal* 85 (1975): 136–49, 145.

36. Bernard Williams, “A Critique of Utilitarianism,” in *Utilitarianism: For and Against*, by J. J. C. Smart and Bernard Williams (Cambridge: Cambridge University Press, 1973), 116–17.

rules will instruct the state to impose some burdens that the losers would not be obliged to accept of their own accord.

This conclusion is also implied by PR_{PPV} . Some writers have objected to Scheffler's agent-centered prerogatives on the grounds that they would permit "forced supererogation." Schefflerian prerogatives permit the agent to give greater weight to her own interests, thus allowing actions that are impartially suboptimal. But since it is always permissible on Scheffler's view to take the impartially best action, it seems that his theory will allow one to force others to behave optimally when doing so makes things go best. Some commentators find this unacceptable.³⁷ In turn, Scanlonian contractualism, which admits that agent-centered reasons can justify rejecting some optimistic principles, is sometimes criticized for permitting too little aggregation.³⁸ We saw earlier that it had trouble explaining how the NHS could justifiably treat ten million headaches rather than pursuing a cure for someone's rare and innate disease. Yet most of us accept that society can legitimately impose such trade-offs. Moreover, Scanlon's acceptance of agent-centered reasons creates, as he himself notes, another problem: if everyone can justifiably accord greater weight to her own interests, then it seems we risk "a moral standoff, in which there is no right answer to the question of what one should do."³⁹

Here is a solution. Scanlonian contractualism holds that we must act consistently with principles that no one could reasonably reject. "The desire to be able to justify one's actions (and institutions) on grounds one takes to be acceptable," Scanlon notes, "is quite strong in most people. People are willing to go to considerable lengths, involving quite heavy sacrifices, in order to avoid admitting the unjustifiability of their actions and institutions." This criterion, Scanlon argues, offers the most phenomenologically plausible account of moral motivation.⁴⁰ Here Scanlon equates a principle's justifiability with its acceptability. Yet when choosing whether to grant or withhold assistance, we ask not whether others could accept our decision, but whether we can justify it to ourselves. In *Mistake*, Bill can hold that he is justified in keeping mum, on the grounds that he could reasonably reject a principle requiring him to speak up.

37. Bashshar Haydar, "Forced Supererogation and Deontological Restrictions," *Journal of Value Inquiry* 36 (2002): 445–54; Eric Mack, "Prerogatives, Restrictions, and Rights," *Social Philosophy and Policy* 22 (2005): 357–93; see also R. H. Myers, "Prerogatives and Restrictions from the Cooperative Point of View," *Ethics* 105 (1994): 128–52, 144–48.

38. See, e.g., Alistair Norcross, "Contractualism and Aggregation," *Social Theory and Practice* 28 (2002): 303–14.

39. Scanlon, "How I Am Not a Kantian," 138.

40. T. M. Scanlon, "Contractualism and Utilitarianism," in *Utilitarianism and Beyond*, ed. Amartya Sen and Bernard Williams (Cambridge: Cambridge University Press, 1982), 103–28, 117; Scanlon, *What We Owe to Each Other*, 187.

That does not, on the other hand, block others from imposing such sacrifices. On what Parfit calls Kantian contractualism, it is wrong to treat people in a way that they could not rationally accept. We always have sufficient reason to accept principles that lead to the impartially best outcome.⁴¹ This suggests a distinction that neither Kant nor Scanlon draws himself: Scanlonian contractualism governs what we owe to others, and Kantian contractualism what we can justifiably do to others.⁴² Because Blue would be so much worse off if we choose to treat many people's headaches instead of her year of pain, we should plausibly give this burden disproportionate weight. But if, even after weighting, treating the headaches remains the optimific choice, Blue could rationally accept it, and society can justifiably impose it. As Roger Crisp observes, "The justification an agent can offer for pursuing her own interests at the expense of the overall good is that these interests are hers. And, as Sidgwick recognised, it does not strike us as unreasonable that an agent give some priority to her own interests. But a similar claim, in support of an alleged constraint, by a *patient* in some case where her interests may be sacrificed to the overall good lacks any plausibility. For the same point could be made by those whose interests are also at stake and in conflict with hers."⁴³

An agent can reasonably reject some sacrifices in Rawls-Scanlon cases that it would nevertheless be legitimate to impose. Though society can justifiably choose to treat the headaches in *Blue and the Many*, Bill is under no obligation to adopt this decision of his own free will in *Mistake*. Nor would we have the duty to impose the sacrifice on ourselves in *Wrinkle*. Neither PR_{RC} nor PR_{PPV} , moreover, is sensitive to the number of people the sacrifice would benefit. To be willing to make the sacrifice in *Wrinkle* would require a harmful degree of self-abnegation regardless of whether it would benefit billions of rich future people, or hundreds of billions, or trillions. Nor

41. Parfit, *On What Matters*, 1:chap. 8, 16–17, 2:250–53. Here Parfit claims that we have sufficient reason to accept principles that would make things go best if they were universally followed—or, on another version, accepted. Elsewhere, however, he rightly suggests that principles should also be optimific however many people comply with them (1:317–19). This implies that the right principles must actually make things go best.

42. This distinction solves a notorious problem for Scheffler's prerogative: that it seems to authorize not only doing less good than we might but also suboptimal acts of harm. Shelly Kagan, "Does Consequentialism Demand Too Much? Recent Work on the Limits of Obligation," *Philosophy and Public Affairs* 13 (1984): 239–54, 251. Because we could reasonably reject principles that required us to treat our own interests on a par with those of strangers, Scanlonian contractualism will permit some omissions that are impartially suboptimal. But since no one could reasonably accept principles allowing her to be treated in a way that was both harmful and impartially suboptimal, Kantian contractualism will prohibit all such actions. See Parfit, *On What Matters*, 2:251.

43. Roger Crisp, "The Dualism of Practical Reason," *Proceedings of the Aristotelian Society, New Series* 96 (1996): 53–73, 65–66; emphasis in the original. See also Parfit, *On What Matters*, 2:250–53.

could this number make a difference to the strength of PR_{ppv} . Unlike conventional pure time preference, which is restricted to weighting, these prerogatives justify absolute agent-centered preference within a limited scope. When others will be in any case much better off, the preference for one's own interests can legitimately be absolute. In other contexts, it should only be weighted.

IV. THE USE AND ABUSE OF AGENT-CENTERED PREROGATIVES

The foregoing arguments for absolute agent-centered prerogatives presuppose that beneficiaries of the sacrifice will be in any case much better off. To see this requirement's intuitive plausibility, recall *Mistake*. There we argued that Bill could reasonably reject a principle that required him to make a great voluntary sacrifice for people who are much better off. Bill would be justified in regarding benefits to these people as irrelevant to his decision. Consider now the following:

Big Mistake: The NHS can either cure Bill's rare disease, saving him from a thousand days of severe pain, or use the same resources to prevent the deaths of ten million people. When Bill receives a letter from the NHS, he realizes that a computer glitch has led it to treat his rare disease.

This time, Bill should clearly speak up. Not only will saving ten million be far better from an impartial point of view, but if he does not, he will forgo the opportunity to prevent a harm that would leave its victims even worse off than he is. While Bill still has a self-regarding reason to give greater weight to his own suffering, it could not be as great as this. Likewise, even defenders of conventional discounting admit that considerable sacrifice would be justified to prevent climate change from impoverishing future people.⁴⁴

This analogy might be challenged. We might think that Bill has a duty to his fellow citizens to speak up, but that mitigating climate change is, in Schelling's words, "very much like a foreign aid program, with some of the foreigners being our own descendants who live not on another continent but in another century."⁴⁵ Commonsense morality holds that it is permissible to give greater weight to benefiting conationals than foreigners, even if the latter are much worse off. We might think we may also give greater weight to benefiting our own generation. J. Paul Kelleher describes a case resembling *Wrinkle*:

44. See, e.g., Nordhaus, *Question of Balance*, 147.

45. Schelling, "Intergenerational Discounting," 397.

Volcano: All members of the present generation enjoy the present average per capita consumption level of Americans today. A volcano will erupt in 2200, permanently reducing everyone's consumption to the US per capita average of 1980. We can prevent the eruption, but at the cost of reducing the consumption of each member of the present and all future generations to the US per capita average of 1985.

In this case the future people will not be better off than the present generation. Nevertheless, Kelleher suggests that "many will take the view that members of the current generation may exercise Schefflerian prerogatives to protect their own consumption levels, especially since they (*ex hypothesi*) do not cause the volcanic eruption."⁴⁶

Perhaps many would. Still, in contrast to *Wrinkle*, it does not seem obviously unreasonable to hold that the present generation would be obliged to intervene. Moreover, the difference that intervention would make in *Volcano* is small, and the US living standard in 1980 may strike us as quite good enough in absolute terms. While intervening would be the impartially best choice, it is hard to imagine anyone losing much sleep over condemning future generations to the consumption level of 1980 rather than 1985. It would be another matter if the volcano threatened to reduce future generations to destitution. Consider the following:

Vicious Volcano: All members of the present generation enjoy the present average per capita consumption level of Americans today. A volcano will erupt in 2200, permanently reducing everyone's consumption to that of today's Haitians. We can prevent the eruption, but at the cost of reducing the consumption of each member of the present and all future generations to the US per capita average of 1985.

Clearly, PR_{RC} will not authorize us to stand by. Any moral code that makes things go best in the long run will surely insist that we intervene against an irreversible catastrophe for humanity at modest cost to ourselves.⁴⁷ Schefflerian prerogatives, in contrast, typically permit agents to give their own goals priority even when by acting otherwise they could save many others from severe poverty. Nevertheless, Scheffler himself holds that PR_{PPV} has finite weight. Hundreds or thousands of times as many people may live in the future as exist today.⁴⁸ It is hard to believe that an agent-

46. J. Paul Kelleher, "Pure Time Preference in Intertemporal Welfare Economics," *Economics and Philosophy* 33 (2017): 441–73, 468.

47. Brad Hooker, *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality* (Oxford: Clarendon, 2000), 169.

48. If we avoid catastrophe, it could be many more than that. See Adrian Kent, "A Critical Look at Risk Assessments for Global Catastrophes," *Risk Analysis* 24 (2004): 157–68, 164; Nick Bostrom, "Existential Risk Prevention as Global Priority," *Global Policy* 4 (2013): 15–31, 18–19.

relative prerogative could justify abandoning such large numbers to poverty. If an American could single-handedly save a whole town from poverty at the cost of reducing her living standard to that of 1985, surely it would be wrong for her to refuse to do so. Nor ought the present generation to reject the sacrifice in *Vicious Volcano*. If Schefflerian prerogatives have limited weight, then catastrophic future losses should outweigh small costs in the present. Scanlon's criterion of relevance reinforces this conclusion. Because personal interests help determine which principles agents can reasonably reject, Scanlonian contractualism permits them to give some extra weight to their own interests. But the damages in *Vicious Volcano* would affect so many people, and leave them so much worse off than we are, that we could not reasonably reject principles requiring a modest sacrifice from us to save them.

In the end, I doubt that many advocates of pure time preference would claim we have an agent-centered prerogative to reduce future generations to poverty. Runaway climate change could leave distant future generations poorer than we are. Even if they are richer, they may not be better off. Severe global warming might so impoverish these people's lives that increased consumption could not compensate for the loss.⁴⁹ Even if we can permissibly give less weight to their welfare than our own, these losses must not be discounted away.⁵⁰

V. IS THERE A RISK?

Is there any reason to worry about such scenarios? Defenders of discounting commonly argue that in spite of climate change, future people will be better off than we are. Yet serious analyses that are optimistic about our ability to manage global warming acknowledge some chance of catastrophe.⁵¹ Robert Mendelsohn writes that "future generations are *likely* to be wealthier than current generations." Bjørn Lomborg claims that emissions cuts will "help people far into the future, where they *will* be much richer," but even he adds that "we are not helping a poor Bangladeshi in 2100 but *much more*

49. Gardiner, *Perfect Moral Storm*, 176–77, 177 n. 48, 283–85, 326; Mathias Frisch, "Climate Change Justice," *Philosophy and Public Affairs* 40 (2012): 225–53, 231–32, 238–40.

50. In Sec. III, I argued that while Scanlonian contractualism governs what we owe to others, Kantian contractualism governs what we can do to them. While the former permits a limited degree of agent-centered preference, the latter does not. It is tempting to argue that since grave climate change damages would be worse both from an impartial perspective and for the personal interests of future people, the latter could not rationally accept any discounting of these damages. Such a claim, however, is complicated by the "nonidentity problem." Ultimately, we seem on firmer ground in holding that even if some agent-centered discounting of such damages can be justified, their impartial badness outweighs it. See Parfit, *On What Matters*, 2:chap. 22.

51. See, e.g., Richard S. J. Tol, "Why Worry about Climate Change? A Research Agenda," *Environmental Values* 17 (2008): 437–70, 447, 450–51.

likely a rather rich Dutchman.” William Nordhaus concedes that “we cannot rule out the potential for catastrophic impacts that might justify trillions of dollars of abatement costs,” before advising that “we should start with the clear and present dangers, after which we can turn to the unclear and distant threats.”⁵²

Should we? Weitzman’s influential warning about “bad tail” scenarios prompted an exchange with Nordhaus, who argued that for a catastrophe to materialize, the climate would have to prove improbably sensitive to greenhouse gas emissions, the damages from climate change would have to be improbably high, and scientists and governments would have to remain improbably complacent in the face of danger signs. Weitzman’s guesstimates of a 5 percent chance of more than ten degrees of warming, or a 1 percent chance of more than twenty, were exaggerated.⁵³ In so arguing, Nordhaus implicitly granted that there was a small chance of catastrophe. Weitzman’s reply was noteworthy. Suppose, he wrote, that a policy maker found Weitzman’s estimates and another economist’s less alarming ones equally credible. The policy maker should give Weitzman’s greater weight: “Suppose one person advises you that a fire insurance policy protecting your house against extreme losses is unnecessary because so few houses of your kind burn to the ground, while another person advises you that a complete fire insurance policy is necessary in your case. Other things being equal, should you flip a coin to decide what to do just because both advisers seem to be giving equally credible guidance?”⁵⁴

Here Weitzman made a simple but subtle point. There are two distinct mistakes we can make in such a case. One is to follow the “maxi-probability method”—determine what is most likely to happen, and do whatever will be optimal if it does. Policy debates sometimes take this approach for granted, but it is foolish.⁵⁵ Nobody expects her house to burn down next year. The maxiproability method would tell us to save our money and not bother with insurance. Calculating expected value, in contrast, allows us to take unlikely contingencies such as house fires into account.⁵⁶ Less widely recognized is that our probability estimates may be wrong. In fact, this may pose a greater risk than does ignoring the small

52. Robert Mendelsohn, “Is the *Stern Review* an Economic Analysis?,” *Review of Environmental Economics and Policy* 2 (2008): 45–60, 52; Lomborg, *Cool It*, 55; Nordhaus, *Question of Balance*, 147; emphasis added.

53. William D. Nordhaus, “An Analysis of the Dismal Theorem,” Cowles Foundation Discussion Paper No. 1686 (New Haven, CT: Cowles Foundation for Research in Economics, 2009), <http://cowles.yale.edu/sites/default/files/files/pub/d16/d1686.pdf>, 16–21.

54. Weitzman, “Fat-Tailed Uncertainty,” 291.

55. Sven Ove Hansson and Mikael Johanneson, “Decision-Theoretical Approaches to Global Climate Change,” in *International Politics of Climate Change: Key Issues and Critical Actors*, ed. Gunnar Fermann (Oslo: Scandinavian University Press, 1997), 153–78, 164.

56. John Broome, *Climate Matters: Ethics in a Warming World* (New York: Norton, 2012), chap. 7.

chance of a genuinely improbable disaster.⁵⁷ Here the error is to adopt the “my-favorite-theory approach”—in this case, assuming that a given set of estimates is correct. Instead, we should assign subjective probabilities to competing estimates (in Weitzman’s example, the house owner assigns 50 percent to each) and use these probabilities to weight the expected value of the actions’ outcomes.⁵⁸ Such probabilities can be hard to assess, but we must do the best we can. Refusing to use them at all would be, as Dominic Roser puts it, like “a visually impaired pilot closing her eyes on account of the impairment rather than using whatever impressions she can glean from her limited eyesight.”⁵⁹ We then choose the policy with the highest *intertheoretic expected value*.⁶⁰

Suppose that on an optimistic assessment there is one chance in a hundred of more than ten degrees of warming, and we assign a subjective probability of 90 percent to the optimists being right and only a 10 percent probability to Weitzman’s guesstimate of one chance in twenty. The intertheoretic probability is $0.9 \times 0.01 + 0.1 \times 0.05$ —fourteen chances in a thousand.⁶¹ But even with the my-favorite-theory approach, which would employ only the optimistic estimate, the expected loss from one chance in a hundred of wrecking the earth’s climate is enormous, unless we discount away the distant future.⁶² Can we justify doing so?

VI. AN ECOLOGICAL FALLACY

In a seminal 1995 paper, Schelling warned against a “fallacy of composition.” While future generations would probably be richer in the aggregate than the present one, this did not mean that action by rich countries to mitigate climate change would transfer resources from the poor to the rich. The future residents of poor countries would likely still be poorer than the residents of rich countries today.⁶³ In fact, Schelling was describ-

57. Toby Ord, Rafaela Hillerbrand, and Anders Sandberg, “Probing the Improbable: Methodological Challenges for Risks with Low Probabilities and High Stakes,” *Journal of Risk Research* 13 (2010): 191–205.

58. Ted Lockhart, *Moral Uncertainty and Its Consequences* (New York: Oxford University Press, 2000), 42; Sven Ove Hansson, “Economic (Ir)rationality in Risk Analysis,” *Economics and Philosophy* 22 (2006): 231–41, 233–36.

59. Dominic Roser, “The Irrelevance of the Risk-Uncertainty Distinction,” *Science and Engineering Ethics* 23 (2017): 1387–1407, 1400.

60. Here I borrow terms from Brian Hedden, “Does MITE Make Right? On Decision-Making under Normative Uncertainty,” in *Oxford Studies in Metaethics*, vol. 11, ed. Russ Shafer-Landau (Oxford: Oxford University Press, 2016), 11:102–28.

61. For a similar example, see Ord, Hillerbrand and Sandberg, “Probing the Improbable,” 203.

62. Mariano Torras, “Orthodox Economics and the Science of Climate Change,” *Monthly Review* 68 (2016): 25–34, 32.

63. Schelling, “Intergenerational Discounting,” 398.

ing an ecological fallacy, not a fallacy of composition.⁶⁴ Some members of future generations are likely to be much richer than the inhabitants of today's rich countries, but others will be poorer. Even if members of future generations are richer on average, this is no reason to discount benefits that would go to these worse-off people.⁶⁵

Ironically, Schelling went on to commit an ecological fallacy of his own. The real question, he argued, was not whether mitigation would be progressive in its effects—it would—but whether the rich could do as much good for the poor through mitigation as they would by devoting the same amount of resources to helping today's poor. "If GDP per capita continues to increase in most of the developing world," he wrote, "as I expect and as the optimization models assume, marginal utilities of the beneficiaries will be much higher during the first 50 years—before [greenhouse gas] abatement benefits become significant—than in the second 50 years," and still higher in the first years than the later ones. This being the case, it would be "logically absurd to ignore present needs and concentrate on the later decades of the coming century"—or, presumably, the centuries and millennia thereafter.⁶⁶

Would it? Consider a new version of *Mistake*:

Vaccination: A computer glitch leads the NHS to decide to relieve Bill's thousand days of severe pain rather than give one chocolate each to ten million healthy young adults. The chocolates contain an oral vaccine designed to confer immunity to a virus which is currently making the rounds. The virus has a 1/100 chance of mutating. If it does, it will kill all ten million people.

Bill might reason that he is very badly off and the others are expectedly well off, even after taking into account the small chance that they will all be killed by the virus. He might conclude that it would be unreasonable to expect him to undergo great pain for people who are expectedly much better off. That would be the wrong conclusion. If it were certain that the other people would be much better off, Bill could reasonably refuse. But while the others are expectedly better off, they might end up worse off. If they do, it will be cold comfort that they would have been better off than Bill in other possible states of the world. "Suppose that unbeknownst to me someone flips a coin," Hugh Lazenby writes, "with the expectation that if it falls on heads I will be made the beneficiary of a vast fortune.

64. The ecological fallacy is related to the fallacy of division—the assumption that what is true of the whole must be true of its parts. The fallacy of composition is to assume that what is true of the parts must be true of the whole.

65. Marc Fleurbaey and Stéphane Zuber, "Climate Policies Deserve a Negative Discount Rate," *Chicago Journal of International Law* 13 (2013): 565–95, 582.

66. Schelling, "Intergenerational Discounting," 399–400.

The coin falls on tails. What have I gained? I might have gained a lot. But I have actually gained nothing. I do not experience what I might have gained. I experience only what I have."⁶⁷ The same will be true for the ten million here.

Instead, Bill should disaggregate the scenarios that could transpire. In most possible states of the world, if Bill alerts the NHS, he will make a great sacrifice that brings many well-off people trivial benefits (a yummy chocolate). He can defensibly disregard these expected benefits. But in the one hundredth scenario, if Bill keeps mum, he will have allowed ten million deaths. Since there is one chance in a hundred that this scenario will materialize, this amounts to the expected loss of one hundred thousand lives. Bill has no excuse for disregarding these expected costs. If they materialize, he will be better off than the victims. Even if Bill can defensibly give more weight to his own interests, it could not be as great as this.

Nor can we justify a similar gamble with climate change. As Schelling advised, we must disaggregate future people into rich and poor. But we must also disaggregate them into different possible states of the world. When economists predict that despite climate change, future people will be richer than we are, they mean one of two things: either that they will be richer in the most likely scenario, or that they will be richer averaging across all scenarios. Since it is obviously foolish to consider only the most likely outcome, let us assume that they mean the latter. While in most of their scenarios future people are richer, in some proportion, they are poorer. If a bad scenario should materialize, that they were richer averaged across all scenarios will be irrelevant.⁶⁸ If there is only one chance in a hundred of leaving the world's population permanently impoverished, this still entails enormous expected impoverishment. The rich countries could bear the entire worldwide costs of stabilizing greenhouse gases without imposing nearly comparable sacrifices on most of their citizens. As Schelling observed, were mitigation to cost 2 percent of world GDP, "the doubled per capita income that might have been achieved by 2060 is reached in 2062. If someone could wave a wand and phase in, over a few years, a climate-mitigation program that depressed our GNP by 2 percent in perpetuity, no one would notice the difference."⁶⁹ By failing to lower their emissions, the inhabitants of the industrialized countries are taking a small chance of leaving an enormous number of people worse off than they are. There is no justification for discounting away these expected losses.

67. Hugh Lazenby, "Broome on Fairness and Lotteries," *Utilitas* 26 (2014): 331–45, 341.

68. Dominic Roser and Christian Seidel, *Climate Justice: An Introduction* (Abingdon, UK: Routledge, 2017), 77.

69. Thomas C. Schelling, "The Cost of Combating Global Warming: Facing the Trade-offs," *Foreign Affairs* 76 (1997): 8–14, 10; Bernward Gesang, *Klimaethik* (Berlin: Suhrkamp, 2011), 159–60.

VII. HOW TO DISCOUNT DEFENSIBLY

Not all treatments of discounting commit the ecological fallacy. Weitzman argues that in choosing among investments we should discount distant future effects at the rate we would use in the least optimistic scenario, since “all of the other states . . . are relatively much less important now because their present value has been reduced by the power of compound discounting at a higher rate.”⁷⁰ Yet with enough pure time preference, the overall discount rate might still be positive, making catastrophic future damages disappear. With any time preference that was short of absolute, on the other hand, we would be obliged to make the sacrifice in cases like *Wrinkle_{DS}*. Dale Jamieson worries that “there seems to be no plausible value for *delta* [the rate of pure time preference]. If the setting is 0 or extremely low, the interests of the present are swamped by the future. If *delta* is set high enough to protect the interests of present people, then the interests of those who will live in the further future lose virtually all significance.”⁷¹

This article seeks to square that circle with an approach to pure time discounting that complements Weitzman’s. Here is how it might be done. First, assess the impartial value of a policy’s costs and benefits in each growth scenario. When costs and benefits would go to people who will in any event be much better off, disregard them—no matter how many stand to gain—unless they are cost-free to provide. Weight the remaining costs and benefits by whatever degree of agent-centered preference we think appropriate. Finally, multiply each scenario’s weighted value by its estimated probability and sum them to obtain the policy’s expected value.⁷²

Under this approach, there would be no obligation to make the sacrifice in *Wrinkle* or *Wrinkle_{DS}*, since here the future people are by stipulation much better off than we are—richer, at any rate—and we can defensibly ignore the benefits that the sacrifice would bring. Suppose, however, that there is one chance in a hundred of our carbon emissions causing global warming that so damages the planet that average welfare remains significantly below that of the inhabitants of rich countries today. Assume also that this would cause a collapse in world population—so that only three billion people would live per century—and that the human race would survive another one hundred thousand years. A simple calculation

70. Weitzman, “Why the Far-Distant Future,” 205; for a helpful discussion, see Eric A. Posner and David Weisbach, *Climate Change Justice* (Princeton, NJ: Princeton University Press, 2010), 152–53.

71. Dale Jamieson, *Reason in a Dark Time: Why the Struggle against Climate Change Failed—and What It Means for Our Future* (Oxford: Oxford University Press, 2014), 128.

72. In their discussion of an “ex post” approach to discounting, Simon Dietz and Geir B. Asheim appear to be proposing something along these lines: “Climate Policy under Sustainable Discounted Utilitarianism,” *Journal of Environmental Economics and Management* 63 (2012): 321–35, 325–26.

shows that in just this scenario our emissions would entail the impoverishment of an expected thirty billion people.⁷³ Even if PR_{RC} or PR_{PPV} were to permit some pure time preference, they could not justify discounting such a catastrophe away. Low-probability catastrophic risks, some of which might be rational to run if only our own lives were at stake, become indefensible when they could permanently mar the planet.⁷⁴

It might seem that the problem of overdemandingness will reemerge. Nordhaus warns that if we were to attach infinite disutility to apocalyptic risks, “we would likely drown in a sea of anxiety at the prospect of the infinity of infinitely bad outcomes.” He suggests that with low discount rates it could be rational to fight preventive wars to prevent possible power shifts in the distant future.⁷⁵ While it is true that nearly anything we do could in principle prove catastrophic, in most cases we have no more reason to believe that acting is riskier than abstaining. It is often more dangerous, for example, to launch preventive wars than to forgo them. With climate change, on the other hand, we have a well-theorized causal mechanism, with good theoretical and empirical reasons for believing that higher greenhouse gas concentrations are more dangerous than lower ones.⁷⁶ There are a small number of other apocalyptic threats that fall into this category—notably thermonuclear war—but for now there are not scores or hundreds.⁷⁷ With technological development, the number is likely to grow. Yet in a nightmare world where such risks had proliferated, what could we reasonably do but exercise extreme caution? Perhaps the problem is not that giving world-wrecking threats their due is too restrictive, but that technological development is too dangerous. We should not take it for granted that the story of industrialization has a happy ending.⁷⁸ If century after century we run such risks, sooner or later we will lose the gamble.⁷⁹

73. $1/100 \times 100,000/100 \times 3$ billion = 30 billion.

74. Olivier Godard, “The Stern Review on the Economics of Climate Change: Contents, Insights and Assessment of the Critical Debate,” *Surveys and Perspectives Integrating Environment and Society* 1 (2008): 17–36, 30.

75. Nordhaus, “Review of the *Stern Review*,” 697; William D. Nordhaus, “The Economics of Tail Events with an Application to Climate Change,” *Review of Environmental Economics and Policy* 5 (2011): 240–57, 254.

76. Sven Ove Hansson, “Decision Making under Great Uncertainty,” *Philosophy of the Social Sciences* 26 (1996): 369–86, 378; Richard K. Betts, “Striking First: A History of Thankfully Lost Opportunities,” *Ethics and International Affairs* 17 (2003): 17–24; Henry Shue, “Deadly Delays, Saving Opportunities: Creating a More Dangerous World?,” in *Climate Ethics: Essential Readings*, ed. Stephen M. Gardiner et al. (Oxford: Oxford University Press, 2010), 146–62, 154.

77. Weitzman, “On Modeling,” 13–14.

78. See Dan Lyons, “Are Luddites Confused?,” *Inquiry* 22 (1979): 381–403.

79. Bostrom, “Existential Risk Prevention,” 25; Nassim Nicholas Taleb et al., “The Precautionary Principle (with Application to the Genetic Modification of Organisms),” Extreme Risk Initiative Working Paper Series (New York: New York University School of Engineering, 2014), <https://arxiv.org/abs/1410.5787>, 2.

Nordhaus argues that the willingness of contemporary societies to accept catastrophic risks such as those of nuclear winter or the destruction of the earth by a particle accelerator, as well as their unwillingness to take more than modest measures to ward off the threat of large asteroids, reveals that we attach limited weight to removing even the most apocalyptic threats.⁸⁰ If Nordhaus's aim were to describe the policy priorities of existing societies, no sensible person could disagree with him. But he seems to think that this tells us what we should do. Nordhaus would not, presumably, hold that other economic policies that states adopt—say, slashing social welfare spending—are justified merely because voters support them. To insist that the preferences of today's consumers determine what we should value is to accept that if the present generation desired massive future death and suffering, there would be no value to preventing it, and indeed value in bringing it about.⁸¹ Nordhaus does not claim that we should discount away damages to foreigners if we care about them less. What then is the justification for discounting away damages to distant future generations?⁸²

VIII. CONCLUSION

It is a mistake to assume that the future will be richer. Instead, we should recognize that there are different possible futures, in some of which our descendants are richer, and in others of which they are poorer. We can defensibly ignore costs and benefits that would require sacrifices on behalf of people who would be in any case much better off. That is not because these benefits would have no value. Rather, it is because we could not reasonably be expected to make the sacrifice needed to produce them. At the same time, cost-benefit analyses should give significant weight to the welfare of everyone else. Given the vast number of people likely to live in the future, this means averting any realistic danger of catastrophic damage to the planet.

80. Nordhaus, "Economics of Tail Events," 253–54.

81. For a related point, see Simon Caney, "Human Rights, Climate Change, and Discounting," *Environmental Politics* 17 (2008): 536–55, 547–48.

82. Jamieson, *Reason in a Dark Time*, 125–26.