

PROGRESSIVE CONSEQUENTIALISM

Dale Jamieson
Environmental Studies and Philosophy, New York University

Robert Elliot
University of the Sunshine Coast

1. Introduction

Consequentialism is the family of theories that holds that acts are morally right, wrong, or indifferent in virtue of their consequences. Less formally and more intuitively, right acts are those that produce good consequences. A consequentialist theory includes at least the following three elements: an account of the properties or states in virtue of which consequences make actions right, wrong, or indifferent; a deontic principle which specifies how or to what extent the properties or states must obtain in order for an action to be right, wrong, or indifferent; and finally, a specification of what is in the domain of the deontic principle. For example, mental state and desire theories provide different accounts of the first element; maximizing and satisficing are distinct deontic principles; and Act and Rule Consequentialism specify different domains over which a deontic principle ranges. A wide range of alternative theories can be generated by modifying these three elements. For example, Hedonistic Act Utilitarianism which requires that each act maximize pleasure, and Perfectionist Lifetime Minimalism which requires that each life satisfy some minimal standard of perfection are both varieties of Consequentialism. The conceptual space Consequentialism describes is vast, versions of Consequentialism vary radically in their plausibility, and few objections count against all versions of Consequentialism.¹

The most ubiquitous objection against Consequentialism generally and utilitarianism in particular is that it demands too much, requiring agents to act in ways that they cannot reasonably be required to act. Different objectors identify different kinds of demands as unreasonable. In some cases it is the sheer scale of Consequentialism's demands that is regarded as unreasonable. In other cases the concern is the iterative, unrelenting nature of the demands. In still other cases the objection is that adherence to Consequentialism would undermine

our integrity by alienating us from our “ground projects” or from those to whom we stand in special relationships.² The core of the demandingness objection in all of its forms is that Consequentialism fails as a moral theory because it implies that moral agents have classes of obligations that they do not have.

Several strategies have been employed to defend Consequentialism against this demandingness objection. The most fundamental is to deny that a theory’s demandingness constitutes any reason at all for rejecting the theory.³ Some argue that the willingness to give up an otherwise plausible theory because it seems too demanding simply reflects attitudes based on false principles or defective moral practices.⁴ Others hold that the demandingness objection, at least in its usual formulations, is unintelligible, incoherent or impotent.⁵

A second strategy is to adopt a “mixed” ethic in which the consequentialist component is tempered by some non-consequentialist components. This strategy has two main forms: one involves a lexical ordering in which a deontological component (for example) absolutely constrains the consequentialist component; the second allows that it is a matter of judgement as to how much to weigh each component in what circumstances.⁶

A third strategy is to muck around with the basic elements of a consequentialist theory in order to frame an account that does not generate overly demanding obligations. Each of the three basic elements provides opportunities for moderating Consequentialism’s demandingness. Philip Pettit and Michael Smith try to moderate Consequentialism’s demandingness by adopting a “global” view of the domain over which the deontic principle ranges.⁷ Some versions of Negative Consequentialism try to moderate the theory’s demands by taking pain reduction as the only property in virtue of which consequences make actions right, wrong, or indifferent.⁸ However, it is the second element, the deontic principle itself, that seems to provide the most promising site for shaping a version of Consequentialism that is not subject to the demandingness objection.⁹

Our purpose in this paper is to develop a response to the demandingness objection that is an example of this third strategy. Progressive Consequentialism, as we shall call this view, rejects maximization as Consequentialism’s most plausible deontic principle. However, Progressive Consequentialism does not suffer from the arbitrariness of Satisficing Consequentialism, a competing view that also rejects maximization. Progressive Consequentialism promises a way of reconciling what is compelling about Consequentialism with the core intuition that motivates the demandingness objection. It is also a view that seems broadly consonant with the attitudes of Bentham, Mill, and the utilitarian tradition. It is thus surprising that Progressive Consequentialism is a neglected option in moral theory.¹⁰

We shall begin by sketching Satisficing Consequentialism, the best known alternative to maximizing views, and point to some of its difficulties. We shall then sketch Progressive Consequentialism, showing how it deals with these difficulties, and suggest that it seems broadly consonant with the utilitarian tradition. We shall then discuss some challenges to Progressive Consequentialism, suggesting

some further specifications of the core view that may help to meet them. Finally, we shall draw some conclusions.

First, some preliminaries. We shall speak indifferently of acts and actions, while recognizing that there is a semantic difference between them that is important for some purposes.¹¹ Unless otherwise noted we shall speak of acts or actions in the generic sense in which obeying a rule or expressing a character trait (for example) can be said to be performing an action. For convenience, we stipulate that an act is right if and only if it is permissible for the agent to perform the act. Numerous other simplifications will be made along the way, but only the most important will be specifically noted.

2. Satisficing Consequentialism (SC)

SC and its attractions can be characterized by contrasting it to Utilitarianism. While Utilitarianism is the view that right acts are all and only those whose consequences are best, SC is the view that right acts are all and only those whose consequences are “good enough.” SC seeks to avoid the demandingness objection by counting a range of actions that produce consequences of different value to be right, in contrast to Utilitarianism which requires the identification and pursuit of a single best outcome or one of a set of outcomes that contain the highest achievable value.¹²

In dialing down the threshold of rightness, SC seems to enjoy intuitive support. Most of us think that people who have donated substantially to famine relief have done at least, if not more than, their duty, even if they haven’t given to the point at which further contribution would hurt themselves more than it would help others. And someone who volunteers six nights a week building houses for Habitat for Humanity, but spends the seventh night watching television, does not seem, on this ground, to be open to moral criticism. Most of us are inclined to think that these people have done enough. Indeed, many would say that they are morally exemplary. One of the advantages claimed on behalf of SC is that it maintains the distinction between the obligatory and supererogatory, a distinction that is difficult to maintain for a utilitarian. An agent who has done enough has discharged his moral obligations; an agent who has done more than enough has acted in a supererogatory way.

The case for SC cannot in the end rest on appeals to favourable examples. In some cases in which we find it morally acceptable for an agent to produce a non-optimal benefit, it is because there is some additional cost associated with producing the optimal benefit. In such cases doing what is “good enough” on a particular occasion may be part of an overall strategy of doing what is best. In cases in which it is clear that both the optimal and non-optimal benefits have the same price tag, it is not obvious that it is morally acceptable to produce the non-optimal benefit. Suppose, for example, that an agent could with the same degree of cost and certainty deliver either benefit N or benefit N+1. For any benefit

N, it seems implausible to suppose that the agent's act is of equal moral value whether she produces N or N+1. How could it be right in such circumstances, everything else being equal, for an agent to produce in a clear-minded way the inferior outcome?

The central problem for SC is that there does not seem to be a principled way of identifying what is "good enough". Any attempt to establish such a baseline appears ad hoc. The utilitarian principle that we ought to do what is best enjoys a clear, non-arbitrary rationale. A principle that says that we ought to bring about (say) 20% of what is best does not. Why not 19.9% or 20.1%? What is the magic of 20%?¹³

At its heart, SC is a response not to the demandingness objection but to what we might call the "onerousness concern." Since Consequentialism takes a global and long-term perspective in delivering its normative assessments, establishing that some action is right, all things considered, requires a serious attempt to assess its impact far from its geographical location and also into the further future. This leads to the problem of computational difficulty, which makes acting on Consequentialism onerous for agents.¹⁴

In response to this concern, Herbert Simon introduced the notion of satisficing as a decision-making heuristic.¹⁵ Since in everyday life there are information costs, transaction costs, and various uncertainties about ourselves, other people, and the world, we sometimes do best by settling for the first alternative that is "good enough" rather than holding out for what may otherwise appear to be the best alternative.¹⁶ But there is little reason to suppose that a principle that may be a good heuristic for decision-making is an adequate account of moral justification. As a principle of moral justification, SC is an ad hoc construction that is mainly deployed in an attempt to protect Consequentialism from the demandingness objection. However, there is little reason to suppose that there is some threshold which is good enough not to demand further action, regardless of other features of the agent and the world. In short, SC was inspired by a theory of rational choice that does not naturally lend itself to a theory of moral justification.

3. Progressive Consequentialism (PC)

As a first approximation we can say that PC is the view that a right action is one whose consequences improve the world. Thus, what PC requires of agents is that they act in such a way as to increase value in the world.¹⁷

PC and SC are similar in that both express obligation as a non-maximising function over value: inserting the relevant values specifies particular obligations. However, PC is not ad hoc in the way that SC is, since the improvement criterion is natural and clear in a way that an arbitrarily drawn line between minimal improvement and maximization is not. Maximization is also a natural and clear criterion but one that is strangely static and unworldly. It requires us to do what

is best all the time, whatever the condition of ourselves and the world. PC reflects a diachronic perspective on morality. Our mission as moral agents is to leave the world better than we found it. This struggle for improvement should be constant. The more we accomplish, the more that is demanded. Ourselves and others are held to ever higher standards as the world improves.

The vision expressed by PC seems more in tune with the work of the classical utilitarians than does any static view about the criterion of right action. For example in Chapter 1, Paragraph 11 of *On Liberty* Mill writes that he regards “utility as the ultimate appeal on all ethical questions; but it must be utility in the largest sense, grounded on the permanent interests of man as a progressive being.” In his book on *Auguste Comte and Positivism*, Mill praises Comte for recognizing the importance of moral progress. Mill goes on to write that “. . .the domain of moral duty in an improving society, is always widening. When what once was uncommon virtue becomes common virtue, it comes to be numbered among obligations, while a degree exceeding what has grown common, remains simply meritorious.”¹⁸ Without wishing to enter the thicket of Mill scholarship, we observe that, in these passages at least, Mill sounds more like a progressive consequentialist than a textbook utilitarian.¹⁹

PC can be thought of as describing a family of theories that fix the second element, the deontic principle, in a particular way, and it should be clear that variants of PC can be produced by supplying different interpretations of the first and third elements. Consider, for example, the following versions of PC: Hedonistic Act PC; Hedonistic Rule PC; Desire-Satisfaction Act PC; Desire-Satisfaction Rule PC. The list could go on. Different versions of PC will seem more or less plausible, and various objections will present more or less serious difficulties for different versions of the theory. In what follows, we shall focus on concerns that center on PC’s deontic principle, and largely ignore problems that stem from particular accounts of the other two elements.

4. Problems with PC

The first problem with PC can be brought out by returning to an objection we lodged against SC. Suppose that an agent could, with the same degree of cost and certainty, deliver either benefit N or benefit $N+1$. Since PC does not require us to do the best we can, but only to improve the world, PC does not require us to deliver $N+1$, so long as delivering N would mark an improvement with respect to the baseline.

Fortunately there seems to be an easy fix to this problem. If a person is willing to allocate a specific degree of effort to improve the world, then we might reasonably require that she use that degree of effort to produce the best result possible. Afterall, PC is motivated in part as a response to the demandingness objection, and delivering N is no more demanding than delivering $N+1$. Thus, PC might reasonably be taken to have an efficiency requirement built into it. The

upshot is that PC generates an obligation not merely to ensure that an action increases value, but also to ensure that no alternative action requiring the same degree of effort increases value more.²⁰

This response raises questions about how to determine the baseline against which improvement is assessed. One approach would be to say that the baseline is the value of the world as it is at the time of an action. For example, if the value of the world at time T1 is N, then in order for an action performed at time T1 to be right, the value of the world at time T2 must be equal to or greater than N+1 (where '1' represents the smallest realizable increment of value). An agent would thus be required efficiently to exert at least the minimum amount of effort that would produce this increment of value.²¹

Unfortunately, however, this account will not work. Suppose that, not yet taking into account the agent's action, due to natural occurrences or to the actions of others the world at time T2 will either be vastly better or vastly worse than at time T1. In the cheery case, the agent is required to do nothing — indeed, he is only required not to act in such a value-reducing way that he drives the value of the world down to its previous state. In the case in which the world catastrophically loses value, the agent's obligations would be incredibly demanding, for he would be required to move heaven and earth, if necessary and possible, in order to improve the world relative to its previous state.²²

One diagnosis of this problem would fault the "impersonal" nature of the version of PC under consideration. This impersonal view makes agents and their projects hostage to factors beyond their control, from natural disasters to the actions of other people, since it requires agents to produce a better world. A "personal" version of PC, on the other hand, would require only that an agent's actions directly produce more good than bad, regardless of the overall state of the world. One way of understanding such a personal view would be to fix the baseline, not as the value of the world at time T1 when the agent initiates her action, but as the value of the world at time T2 on the counterfactual assumption that the agent does not act at all. Thus, in the cheery case the agent is required to make the world even happier than it would otherwise be were she to "do nothing," and in the catastrophic case she must act in such a way that the world is less bad than it would be were she to do nothing. This view, in supposing that if the world is already bright one is required to add one's candle to the light, and that, if the world is dark, one is required to light one's candle against the darkness, seems to reflect PC's guiding intuition.

Unfortunately there are problems with this view. Any account that puts significant weight on the distinction between an agent acting and doing nothing will have some explaining to do. Consequentialists for good reasons have been hostile to views that suppose that actions can be clearly distinguished from omissions, refrainings, failures to act, and the like. The point that has mattered most to consequentialists is to deny that moral relevance attaches to an act, omission, or refraining simply in virtue of it being the kind of thing that it is.²³ Indeed, this claim is at the center of what distinguishes consequentialists (and

others) from deontologists. In addition, this version of the personal view produces results that would be unpalatable to a clear-minded consequentialist. Suppose, for example, that an agent is in a situation such that if he does nothing the world will improve, but if he does anything the world will become worse. This version of PC implies that an agent should perform that action which has the least bad consequences rather than the obviously correct response (to a consequentialist at least), that the agent should do nothing at all.²⁴

In response to such difficulties it is natural to turn back to an impersonal version of PC. Instead of fixing the baseline as the value of the world at time T2 on the counterfactual assumption that the agent does not act at all, we might suggest instead that the baseline is the value of the world at time T2 on the counterfactual assumption that that the agent does not exist.²⁵ This would free us from placing moral weight on categories of action and inaction, and would also give us resources for dealing with the second objection. If the baseline is the value of the world at time T2 were the agent not to exist, and if there is nothing that the agent can do to improve upon this baseline, then the agent should do nothing.

While this view of the baseline seems to handle this objection, it faces some difficulties of its own. Imagine a case in which, were an agent not to exist, the value of the world at time T2 would be N, but the agent's natural inclinations and dispositions would lead her to act in such a way that at time T2 the value of the world would be N+2. On this version of PC, an agent who suppresses her naturally benevolent dispositions and instead produces a world whose value is N+1 will have acted rightly. For in this case the PC threshold for right action is lower than what an agent would naturally do without the guidance of PC. This result may lead some back to a further exploration of a personal version of PC.

A version of a personal view that would seem to handle such a case is one which built into the baseline what an agent would do acting on her natural inclinations and dispositions. Thus we might say that the baseline from which improvement is assessed is the value of the world at time T2, which includes what the agent would have done following her own inclinations and dispositions. In a domain populated by sketchy accounts this one may try one's patience to excess. How do we know what an agent would do? What if the agent would naturally act on the basis of moral motivation? How can we distinguish moral motivation from dispositions and inclinations? Should we try to distinguish them? The most serious objection to this view, however, is that it risks making morality extremely demanding for the highly moral and extremely undemanding for the morally indifferent. If the value of the world at time T2 includes Ghandi's good works, then he will have to do something very good indeed to make the world even better. On the other hand if the consequences of Madoff's unscrupulous behaviour are built into the value of the world at time T2, then showing a little common decency would be enough for him to act rightly. Of course Ghandi may experience the opportunity to improve the world as a gift that does not require steely determination to carry out, while Madoff may have to summon all

of his determination and commitment to do the decent minimum, but this does not affect the basic point that morality demands much more of Ghandi than of Madoff. And while there may be some truth in the adage “to whom much is given, much is demanded,” this version of PC would seem to incorporate an extremely implausible version of it. In response, we might try to normalize the actions of different agents, and build into the baseline not the actions that individual agents would perform, but some reasonable expectation about how a generic agent would act. However, the search for the reasonable expectation risks both circularity and introducing an arbitrary standard into an account, one of whose strengths is the non-arbitrariness of its deontic principle. Indeed, it risks collapsing PC into a version of SC.

A better approach for a defender of PC is to stay the course with the impersonal view: then two kinds of plausible response can be given to the objection posed above. First, it could be claimed that this case runs afoul of the efficiency condition. If an agent’s natural inclinations would lead her to produce a world whose value is $N+2$, it seems plausible to suppose that it would take at least as much effort to suppress these inclinations in order to produce an inferior world. If this is the case, then the agent is required to produce a world whose value is $N+2$, since that would be the best result she could produce given the amount of effort she is willing to allocate. A second response would be to bite the bullet and say that all that morality demands is that the agent produces the world whose value is $N+1$, but like other cases of supererogation, it would be so much the better if the agent produced a still better world than she is obliged to do. So, a defender of PC should stick with the impersonal view that fixes the baseline as the value of the world at time T_2 on the counterfactual assumption that the agent does not exist.

5. Conclusion

In our view PC is a promising but neglected option within the family of consequentialist theories. Unlike SC, PC is not an ad hoc construction; maximization and improvement are both natural criteria for right actions. PC is also in the spirit of the classical utilitarian tradition, and helps us to recover the optimistic, progressivist dimension of the thought of Bentham and Mill.

Objections can of course be raised against PC, but it will be important to distinguish those which attach to a large range of consequentialist theories from those which are specific to PC. In some cases what may seem to be an objection may actually be viewed as an advantage of the theory. For example, in most cases, from the perspective of PC, there will be a wide range of permissible actions, reflecting the fact that there are a variety of different ways in which an agent can improve the world. Some may see this as a fault of the theory — as a failure to be sufficiently action-guiding and directive. This same feature also tends to relieve an agent from the excruciating trade-offs faced by a utilitarian. In the end, some

may think PC is not demanding enough, while others may think that it is too demanding. We may have to accept that from the perspective of any plausible consequentialist morality, our obligations are very strenuous by the standards of common sense, and that almost all of us fail to live up to them almost all of the time. The problem may not be with the theory, but with ourselves.²⁶

Notes

1. While there are many different accounts of Consequentialism, for further elucidation of the one sketched here see Dale Jamieson, *Ethics and the Environment: An Introduction* (Cambridge: Cambridge University Press, 2008), pp. 77–85.
2. For the most influential statement of the latter objection, see Bernard Williams, “A Critique of Utilitarianism,” in J.J.C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973). For general discussion, see Tim Mulgan, *The Demands of Consequentialism* (Oxford: Oxford University Press, 2001), Garrett Cullity, *The Moral Demands of Affluence* (Oxford: Oxford University Press, 2004), and especially Samuel Scheffler, *The Rejection of Consequentialism*, revised edition (New York: Oxford University Press, 1994).
3. Some (but not at least one of us) would qualify this claim by adding the provision, “at least so long as the principle that ‘ought’ implies ‘can’ is not violated.”
4. See, for example, R.M. Hare, “The Argument From Received Opinion,” in his *Essays on Philosophical Method* (Oxford: Oxford University Press, 1971), Peter Singer “Sidgwick and Reflective Equilibrium”, *Monist* 58 (1974): 490–517, and Peter Unger, *Living High and Letting Die* (New York: Oxford University Press, 1996).
5. For discussion of these responses see David Sobel, “The Impotence of the Demandingness Objection,” *Philosophers’ Imprint* 7.8 (September, 2007): 1–17, and Liam Murphy, *Moral Demands in Nonideal Theory* (New York: Oxford University Press, 2000).
6. There are many versions of such pluralist theories. For a particularly clear example see Jonathan Glover, *Causing Death and Saving Lives* (New York: Penguin Books, 1977).
7. “Global Consequentialism”, in Brad Hooker, Elinor Mason and Dale E. Miller, (eds.), *Morality, Rules, and Consequences: A Critical Reader* (Edinburgh: Edinburgh University Press, 2000), pp. 121–133. Holistic Consequentialism still faces the problem of deciding, on particular occasions, what to take as the evaluand, and this in itself could be quite onerous. A more straightforward approach to moderating Consequentialism’s demands through specifying a particular domain over which the deontic principle ranges is Brad Hooker’s Rule Consequentialism, developed in his *Ideal Code, Real World* (Oxford: Oxford University Press, 2000).
8. The origin of this view is usually located in Karl Popper, *The Open Society and its Enemies* (Princeton N.J.: Princeton University Press) Vol. I, Ch. 5, note 6. For criticism, see R. N. Smart, *Mind* 67.268 (Oct., 1958): 542–543. In any case, since a duty to maximally reduce pain would be extremely demanding it is not at all clear that limiting what is of value independently of restraining the deontic principle is sufficient for responding to the demandingness objection.

9. Frank Jackson has tried to defeat the demandingness objection by going beyond these three elements and exploiting the distinction between the actual and foreseeable consequences of our actions; see his "Decision Theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* 101: 461–482. For a response, see Michael Smith, "Consequentialism and the Nearest and Dearest Objection," forthcoming in Ian Ravenscroft (ed.), *Minds, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson* (Oxford: Oxford University Press).
10. So far as we know this view was first suggested in Robert Elliot, *Faking Nature: The Ethics of Environmental Restoration* (New York: Routledge, 1997), pp. 46–48, and further developed in an earlier version of this paper presented at International Society for Utilitarian Studies meeting in New Orleans in 1997. The views put forward in this paper have been referred to or discussed in Ben Bradley, "Against Satisficing Consequentialism," *Utilitas* 18.2 (June, 2006): 97–108, and in Walter Sinnott-Armstrong, "Consequentialism," *Stanford Encyclopedia of Philosophy*, available at <http://plato.stanford.edu/entries/consequentialism/>.
11. E.g., there are mental acts but no mental actions.
12. For an influential statement and defense of satisficing, see Michael Slote, *Common-sense Morality and Consequentialism* (London: Routledge & Kegan Paul, 1985). For further discussion, see the papers collected in Michael Byron (ed.), *Satisficing and Maximizing* (Cambridge: Cambridge University Press, 2004).
13. But cf. Alastair Norcross who thinks that any threshold for distinguishing right acts from others is arbitrary ("Reasons Without Demands: Rethinking Rightness," in James Dreier (ed.), *Contemporary Debates in Moral Theory*, Malden, MA: Blackwell, 2006).
14. James Lenman thinks this leads to more than computational difficulty. See his "Consequentialism and Cluelessness," *Philosophy and Public Affairs* 29 (2000): 342–370. It is an open question whether consequentialists can deal with some of the computational difficulties by adopting a subjectivist form of the theory, or whether there is any principled way of limiting the horizon of relevant consequences.
15. The original paper is his "A Behavioural Model of Rational Choice," *Quarterly Journal of Economics* 69 (1955). See also his *Models of Bounded Rationality* (Cambridge, MA: The MIT Press, 1982).
16. Or so it is claimed. But there are interesting questions as to how we could come to know that this is true. For discussion of this and related issues, see Dale Jamieson, "When Utilitarians Should be Virtue Theorists," *Utilitas* 19.2 (June, 2007): 160–183.
17. What it means to increase value will depend on, among other things, substantive value-theoretic commitments. In this paper we do not take sides regarding what is valuable or on the nature of the correct function for describing increases in value.
18. *Collected Works of John Stuart Mill, Volume X*, edited by J.M. Robson (Toronto: University of Toronto Press, 1969), p. 338. We thank Henry West for these references.
19. Both PC and SC seem to reflect a range of widely held moral intuitions that seem at odds with Utilitarianism. For example, the Lockean Proviso only requires us to leave "as much and as good" for others, not to use a resource in the best

- possible way; and sustainability is often characterized as the preservation of resources, opportunities, or welfare, rather than their maximization.
20. There are complications here that we shall ignore, including what exactly is meant by 'effort', and the relation between efforts that agents are required to make and those they are willing to make.
 21. There are some technical issues here that would have to be sorted out on a full account. For example, the fact that actions have consequences that both ramify into the future and have different values at different points in the future would have to be accounted for. Moreover, a total view of the value of the world would have different implications for action than other views, such as an average view. We shall ignore these and similar points for present purposes.
 22. The fact that cases in which the world is going to be worse whatever an agent does present some of the most difficult challenges for PC, may (bizarrely enough) provide some indirect evidence for this being the correct reading of the utilitarian tradition. Bentham and Mill were generally optimistic about the future, and so may not have considered the plausibility of their theory for a world in which things are getting worse. Perhaps optimistic PC is a more plausible view than pessimistic PC.
 23. See, for example, James Rachels' classic, "Active and Passive Euthanasia," *The New England Journal of Medicine* 292 (1975), and the papers collected in Bonnie Steinbock and Alastair Norcross (eds.), *Killing and Letting Die*, 2nd edition (New York: Fordham University Press, 1994).
 24. Similar cases are discussed by Bradley, *op. cit.* and the literature cited therein.
 25. Since agents leave marks on other people and the world in various ways throughout their lifetimes, the case in which the agent never existed will be very different from the case in which the agent existed prior to time T2. What we are imagining is that the world bears the marks of the agent, but that the agent does not exist in the baseline, nor are there any particular effects engendered by the agent dying, popping out of existence, or otherwise not being present at time T2.
 26. Earlier versions of this paper were presented in 2000 at the International Society for Utilitarian Studies at Wake Forest University, at Ohio University in 2009, and in graduate seminars at New York University in 2007 and 2009. We thank all those who participated in these discussions, and Douglas Portmore for his comments on an earlier draft.