

**Responsibility and the Demands of Morality**  
Stephen J. White

In thinking about what morality requires of us—especially in terms of beneficence—we seem to be pulled in different directions. On the one hand, we are not egoists. We recognize that the needs and interests of other people matter—indeed, that at some level they matter just as much as our own needs and interests. Hence we are not indifferent to the misery and suffering of other people, nor to the ways in which we might help them. On the other hand, there is the sense that if morality required us to aid those worse off than ourselves whenever possible, it would require too much. And it is a common objection to moral theories that incorporate extremely demanding requirements of beneficence (utilitarianism being the prime example) that they are too demanding to be plausible. This is the so-called *demandingness objection*.

In this paper, I take up the issue of how we should understand this “demandingness” objection, and whether it is a good one. Several questions arise. One is the question of what it means for a theory to be demanding in the first place. What, in other words, is the nature of the demands we are concerned with and that, intuitively, we think must be limited? This is the main question I attempt to address in this paper. How we answer it will make a difference to how we go on to address two other important questions. First, what are the underlying philosophical grounds for thinking that morality is itself sensitive to the demands individuals face in virtue of the obligations they have to others? And second, what are the possibilities for incorporating limits on what morality can demand into moral theory?

A second criticism of utilitarianism and other forms of (act-)consequentialism often goes

hand-in-hand with the demandingness objection. This is that such theories imply that we are responsible for outcomes we would not normally take ourselves to be responsible for. Act-consequentialist views imply that we are just as responsible for what we fail to do or prevent as we are for what we actively bring about. Some philosophers argue that this is in itself unacceptable, or at least counterintuitive.

But the two criticisms are also often seen as connected in the following way. The doctrine of “negative responsibility” associated with (e.g.,) utilitarianism explains why it has the potential to take such a toll on the lives of individual moral agents.<sup>1</sup>

One of my central aims here is to argue for a different connection between the demandingness of certain moral theories and their implications concerning responsibility. Rather than understanding the doctrine of responsibility as being, so to speak, the cause of utilitarianism’s demandingness, where this is in turn understood in terms of its potential to require large sacrifices on the part of moral agents, I will instead suggest that the expansive doctrine of responsibility provides a better interpretation of the sense in which the theory is extremely demanding. The welfare costs that are imposed on those who act as the theory requires are a red herring. If there is a way in which utilitarianism is too demanding, it is that it holds us responsible for too much.

To make the case that this is the best way to understand the demandingness objection, I first argue against the high-cost-of-compliance interpretation of demandingness (§2). What I hope to show is not so much that utilitarianism, for instance, remains plausible despite the objection, but that once we probe this interpretation of demandingness, it loses its intuitive force as an objection to begin with. I will then argue that it does not help to shift from a concern with

---

<sup>1</sup> See, especially, Bernard Williams, “Critique of Utilitarianism,” in J. Smart and B. Williams, *Utilitarianism: For and Against*, (Cambridge: Cambridge University Press, 1973).

the *degree* of sacrifice that may be required to a concern with the *fairness* of increasing the morally required level of sacrifice in conditions partial compliance (§3). To articulate a clear objection, we will need to turn our attention to something other than the welfare costs associated with fulfilling one's moral obligations. The rest of the paper then defends a reinterpretation of the objection. I first make precise the sense in which utilitarianism and views like it may be thought to have extreme implications for moral responsibility (§4). Second, having clarified the relevant notion of responsibility, I attempt to make intuitive sense of a way in which it can be demanding or burdensome that does not reduce to welfare costs associated with fulfilling one's moral obligations (§5). Finally, I spell out some implications of this way of understanding the objection (§6), before concluding (§7).

A brief note of clarification before continuing. I will generally be discussing the demandingness objection as it applies to utilitarianism—indeed, a fairly simple and unsophisticated form of act-utilitarianism. This is not because I take utilitarianism to be the only or the most plausible theory that is potentially very demanding. Rather, it is because it has generally been treated as a theory paradigmatically subject to the demandingness objection. Since I'm interested in whether it ever makes sense to object to a moral theory on the grounds that it is excessively demanding (and, not primarily, whether we should accept or reject any particular theory), it seems wise to raise the question in relation to the theory most likely to draw such an objection.

### §1. The Demandingness Objection: Its Basic Form

Let's begin with a couple of cases meant to provoke the intuition that utilitarianism is overly demanding. Imagine, first, a relatively affluent person—your typical middle-class American, for

instance. She earns a decent salary and donates a sizable portion of it to worthwhile charities. But she is unhappy with her career and has always dreamed of being a novelist. She has enough savings to make it feasible for her to quit her job and make a go at a writing career. However, we can plausibly suppose the utilitarian calculation recommends against it. She can contribute more to the total welfare if she remains in her hateful job and gives away most of her salary. In fact, we can further suppose the utilitarian would have it that what she is morally required to do is work longer hours and take fewer vacations and give up seeing first-run movies and live music, etc., etc. At this point we may be inclined to ask: Is it really reasonable to ask this person to give up so much for the sake of helping others? Shouldn't she be allowed some space to live her own life?

David Sobel offers a different sort of case that is meant to provoke a similar intuition.<sup>2</sup> He tells us about Sally, who needs a kidney transplant in order to live. Joe has two healthy kidneys and is the only person present who is a match for Sally. If Joe doesn't give Sally one of his kidneys, she will die. Sobel assumes that, in this situation, Joe will maximize aggregate welfare by giving one of his kidneys to Sally (who, we might as well add, is a total stranger to Joe). According to utilitarianism, then, Joe is morally required to do so. But again, we may feel that this demands too much of Joe. It seems unreasonable to suppose that, in this situation, it would be wrong of Joe to keep both of his kidneys.<sup>3</sup>

Now I think that in these cases utilitarianism does seem to demand more than we feel is really required of the individuals involved. But what exactly is the problem? We need both an interpretation that provides some way of establishing that a theory is very demanding, as well as some explanation of why this should matter to us in our thinking about morality.

---

<sup>2</sup> David Sobel, "The Impotence of the Demandingness Objection," *Philosophers' Imprint* 7 (2007), 1-17.

<sup>3</sup> Another way to make the same point would be to say that, if Joe were to give Sally his kidney, his action would be supererogatory. We would naturally credit him with doing far more than simply meet his basic obligations.

Donating an organ to a stranger—even if she needs it to live—is surely to go above and beyond the call of duty.

Here we need to draw a distinction between two forms the demandingness objection can take. Often the two are run together in ways that lead to confusion. First, there is what we might call the “mismatch” form of the objection. On this version, the sense in which utilitarianism, for instance, is overly demanding is that it implies that people are morally required to do things for others—make certain sacrifices—that they are not, in fact, morally required to do. The intuition here is simply that the theory does not match moral reality. Thus, for example, the utilitarian conclusion that Joe is required to give his kidney to Sally, may conflict with our ordinary, commonsense conviction that no one is obligated to donate an organ to a total stranger (at least while one is still alive). The mismatch version of the demandingness objection is just an instance of a familiar form of argument to the effect that the theory under consideration is at odds with our considered first-order judgments concerning what is and is not morally required.<sup>4</sup> This form of criticism is fine as far as it goes; but it is not the form of the objection I am going to focus on here.

The fact (if it is a fact) that there is a mismatch between what morality actually requires and what utilitarianism says it requires, with respect to, say, organ-donation, does not in itself give us any clue as to *why* morality does not require that we donate our organs to strangers who need them. This is not so with regard to the second form of the demandingness objection, which we might label “demandingness-as-independent-concern.” On this version of the objection, we treat it as a desideratum, or *prima facie* criterion of adequacy, for a moral theory that it not be extremely demanding. This is a more general and abstract conviction about morality itself—that it must not be too demanding (in some sense of “demanding”). Understood in this way, the intuition in Sobel’s example is not merely that Joe is permitted to keep his kidney, contrary to the

---

<sup>4</sup> Cf. Rawls’s discussion of reflective equilibrium and the associated method of moral argumentation in *A Theory of Justice*, (Cambridge, MA: Harvard University Press, 1971).

utilitarian verdict. It also purports to be an intuition about what explains *why* Joe is permitted to keep his kidney—namely, that a morality that *required* Joe to donate his kidney in these circumstances would be overly demanding.

An analogy may be helpful. Utilitarianism is not only criticized for being too demanding. It is also criticized for being, in certain respects, too permissive. It may be thought to permit, for example, doing serious harm to one person if this will bring small benefits to a large number of other people. If this constitutes a good objection to utilitarianism, however, it is clearly of the “mismatch” variety. The relevant objection is that utilitarianism implies that a certain act is permitted, though we are convinced it is not. But the intuition is not that we are in general concerned to avoid excessive permissiveness in our moral outlooks. (Not to mention that no one would accuse utilitarianism of being overly permissive in this general way, since nearly every permissible act is for the utilitarian also obligatory). The demandingness objection seems different, however. There seems to be an independent concern or principle behind the objection, in a way that there does not when it comes to the “permissiveness” objection. It is this demandingness-as-independent-concern form of the objection I propose to investigate.

Once we distinguish the demandingness objection proper from the mismatch version, the need for some account of the measure of demandingness becomes salient. We cannot simply point to examples like those given above.

The most straightforward and familiar way of measuring the demands associated with a particular theory would seem to be in terms of the costs imposed on individuals who comply with the theory's requirements.<sup>5</sup> In the next section, I will raise several objections to this way of framing the objection.

---

<sup>5</sup> This is the concern taken up, for example, in most of the essays found in a recent volume on “the problem of moral demandingness.” See Timothy Chappell (ed.), *The Problem of Moral Demandingness: New Philosophical Essays* (Palgrave Macmillan, 2009). An interesting exception to this is the contribution to that volume written by Onora O’Neill.

## §2. The High-Cost-of-Compliance Interpretation

On the current understanding, a theory will be seen as excessively demanding if, under normal circumstances, an agent must sacrifice a great deal of her own well-being in order to act as the theory says she ought to act. In Sobel's example of Joe and Sally, and in my example of the hopeful novelist, it is natural to think that the welfare costs imposed by compliance with what utilitarianism would require are simply too great. It is not reasonable to demand that one give up one's dreams or one's organs, even if doing so would benefit others a great deal. Or so the objection goes.<sup>6</sup>

Is it a good objection to a theory that compliance with it is extremely costly in terms of agents' overall well-being? I think it is not. There are three outstanding questions about how the high welfare-costs associated with complying with a given theory could constitute a real objection to that theory. I will argue that the difficulty of finding satisfactory answers to these questions should motivate us to consider alternative interpretations of the objection.

The first question is whether the demandingness objection applies only to the claims a moral theory makes about our positive duties to aid others and to promote their well-being, or whether it also applies to other kinds of moral requirements—requirements against assault, theft, murder, and so on. Compliance with these negative duties—prohibitions against the use of violence, for example—may after all be very costly for a person.<sup>7</sup> Perhaps one's best opportunity to extract oneself from very unhappy circumstances is to threaten an innocent person's life, or

---

<sup>6</sup> Of course, one line of response offered by utilitarians is that utilitarianism is not as demanding as it may initially seem to be. For two classic examples of this strategy, see Henry Sidgwick, *The Methods of Ethics*, (Hackett, 1966), and Peter Railton, "Alienation, Consequentialism, and the Demands of Morality," *Philosophy & Public Affairs* 13 (1984), 134-171. But, since I'm interested in what sort of objection it would be to a theory that it is extremely demanding, I will assume for the sake of argument that utilitarianism is extremely demanding.

<sup>7</sup> For versions of this point, see Shelly Kagan, *The Limits of Morality*, (Oxford: Clarendon, 1991), Liam Murphy, *Moral Demands in Nonideal Theory* (Oxford: Oxford University Press, 2000), and Sarah Stroud, "They Can't Take That Away from Me," in M. Timmons (ed.) *Oxford Studies in Normative Ethics*, vol. 3. (Oxford: Oxford University Press, 2014).

falsely accuse someone of a crime.<sup>8</sup> And yet, the fact that forbidding such acts may require real sacrifice seems not to raise any serious objection. The concern that moral requirements not be excessively demanding—in the sense of imposing very high welfare costs on compliant agents—appears, therefore, not to be a concern that arises for just any moral requirement. Rather it is the demandingness of beneficence in particular that we object to. And yet this restriction on the force of the objection remains mysterious.

The second question concerns our narrow focus on the costs incurred by compliant agents. If we are evaluating the acceptability of a theory in terms of its effects on persons' well-being, what justifies us in ignoring, for instance, the benefits the poor and desperate will receive through the compliance of those better off than they are? Why shouldn't we think it equally objectionable to ask those who are very badly off to accept that, as a matter of moral principle, their misery and suffering and poverty do not generate a claim on others—who are already much better off—to make very significant sacrifices in order to help them out of their desperate circumstances? After all, theories that do not require very much in terms of beneficence will thereby permit people to refrain from helping others who may very much need the help. However, if we take into account the costs that a theory permits as well as those that it requires, then we are no longer in a position to claim that utilitarianism is especially demanding when compared with other moral views.

In order to vindicate the demandingness objection then, we need some justification for discounting the burdens that the less fortunate would be left to bear were we to adopt a less stringent principle of beneficence.<sup>9</sup> Obviously, one cannot reply at *this* point that it would

---

<sup>8</sup> For a compelling discussion of a real-life case that illustrates this, see Stroud, "They Can't Take That Away from Me."

<sup>9</sup> Sobel argues that if there is an argument establishing a morally significant distinction between the costs a theory requires and the costs it permits, then this argument will itself already imply that consequentialism is false. If he is right about this, then the demandingness objection can serve to support non-consequentialist moral views over



demand too much of individuals to insist that they consider the needs and interests of those they might aid as being on a par with their own. For there must be some independent argument against this if we are to get the demandingness objection off the ground in the first place.

One response here is that the focus on the costs for compliant agents is justified by a more basic concern with what is motivationally realistic for most people. Thus, in a recent paper, Brian McElwee argues that what justifies treating the costs to agents differently from the costs to “patients” is that the former are self-imposed and thus engage the will of those subject to them.<sup>10</sup> Any theory that requires people to impose large burdens on themselves will, McElwee thinks, be beyond the motivational reach of most people. And this provides a basis for skepticism that does not apply to a moral theory which simply allows that people may be left to bear large burdens they have in no way chosen.

According to this way of developing the objection, then, the fundamental complaint is not so much that a theory like act-utilitarianism fails to take sufficient account of the costs borne by compliant agents. (To this the utilitarian will reply that this will only seem to be so from a standpoint that fails to take sufficient account of the burdens that would thereby be prevented from befalling others.) It is rather that utilitarianism fails to take sufficient account of “what humans find it easy or difficult to be motivated to do.”<sup>11</sup>

The engagement of agents’ motivational resources clearly marks an important difference between the burdens a theory requires agents to take on and the burdens a theory permits people

---

consequentialist views only if it is based on some other argument that already provides sufficient grounds for rejecting consequentialism. See Sobel, “The Impotence of the Demandingness Objection.”

<sup>10</sup> Brian McElwee, “Demandingness Objections in Ethics,” *Philosophical Quarterly* (forthcoming). See also, Samuel Scheffler, *Human Morality*, (Oxford: Oxford University Press, 1992).

<sup>11</sup> McElwee, “Demandingness Objections,” p. 16. After drawing a connection between moral obligation and the appropriateness of blaming a person, he continues:

...our judgments of what is morally obligatory must likewise take into account standard patterns of human motivation... In considering when it is appropriate to feel blame toward such agents [hence when it is right to hold that they are under a moral obligation], we must evaluate the quality of their willing relative to natural or common patterns of human willing and motivation (16-17).

to suffer, in the sense that it includes no duty to prevent them. There are serious problems, however, with tying moral obligations too closely to the motivational tendencies of ordinary people. The issues here run deep and I will not be able to do full justice to them. But I will highlight three areas of concern about this approach.

- 1) Whether a person is motivated to behave in a certain way—especially when it involves making sacrifices—often depends on whether she is convinced that it is the right thing to do (that it would be wrong to behave otherwise). Attempting to determine whether one has a duty to act in a certain way by asking how motivated one is to do so would therefore seem to be putting the cart before the horse. A similar problem arises if, in working out a theory of moral obligation, one is asking how far the average person is willing to behave in a certain way. In general, it is not clear that the answer will be all that significant if the common absence of motivation corresponds to doubts about or outright disbelief concerning whether there is a duty to do the kind of thing in question.
- 2) There is reason to be skeptical of a view that has our moral obligations depend on our actual motivational patterns and tendencies, given that those tendencies themselves depend on a variety of factors which may properly be called into question in the course of our moral theorizing. What we find ourselves easily motivated to do depends on our moral education, on the culture in which we're brought up, on the economic, political, and social institutions and practices that structure our lives. But, arguably, these things should not constitute *constraints* on a theory of right and wrong—features of moral life that must be taken as given. Instead, we might reasonably think, to the extent that these factors establish patterns of motivation at odds with what our best moral theories imply we ought to do, they should be objects

of criticism.

- 3) There is something independently counterintuitive about the idea that we can rule out, *a priori*, the possibility of widespread moral failure, except where accounted for by widespread moral ignorance. But this seems to follow if we accept the motivational constraint.<sup>12</sup> This is not just counterintuitive in the abstract. There are many concrete instances where we recognize common motivational impediments, but are nevertheless convinced that there are genuine wrongs being committed. We know, for instance, that it can be very difficult for people to refuse the directives of those in positions of authority, even when those directives are profoundly evil. Do we really want to say that there is no obligation to disobey in such cases? Or again, most of us have experienced how hard it can be not to lie and dissemble in situations where, for example, we have been caught doing something we should not have been doing. But should this by itself lead us to doubt that, in such situations, our duty is to tell the truth? In general, it is hard to believe that there are no contexts or conditions under which doing the right thing will be very hard.

These concerns should, I think, raise serious doubts about whether our intuition that utilitarianism is too demanding is best understood as grounded in the view that moral obligations generally must be such that the average person can fairly easily bring herself to satisfy them.

This brings us now to the third question about the high-cost-of-compliance interpretation. Let's grant, for the sake of argument, that there *is* something especially problematic about the costs agents are required to self-impose, according to a particular moral theory. The question is whether there should therefore be a *limit* to what morality demands.

---

<sup>12</sup> And if we commit to the idea that where there is widespread failure to live up to some moral obligation, this is necessarily due to ignorance, this would seem to make the first problem—about the dependence of motivation on moral conviction—that much more pressing.

I suggest we face a dilemma here. Suppose that there is a strict limit to what any reasonable moral view will require people to do for others. This implies that what one is morally required to do for others is insensitive, at least beyond a certain point, to how much people actually need and how much one could do to help given the circumstances. This is an odd result. The demandingness objection loses much of its plausibility if it forces us to accept that how dire others' straits are could make *no* difference to what it is reasonable to demand of those in better circumstances.

On the other hand, if we adopt a view that attempts to make the obligations of beneficence sensitive to the situations of those we are required to help, this will fundamentally alter the nature of the demandingness objection in a way that also robs it of its intuitive force. Suppose we do think that what we are morally required to do for others, by way of providing aid or preventing harm, and thus what we are required to sacrifice of ourselves, depends on how badly off others are—or will be without our intervention. Our sense that a particular theory is too demanding will now need to take a specific form. We can no longer object that this or that theory simply asks too much—that is, requires complying agents to sacrifice more than it is reasonable for a person to have to sacrifice. Rather, the objection will have to be that a theory requires a sacrifice that is disproportionate to the benefit provided. In other words, a theory will appear overly demanding inasmuch as it requires individuals to make sacrifices that do not seem worth the resulting benefits to others.

This form of the objection, however, has considerably less appeal. It requires us to make certain comparative assessments of gains versus losses that lack a good deal of the plausibility of the initial idea that a theory like utilitarianism is too demanding to be correct. Should one be permitted to buy a new pair of shoes, or go to first-run movies, when one could instead contribute that money to help someone get life-saving medicine they would otherwise not have

access to? Perhaps we think that morality would demand too much if it never permitted one to buy a new pair of shoes for oneself. But if we reject a strict limit on what can reasonably be demanded of us in favor of more flexibility, we will have to construe the objection differently: it would be too much to require a person to forgo up-to-date fashion merely in order to help people get the medicine they need to live. The objection would have to be construed in this way because, having rejected the idea of a limit, we must be prepared to say that in some circumstances a person should *not* be permitted to spend her money on new shoes, given the good she could otherwise do with that money. Thus, our objection must amount to the judgment that, in the actual circumstances, the benefit to others her money could provide (e.g., the medicine) is not sufficient to offset the sacrifice of living without decent footwear (or without first-run movies, or beach vacations). I think for most of us, such judgments hold little intuitive appeal.<sup>13</sup>

To summarize, we seem to lack, first, a way of explaining why the objection applies specifically to principles that concern the duty of beneficence and not other moral principles. Second, it seems we do not have any very convincing explanation as to why the costs that persons are made to bear because of what a theory requires of them should matter more to us than those costs that persons are made to bear because of what a theory permits others to do. And finally, it seems that the demandingness objection remains compelling only as long as we do not specify whether or not it implies that there should be a limit to what morality can demand of us.

### §3. Murphy's Argument

---

<sup>13</sup> Much of the force of Peter Singer's famous article, "Famine, Affluence, and Morality," for instance, stems from the way he focuses the reader's attention on such comparative assessments of the worth of what might be sacrificed and what might be gained. The principle he defends and elaborates in that article will surely strike most of us as excessively demanding. And yet the sorts of comparative assessments of worth that Singer suggests we are committed to by our current practices do seem hard to justify. See Peter Singer, "Famine, Affluence, and Morality," *Philosophy & Public Affairs* 1 (1972), 229-243.

The problems discussed in the previous section do not, at least for me, dislodge the intuition that a moral theory or principle that makes extreme demands is for that reason objectionable. I suggest, therefore, that we should look for a different interpretation of the objection. I think the arguments I've rehearsed indicate that if a moral view like utilitarianism is objectionable for being overly demanding, this is not simply because the cost of complying with its requirements is likely to be very high.

In this section, I discuss a different sort of objection to the costs associated with theories like utilitarianism, the most fully developed version of which has been offered by Liam Murphy.<sup>14</sup> According to Murphy, what is problematic about the utilitarian principle of beneficence (what Murphy calls the “optimizing principle”) is not that compliance with it would be very costly for most of us given our present circumstances. Rather, the problem is due the difference in the levels of sacrifice the principle requires in conditions of partial compliance versus conditions of full compliance. Murphy’s objection is to the fact that the demands the optimizing principle makes on those who comply with it increase as overall compliance with the principle goes down. That is, in circumstances (like those at present) in which not everyone engages in a full time effort to promote aggregate well-being, those who do will have to do more and take on greater burdens than they would have to if everyone complied with the principle. This, Murphy says, is unfair. One should not be assigned a greater responsibility for promoting overall well-being simply because other people fail to act as they should. In general, Murphy thinks, a moral principle should not increase the demands it makes on agents merely because, over all, compliance with that principle has decreased.<sup>15</sup>

---

<sup>14</sup> Murphy, *Moral Demands in Nonideal Theory*. Cf. David Miller, “Taking Up the Slack: Responsibility and Justice in Situations of Partial Compliance,” in C. Knight and Z. Stemplowska (eds.), *Responsibility and Distributive Justice*, (Oxford: Oxford University Press, 2011).

<sup>15</sup> This is what Murphy calls the “compliance condition.” See *ibid.*, p. 77. For the official version of the condition,

Murphy's argument for this “compliance condition” can be put as follows.

1) If the total effects of compliance with a principle of beneficence are fairly distributed in circumstances of *full* compliance, then the allocation of individual responsibility for promoting well-being in those circumstances is fair.

2) One's fair share of responsibility for promoting well-being just is the share of responsibility assigned to one in a fair allocation of individual responsibility under full compliance.

3) One should not be required to take on more than one's fair share of responsibility for promoting well-being.

4) By (1) and (2), any principle of beneficence that allocates responsibility fairly in conditions of full compliance but increases its demands on individuals in conditions of partial compliance will require people to take on more than their fair share of responsibility under partial compliance.

5) So a principle must not demand more of individuals in circumstances of partial compliance than it does in circumstances of full compliance.

In evaluating this argument, one might wonder, first, why we should think that what constitutes a fair share of responsibility in conditions of full compliance should imply anything about what counts as a fair share in other situations (premise 2). But I won't press this question. The question I want to raise is why, if we accept premise 2, we should accept premise 3. That is, why shouldn't one be required to take on more than one's fair share in situations of partial compliance?

The answer may seem obvious. One shouldn't be required to take on more than one's fair

---

see p 85.

share of responsibility because to require this would be unfair. But this is not necessarily so. It may be that, if one is required to take on more than one's fair share of responsibility for something, this implies some unfairness somewhere. It need not imply that the requirement itself is, in the circumstances, unfair. Someone who has to do more than her fair share of a burdensome task, because others have failed to do their part, no doubt has a legitimate complaint. But such a complaint would normally be directed at those who failed to pull their weight, not at the rule or normative principle dictating how she ought to proceed, given others' irresponsibility.<sup>16</sup>

In determining whether a requirement of beneficence, in particular, unfairly distributes costs to compliant agents, we cannot neglect the fact that alternative principles may simply shift these burdens onto others, in a way that will ground similar charges of unfairness. If there is less than full compliance with the principle, then it is a reasonable assumption that some people will be less well-off than they would have been under full compliance. One possibility is that some people will not be helped in ways they should have been. Another possibility is that some will have to undertake greater sacrifices in order to help those who need it. A principle of beneficence that meets Murphy's compliance condition and does not increase its demands in conditions of partial compliance will, it seems, allow these additional burdens to fall disproportionately on those who need the help rather than requiring others who are better positioned to make sacrifices in order to provide that help. It appears, then, that principles that meet the compliance condition may themselves be subject to a charge of unfairness. In any case, this casts doubt on the third premise of Murphy's argument. That is, we should not think that a principle that requires persons to do more than their fair share is for that reason unacceptable.<sup>17</sup>

---

<sup>16</sup> Cf. Anja Karnein . "Putting Fairness in Its Place: Why There is a Duty to Take up the Slack." *The Journal of Philosophy* 111 (2014), 593-607. Karnein draws a distinction between the obligations we owe to third parties and the obligations we owe to one another as fellow bearers of duties to provide aid. The issue of fairness, she argues, is relevant to the latter issue, but not the former.

<sup>17</sup> Murphy offers a response to a related objection on p. 92 of *Moral Demands*: He claims that the objection



Murphy, in my view, has not adequately explained why the demands of the optimizing principle appear absurd. If there is an objection to a principle that increases its demands as compliance decreases, it's doubtful that it is grounded in concerns about fairness.

#### §4. Responsibility, Answerability and Demandingness

If the problem with utilitarianism (or the optimizing principle) is not a matter of the sacrifices it insists people may be required to make, does this show that our initial intuitions were off-base—that there is no real problem here at all? We should not be too quick to draw this conclusion. In thinking about the limits of what can be morally demanded of us, what we appear to be responding to is the peculiar significance for a person of her own life and well-being. But this does not by itself imply our primary concern is simply with the sacrifices, in terms of welfare, agents may be required to make. We have already seen that it is difficult to make sense of our intuitions about demandingness—in regard both to degree, as well as fairness—on the assumption that the demandingness of a theory is captured by the costs of complying with it. I want to suggest, therefore, an alternative that seems to me plausible and that avoids the problems of trying to understand demandingness in terms of the sacrifices agents may be obliged to make.

Consider first the link between moral obligation and the justification of action. If, for instance, one adopts a straightforward understanding of utilitarianism as a guiding moral principle, one will then see one's actions as morally justified only if one can expect the consequences of those actions to be better, in terms of promoting overall well-being, than the

---

"assimilates a concern with the fairness of the way a principle of beneficence imposes responsibility on agents to a general concern about the fairness of the distribution of well-being." He goes on to explain that "though the collective principle of beneficence [his preferred principle] leaves the victims of noncompliance worse-off than they would be if the compliers took up (some of) the slack, it cannot be said that the victims have been required to take on (either actively or passively) responsibilities that rightly belong to others" (p. 92). However, this does not answer the objection in the text, which is not about whether the "victims" are required to take on more than their fair share of responsibility. Rather, it concerns the unfairness to the victims of not requiring those who are better off to take on more responsibilities.

consequences of anything else one might have done instead. Philosophers have sometimes claimed that this view about the grounds of our obligations has unacceptable implications for what we can rightly be held responsible for.<sup>18</sup> My interest here is to develop this claim in connection with the demandingness objection in the hopes of shedding light on what exactly is supposed to be troubling about this aspect of utilitarianism.

One might be wondering how a theory like utilitarianism, which is a substantive view of what makes actions right or wrong, could tell us anything by itself about the conditions under which a person is responsible for what she does.<sup>19</sup> My task in this section will be, first, to explain the specific ways in which a first-order theory of what we are morally obligated to do can have implications concerning when a person is responsible for something. Second, I hope to show that the implications of utilitarianism, in particular, are extreme in this regard. I will then turn, in section 5, to argue that this is a way in which utilitarianism is intuitively overly demanding.

The relevant notion of responsibility is a matter of what one is *answerable* for. Consider, for example, Angela Smith's account of this sense of responsibility:

...to say that 'A is morally responsible for X' is to say that A is an intelligible target of requests to justify X and that she is eligible, in principle, for a range of moral responses (from positive and negative appraisal, to the reactive attitudes, to attitudinal and behavioral expressions of praise and blame) based on the nature of the thing in question and the quality of the reasons she could give in response to these justificatory requests.<sup>20</sup>

---

<sup>18</sup> The classic discussion of utilitarianism's doctrine of "negative" responsibility in general is in Williams, "Critique of Utilitarianism." See note 29, below.

<sup>19</sup> My thanks to an anonymous reviewer for pressing this question in an especially clear way.

<sup>20</sup> Angela Smith "Responsibility as Answerability." *Inquiry*, 58 (2015), 99-126. Compare Pamela Hieronymi, "Reflection and Responsibility." *Philosophy & Public Affairs* 42 (2014), 3-41: "One is answerable just in case a request for one's reasons is given application" (p.12). Smith and Hieronymi both argue that answerability is the fundamental notion relevant to our moral responsibility practices. For arguments against this latter claim, see Gary Watson, "Two Faces of Responsibility," in *Agency and Answerability* (Oxford: Oxford University Press, 2004), and "The Trouble with Psychopaths," in R. Jay Wallace, Rahul Kumar, and Samuel Freeman, eds., *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon* (Oxford: Oxford University Press, 2012); and David Shoemaker, "Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility." *Ethics* 121 (2011), 602-632. I will not take a stand on this debate here, since even those who think that there are senses of responsibility, corresponding to different aspects of our responsibility practices, that cannot be grounded in answerability, generally acknowledge that answerability is a genuine and important concept of responsibility. I will also set aside the question of whether, in addition to one's actions (and states of affairs one can affect via one's actions) one can be answerable, and thus responsible, for attitudes, e.g., beliefs

A person is thus answerable for her conduct if the person can be expected to account for or justify her conduct by citing the considerations that counted in favor of it (and how they stood up in the face of considerations that counted against it). But we need to be careful about how we understand what exactly one is answerable for in a given context. It is a familiar point from the philosophy of action, going back to Anscombe and Davidson, that only some of the ways in which we might accurately describe what a person is doing at a certain time will be descriptions under which what she is doing is intentional.<sup>21</sup> The same point applies to what a person is answerable for. That is, a person is answerable—the intelligible target of requests for reasons—under some descriptions of what she has (or has not) done, but not others. This is not to say that the descriptions under which one is answerable for something are just the same as the descriptions under which it was intended. Sometimes one is responsible for what one does recklessly or negligently, (and therefore unintentionally). In such cases the problem may be that one lacks reasons altogether, where reasons are required.<sup>22</sup>

With this point in mind, we can begin to make sense of how a theory of what our moral obligations are can bear on what a moral agent is responsible for. The reason is that the descriptions under which a person is answerable for what she does in a given situation will often depend, in part, on what is morally expected of that person in that situation.

To illustrate: We normally think that lifeguards, in virtue of their role, have obligations to rescue struggling swimmers that average beachgoers do not have. And it seems that, because of

---

and desires, which are not directly subject to one's voluntary control. On this question, see Angela Smith, "Responsibility for Attitudes: Activity and Passivity in Mental Life," *Ethics*, 115 (2005): 236-71; and Hieronymus, "Reflection and Responsibility."

<sup>21</sup> G.E.M. Anscombe, *Intention*, 2<sup>nd</sup> ed. (Cambridge, MA: Harvard University Press, 2000); Donald Davidson, *Essays on Actions and Events*, (Oxford: Oxford University Press, 1980).

<sup>22</sup> Here, I follow Hieronymi, who holds that the sense in which a request for one's reasons is "given application" does not imply that one must have any reasons (even bad reasons) to offer in reply. See "Reflection and Responsibility," p. 13. Cf. Graham Hubbs, "Answerability without Answers," *Journal of Ethics and Social Philosophy* 7 (2012).

this, a lifeguard who declines to rescue someone who is drowning is responsible for the death in a way that others who were present are not. Certainly, this seems true if the reason the person wasn't rescued was that nobody was paying attention. Here, the death can be attributed to the lifeguard's failure to pay attention, but not to the others'.<sup>23</sup> Though he might not have intended it, because he has certain duties in his role as lifeguard, the person's death is, at least in part, his fault. To borrow a phrase from Frances Kamm, we can lay it at his doorstep.<sup>24</sup>

In this case, we assume, the lifeguard has a general duty to ensure swimmers' safety, in light of which we expect him to prevent this kind of thing from happening. But here we have to be careful—for it's not quite right to say that we expect lifeguards on duty to rescue any drowning person on their watch. It is more accurate to say that, with respect to events of this kind, we expect the lifeguard to either prevent it *or* have an adequate reason or justification for not taking steps to prevent it. After all, if he were busy rescuing some other person at the time, we would not say that he had failed to do what was expected of him as a lifeguard. Thus, given our assumptions, the person's drowning is something the lifeguard can be asked to account for—he is subject to a demand for explanation as to why he allowed or risked allowing it to happen. And because, in the case we imagined, he has no adequate reasons to give, all else equal, he is appropriately blamed for the person's death. Things are different, however, with the ordinary beachgoers. It's not that they acted on *good* reasons in not rescuing the person—they were not paying attention either. It's that they are not subject to the same expectations, and are therefore not subject to the same demand for justification. Although we might describe the beachgoers as “failing to prevent the drowning,” or, perhaps, as “relaxing on the beach *rather than* pulling the swimmer to safety,” these are not descriptions of their conduct under which they

---

<sup>23</sup> If you disagree with this judgment, ask yourself whether this is because you believe that an ordinary beach-goer does have some obligation to be on alert for swimmers in trouble. If so, then the link between obligation and responsibility is sustained.

<sup>24</sup> Frances Kamm, “Responsibility and Collaboration.” *Philosophy & public affairs* 28 (1999), 169-204.

are answerable. The ordinary beachgoers stand in the same causal and intentional relations to the person's drowning as the lifeguard does. But given the assumption that different obligations apply to them,<sup>25</sup> they are not responsible for that event as the lifeguard is.

Consider, now, what utilitarianism—or at least certain versions of it—implies about our responsibility for various states of affairs. In order for one's conduct to be justified in utilitarian terms, one must be able to justify a course of action by explaining how the considerations that counted in favor of it defeated all the considerations counting against it, where, specifically, the latter category includes any welfare gains that might have been had by doing something else.<sup>26</sup> Utilitarianism thus imposes on agents a kind of responsibility which (recalling Williams) we can label “negative answerability.” A person is negatively answerable for conduct insofar as she is answerable for what she does not do. And she is negatively answerable for a state of affairs insofar as she is answerable for not doing what she could have done to alter that state of affairs. Given the content of the utilitarian or optimizing principle, then, and the resulting structure of justification that is required of a person's actions, the range of descriptions under which one may be held to account for one's conduct will be enormous. For it will include all the ways in which one might have, but did not, intervene to promote someone's welfare (at least given that it was feasible for one to have discovered the opportunity to do so). And this implies that the utilitarian should see herself as being responsible, to an enormous extent, for how other people are faring in their lives.<sup>27</sup>

---

<sup>25</sup> This is not, of course, an assumption I have argued for, though it seems commonsensical enough. Its role in the argument is merely to illustrate the way in which assumptions about what obligations people have can have implications for what they are answerable for.

<sup>26</sup> Of course, in the specific case of utilitarianism, the relevant considerations will simply be the ways in which one's conduct either enhances or detracts from aggregate well-being. Other theories will include other considerations as relevant to the justification of conduct.

<sup>27</sup> Given the emphasis on a kind of “negative responsibility,” it may be useful to say something about how I understand the relation between Williams's argument and the demandingness objection. In his “Critique of Utilitarianism,” Williams objects to the utilitarian doctrine of negative responsibility, on the grounds that it undermines the integrity of the moral agent. Williams is sometimes read as claiming that there are certain

## §5. Too Much Responsibility?

My suggestion, then, is that we should understand the demandingness objection in terms of the notion of responsibility understood specifically as answerability. The point, of course, is not to deny that we have, and take ourselves to have, significant responsibilities regarding others' welfare. But I suspect that the intuitions underlying the demandingness objection reflect our sense that there must be limits and, further, that a person's answerability for her own well-being is more extensive than what she may appropriately be held to answer for when it comes to the lives of others. If this is right, then there is an important respect in which a theory will seem overly demanding if it insists that we take into account, in deciding what to do and how to live our lives, all the various ways in which we might act to enhance the well-being of other people. My claim is that the type of responsibility that such a theory would require us to take for other people would itself constitute a significant burden—one not reducible to ordinary welfare costs

---

sacrifices that it is simply not reasonable to expect an agent to make—viz., the sacrifice of what he later called the agent's ground projects (Bernard Williams, "Persons, Character, and Morality," in *Moral Luck*. (Cambridge: Cambridge University Press, 1981). On this reading, he would be pressing a particular version of the cost-of-compliance interpretation of the demandingness objection, and, as such, I believe it faces the problems I raised in section 2.

I read Williams's integrity objection differently, however. As I understand it, the problem is that the utilitarian agent will, in general, be forced to take incompatible views on the value of her own projects and commitments. As a utilitarian, she is supposed to base her decisions solely on consideration of aggregate human happiness. And this requires her to view her own projects and commitments as items whose importance consists in their status as preferences or desires whose satisfaction contributes to her (and perhaps others') happiness. But, for most of us, when it comes to the most significant projects we pursue in our lives, we do not see their value as consisting solely in their being objects of satisfaction to us. Indeed, if we did conceive of our projects like this, many would cease to be meaningful to us in the way required to make any significant contribution to our happiness. But if our projects' value to us lies in something other than their contribution to our happiness, then, Williams claims, an agent who simply ignored this or pretended otherwise in actually deciding what to do would lack integrity. On this interpretation of Williams's critique, the problem is not fundamentally about excessive *demands* placed on moral agents. Rather, the problem is with utilitarianism's exclusive focus on human happiness (understood roughly as desire- or preference-satisfaction) as its sole object of pursuit. This accounts for what Williams sees as the resulting tendency for what is deep in a person's life to be converted into something shallow when viewed from a utilitarian moral perspective. Though important and insightful, this argument is somewhat to the side of my primary concern. For on this reading it is not an objection that will directly apply to theories which, though they may be very demanding, do not take aggregate welfare, understood as satisfaction, to be the sole moral aim.

that might accrue from performing one's obligations.<sup>28</sup> It is a burden we should be sensitive to in formulating a theory of what morality requires.

Why should we accept this claim? We can, I think, distinguish two different aspects of the burden of responsibility associated with a theory like utilitarianism. The first is what we might label the *crushing moral pressure* aspect. The thing to note here is how extraordinarily high the standard of justification is for the most routine actions on the act-utilitarian picture, given its implications for negative answerability. In order for a person to adequately defend her behavior—in any context—she must in principle be able to defend her choice on the grounds that there was nothing more she could have done at the time to contribute to aggregate well-being. To think that it is morally incumbent on us to subject everything we do to this kind of scrutiny is, I think, to place a heavy moral burden on ourselves—even if we set aside the likelihood that living up to such a standard would require significant sacrifice. In light of this, we need to take seriously the possibility that what we are tracking in objecting to the demandingness of utilitarianism is the intuition that being subject to such a high standard of justification for everything we do is itself an excessive burden to bear—one that is distinct from the material sacrifices that might be required to meet that standard.

Indeed, this way of cashing out the basic intuition seems more plausible when applied to many of the contexts in which the demandingness objection seems to have some pull. If, for instance, we think that a theory is too demanding if it implies that, in all likelihood, one is never permitted to buy a new pair of shoes for oneself, or go to first-run movies, I doubt this is because we think there is something deeply important, even from one's "personal point of view," about

---

<sup>28</sup> In saying this, I do not mean to commit to a theory of well-being that rules out the type of moral burden I am discussing as a component of a person's well-being. The point is to distinguish this sort of burden from the sacrifices an agent might make in order to comply with her moral obligations. If the present "burden of responsibility" falls under the heading of well-being, in some broad sense, its impact is via the mere subjection to obligations of a certain sort—not via the things one might give up to comply with them. I am grateful to an anonymous reviewer for pressing me to clarify this issue.

having new shoes and seeing movies (why would we have thought that?). It is more natural to think that, in considering whether, say, to buy a pair of shoes, one is not always required to take into account everything else one might do with that money instead and compare all of these alternatives to the new-shoes option. On this view, it's the insistence that we conceive of what we are doing as *buying shoes as opposed to contributing to famine relief* (for instance)—and that we must justify ourselves in these terms—that we find suspect. Though we are not permitted simply to ignore it, we are also not always required, at every moment, to treat all of the suffering and need in the world as something for which we are responsible, and which can be attributed to us insofar as the potential existed for us to reduce it (or so we tend to think).

The second, and perhaps more significant aspect of this type of burden is what we might label the *ownership of misery* aspect. I think most of us will, once we get it into focus, recognize that the moral relation of being responsible for another person's misery or suffering constitutes a distinctive kind of moral burden. Of course, we don't like to be blamed for things, or to feel guilty; but it's not just that. The point is related to the phenomenon Bernard Williams termed "agent-regret."<sup>29</sup> This is regret, not just over what has occurred, but over one's contribution to what occurred through the exercise of one's agency. And as Williams point out, agent-regret in this sense clearly ranges over objects for which one is not to blame, morally speaking. Indeed, he first introduces the notion of agent-regret in discussing its appropriateness even in cases where one's agential contribution to the regrettable outcome is just brute bad luck. But there is also another kind of case, which is perhaps more relevant here. This is the case where, knowing there will be harmful consequences, one nevertheless intentionally does or refrains from doing something, but for good reason.

Consider, for example, a doctor who has to refuse her patient a spot in the ICU, thereby

---

<sup>29</sup> Bernard Williams, . "Moral Luck," in *Moral Luck*.



putting him at greater risk. There is a shortage of beds, and she judges that another patient has a greater need for the spot available. Everyone involved, including the doctor herself, may be agreed that this was the right decision to make. And yet it would be reasonable for the doctor to acknowledge that a deterioration in her first patient's condition may be partly attributable to her decision. There is a sense in which she is morally responsible—answerable—for the person's worsening condition. Assuming she had good reasons for making the triage decision in the way she did, she is not blameworthy. Still, agent-regret is a natural and appropriate reaction for the doctor to have. Her relation to her patient, who may now be facing death, is of ethical significance for her, and is something to regret, even if it's not true that she should have acted differently.<sup>30</sup> Indeed, we might expect the doctor to experience her responsibility in this regard as one of the great burdens of her chosen career.

This is enough to suggest that we are not indifferent to what we think of ourselves as responsible for. And the concern does not seem to be mainly about what we might have to do or sacrifice as a result. Recall the hopeful novelist I introduced at the beginning of the paper. This case was meant to elicit the intuition that a theory like utilitarianism is overly demanding. We can now see more clearly what the problem is. It is not that she would not get to pursue her writing career. The problem is that it would be incumbent on her to justify her career choice in light of all the ways alternative career paths might have benefitted others. Even if, all things considered, she were permitted to pursue the career she finds most satisfying, she must see her satisfaction as playing the role of supplying a justifying reason for others to be worse off than they might be. What I have tried to bring out is that this is itself a demanding relationship to stand in to other people, quite apart from any sacrifices one might have to make as a result. It

---

<sup>30</sup> Here it matters that regretting the way in which your agency contributed to a particular outcome does not require wishing that you had not acted as you did. See *ibid.*, p. 31. Cf. Amelie Rorty, "Agent Regret," in A.O. Rorty (ed.) *Explaining Emotions* (University of California Press, 1980).

should be no surprise, therefore, that we find ourselves resistant to—even repelled by—moral theories that demand we take responsibility every situation we might have done something to alter or improve, but did not.

This is, at any rate, how I propose to understand the thought that a strict utilitarian morality would be too demanding. I have not attempted to vindicate the objection based on this intuition. Perhaps it can be overcome. In any case, we would certainly want to say more about the source and legitimacy of the intuition itself. But in order to evaluate the objection, we need to understand the sense of demandingness that should be our proper concern. I have argued that cost-of-compliance interpretations do not allow us to do so much as formulate an intuitively plausible objection. In the remainder of this section, I will show that the responsibility interpretation does not give rise to the same problems.

First, if the demandingness of a moral theory is a matter of its implications for what we can be held responsible or answerable for, it is easy to see why it is only positive duties of aid and beneficence that appear problematic. Negative duties—such as the duty not to harm others—normally imply that one must take responsibility for the (foreseeable) harmful consequences of one's action. If one acts in a way that harms others, one needs some special reason for doing so—a reason sufficient to override the reasons not to cause harm. But such answerability for harms applies only to the limited range of cases in which one foreseeably causes harm through one's actions. Thus, while compliance with negative duties may, in certain circumstances, be very costly, such negative duties do not impose the level responsibility for others' welfare that very demanding positive duties do.

Second, this account makes it clear why we are focused specifically on the burdens associated with obligations rather than with permissions. Different principles of beneficence may have differential effects on the well-being of those who are better off versus those who are less

well off, depending on the degree to which such principles allow people to pay special attention to their own needs and interests. But the burdens of responsibility associated with a given principle of beneficence will not vary in this way according to individual circumstances. Even a very unfortunate person, whose welfare would likely improve under general compliance with the utilitarian principle, would nevertheless have to bear an extremely demanding level of responsibility for others' welfare. For she would also be required to take responsibility for any feasible, alternative pattern of action and benefit she might have effected as compared to what she in fact did.

Finally, because the sense of demandingness at issue on the proposed interpretation does not concern cost or degree of sacrifice, the question of whether there should be a strict limit to the welfare costs that a person can be morally required to take on is simply not relevant. If I am right, the problem of demandingness cannot be solved by incorporating into our moral theory limits on the costs a person can be required to bear for the sake of benefiting others. I will expand on this point in the next section.

#### §6. An Implication of the Account

Suppose we accept that morality must not be extremely demanding. What are the possibilities for incorporating this fact into a sound moral theory? The interpretation of moral demandingness I have been developing has important implications regarding this question. In particular, one especially influential way of dealing with intuitions about demandingness—the incorporation of an “agent-centered prerogative”—turns out to be inadequate.

An agent-centered prerogative is conceived of as a permission to give more weight to one's own aims and interests than to those of other people in determining the best course of

action to take.<sup>31</sup> There are limits, of course, to how much more weight one is allowed to give one's own interests. But if one exercises this option, one may justify acting in a way that would be better for oneself, as opposed to doing what would be best, considered wholly impartially.

If the problem of what morality demands were simply a problem with having to sacrifice one's own interests for the sake of the common good, positing an agent-centered prerogative would be a plausible response. Such a prerogative would allow one to sacrifice less than one would otherwise have to.

However, if our central concern should not be the costs agents must bear, but the level of responsibility they must assume for how others' lives are going, then having a prerogative to give more weight to one's own interests will not help. This is because such a prerogative does nothing to reduce the extent to which one may be held to answer for others' well-being. One would still be required to justify oneself by explaining how acting in one's own interest is preferable to acting in ways that would have various benefits for others. It is just that, in providing the explanation, one is allowed to cite the fact that one's interests are one's own, and thus have more weight for one than the interests and needs of other people. But one must, nonetheless, see oneself as responsible (if not to blame) for countless people being worse off than they otherwise would have been, given the variety of things one could have done, but chose not to. Insofar as we think that a person's distinctive relation to her own life and well-being is a matter of what she is specially responsible for, granting that we have an agent-centered prerogative to favor ourselves simply misses the point.

## §7. Conclusion

---

<sup>31</sup> Samuel Scheffler, *The Rejection of Consequentialism*, Rev. ed. (Oxford: Oxford University Press, 1994).

I have attempted to offer a plausible account of the intuitions underlying the demandingness objection, one that does not reduce to a concern over the costs associated with meeting our obligations. The idea of responsibility as negative answerability gives us a way to articulate our sense that a viable moral theory cannot be overly demanding on individual agents in a way that is coherent and appears not to beg any questions. My aim in this paper, however, has been somewhat limited. I have attempted only to spell out and render intelligible our commonsense intuitions and reactions. Beyond this, I do not claim to have vindicated these intuitions. Nevertheless, it seems to me a useful step toward a deeper understanding of the special significance our own lives have for us and how this impacts our moral relations with other people.<sup>32</sup>

---

<sup>32</sup> For extensive discussion and feedback on this paper, I owe thanks, first and foremost, to Barbara Herman, AJ Julius, and Pamela Hieronymi. For helpful discussion along the way, thanks also to Mark Alznauer, Arudra Burra, Lee-Ann Chae, Jorah Dannenberg, Kyla Ebels-Duggan, David Ebrey, Sam Fleischacker, Jeff Helmreich, Louis-Philippe Hodgson, Rob Hughes, Brian Hutler, Tom Hurka, Kristin Inglis, Aaron James, Herb Morris, David Plunkett, Karl Schafer, Kieran Setiya, Steve Yalowitz, and two anonymous reviewers for *Journal of Moral Philosophy*.