

Moral Limits on the Demands of Beneficence?

Richard J. Arneson

Forthcoming in *The Ethics of Assistance: Morality, Affluence, and the Distant Needy*, ed, Deen K. Chatterjee
(Cambridge: Cambridge University Press)

If you came upon a small child drowning in a pond, you ought to save the child even at considerable cost and risk to yourself. In 1972 Peter Singer observed that inhabitants of affluent industrialized societies stand in exactly the same relationship to the millions of poor inhabitants of poor undeveloped societies that you would stand to the small child drowning in the example just given. Given that you ought to help the drowning child, by parity of reasoning we ought to help the impoverished needy persons around the globe. To capture this intuition Singer proposed this principle of benevolence: If one can prevent some significant bad from occurring, without sacrificing anything of comparable moral importance, one ought morally to do so.¹ Premature death caused by preventable disease, injury, and poverty is uncontroversially a significant bad. Donations to charitable organizations such as Oxfam can prevent many of these deaths around the world, so Singer's principle holds that we ought to donate (or take some action that is comparably efficient at saving lives).

But this principle of benevolence is far more stringent than common-sense opinion, for even after one has donated most of one's income each month to world poverty relief, one could still donate more, and should do so according to the principle. For after all, the further reduction in one's available spending money does not incur anything that is comparable in badness to the loss that occurs to those in need of charitable relief if one's extra monthly donation is not forthcoming. The Singer Principle thus entails that one should continue to give until the point at which the marginal value of

the next bit of money one might give would do equal good as famine relief and as an increment to one's available spending money. Very few inhabitants of affluent industrialized societies today act as though they believe that the morality of benevolence is anywhere near this demanding. Hence the puzzle: What distinguishes the case of the drowning child nearby from the case of the distant starving stranger?

This essay considers three strategies of argument that aim to set moral limits on the requirements of beneficence. In the end none of these strategies proves viable, so the Singer Principle is embraced. But of course from the fact that some arguments for rejecting a moral principle are unsuccessful it does not follow that the principle is correct. My aim is limited to showing that it is harder than one might think to translate our intuitive revulsion against some implications of broadly consequentialist principles of beneficence such as the Singer Principle into a theoretically satisfying rationale for rejecting the lot. Still, the revulsion persists. Toward the end of this essay I attempt to lessen the counterintuitive sting of consequentialist beneficence principles by distinguishing between what one morally ought to do and what one is morally obligated to do and should be punished in some way for failing to do.

The Singer Principle is a close cousin to act-consequentialism, the doctrine that one morally ought always to do an act whose consequences are not less good than those of any available alternative. The difference is that whereas act-consequentialism denies the existence both of deontological constraints and of morally permitted options, the Singer Principle denies options but not constraints. It would be consistent with the Singer Principle to deny that one should provide aid to many distant strangers whose lives are at risk if the only effective way to provide help would involve violating a moral

constraint against deliberately inflicting harm on some innocent by-stander. In this essay I want to examine the merits of act-consequentialism while bracketing the issue of constraints. This amounts to an assessment of the Singer Principle.

The worry that consequentialism is too demanding proceeds from the assumption that the scope of moral demands is limited and that while respecting these demands an individual is (except perhaps in unusual extreme circumstances) free to lead her life as she chooses in any of a wide variety of ways. But if morality requires that one always should do what will produce the best, one is only free to choose among those actions whose outcomes are tied for best. The moral freedom to live as one pleases drastically shrinks on this conception of morality.

Actually we should distinguish several different ways in which a morality might be thought to impose requirements on conduct that are too demanding. One is that the morality unduly limits moral freedom as just characterized. Here are some others:

2. The morality imposes requirements that in many circumstances will turn out to be strongly opposed by the desires and inclinations that agents come to have, so it is psychologically difficult and burdensome to comply with these requirements.

3. The morality imposes requirements that are strongly in conflict with what is in the self-interest of the agent. Any morality will impose requirements that conflict with the course of action that would be to the maximal advantage of the agent in many circumstances and with a course of action necessary to avoid disaster for the agent in extreme circumstances. But a morality will seem excessively demanding if its requirements standardly require an agent to act in ways that are enormously to her disadvantage.

4. The morality imposes requirements that give no special weight to the personal concerns of the agent, the concerns that especially matter to that agent, with the result that doing what morality requires standardly and enormously conflicts with doing what would promote the agent's personal concerns. Here we suppose that an agent's personal concerns might not be exclusively self-interested, so demandingness 3 is not the same as demandingness 4.

When we characterize a morality as excessively demanding, we are at least implicitly comparing it to another morality that seems more plausible in part because it is less demanding either in the terms of limits on moral freedom, opposition to inclinations, conflict with self-interest, or conflict with personal concerns.

FAIRNESS AS CONSTRAINT ON REQUIRED BENEVOLENCE

Liam Murphy explores the interesting suggestion that what is most problematic about the consequentialist account of the moral wrongness of failing to provide aid when doing so would produce the best outcome is that the requirement of beneficence that falls on the individual may increase in virtue of the failure of others to do their part. To avoid this result, which he finds unpalatable, Murphy introduces a "Compliance Condition: a principle of beneficence should not increase its demands on agents as expected compliance with the principle by other agents decreases."² This Compliance Condition is satisfied by a principle of beneficence that fixes the upper bound of what is required of any individual by the cost the individual would have to undergo if there were full compliance by all agents with the principle in question. In this spirit Murphy proposes a Cooperative Principle of beneficence:

Each agent is required to act optimally—to perform the action that makes the

outcome best—except in situations of partial compliance with this principle.

In situations of partial compliance it is permissible to act optimally, but the sacrifice each agent is required to make is limited to the level of sacrifice that would be optimal if the situation were one of full compliance; of the actions that require no more than this level of sacrifice, agents are required to perform the action that makes the outcome best.³

Murphy's Cooperative Principle gives a straightforward response to the problem of distinguishing the case of helping a drowning child close at hand and providing lifesaving aid to needy persons threatened with premature death around the globe. We assume I face the drowning child alone, so one should do the optimal,⁴ helping act. But many people are in a position to help the global needy, and according to the Cooperative Principle, the requirements on an individual to offer aid do not increase because others are not doing their share. One is required to give aid at the level that would produce the best outcome if everyone complied with moral requirements, no more and no less.⁵

Notice first of all that Murphy's Cooperative Principle renders a person's obligations of beneficence constant in the face of changes in the ratio of the cost she must suffer if she gives aid to the benefit that others will gain if the aid is forthcoming (so long as the required act is optimal in consequences). Suppose a famine is in progress and can be interrupted, with a consequent reduction in the misery of the famine victims that just barely outweighs the losses that aid-givers must incur to produce this result. If famine relief, though barely worth its costs, yields the best outcome that can be reached, and an optimal arrangement is set in place, with everyone doing his part, then the agent in turn must do her assigned part in the aid scheme, even if the good she would do by her

contribution just barely exceeds the good she could do by going to the movies instead. In this scenario, the Cooperative Principle requires the agent to contribute to the famine relief scheme.

In contrast, consider a case in which others are not doing their part in what would be the optimal arrangement for bringing about the best outcome, but the cost to the agent/gain to beneficiaries ratio is extremely favorable if the agent provides aid above the level of sacrifice she would have to incur if all had done their part. Suppose a thousand people have fallen off a ship and are in danger of drowning. One hundred persons including the agent are standing on the deck and could save lives by tossing life preservers to the drowning. If the hundred were to divide the lifesaving effort in an optimal fashion, each would incur a trivial cost, say of one dime's worth of inconvenience, and 1000 lives would be saved. Unfortunately, apart from the agent, none of the hundred persons standing on the deck is fulfilling the obligation to aid, but nonetheless the agent could undertake the entire lifesaving operation herself, undergoing let us say eleven dollar's worth of inconvenience to save 1000 lives. Here act consequentialism yields the plausible verdict that the agent should undertake the entire rescue operation and save 1000 lives at small cost of inconvenience to herself. Murphy's Cooperative Principle yields the distinctly implausible verdict that the agent is required to do no more than act in an optimal way to save lives up to the level of sacrifice she would have sustained had there been full compliance. The agent under the Cooperative Principle is then required to undertake a dime's worth of sacrifice that saves ten lives, leaving 990 persons to perish who could have been saved at very modest cost to the agent.

This result to my mind shows that the Cooperative Principle is unacceptable. The mere fact of noncompliance by some does not automatically set an upper limit on the amount of sacrifice it is reasonable to demand of others who can provide cost-effective aid.

My route to this conclusion may be wayward. At any rate, there is a flaw in the example of the 1000 lives at risk that I used to illustrate the unpalatable implications of the Cooperative Principle. An example of a case in which some can rescue others near at hand at severe risk of death or injury tends to elicit from ordinary common sense the judgment that the obligation to effect easy rescue is stringent. But rescue cases elicit the response that help is obligatory even from people who do not think there is any stringent duty to provide assistance to needy distant strangers. The salience of the imagined fact that in rescue cases the person in distress is about to perish in one's sight and that one could easily prevent this catastrophe distorts ordinary judgment. If one redescribes the case so that physical proximity does not seem to put the person at risk and the potential rescuer into an I-Thou relationship, the strong judgment that aid must be forthcoming tends to evaporate.⁶

My response is that noting this psychological proclivity regarding easy rescue does nothing to diminish the unattractiveness of the Cooperative Principle. Let it be specified that in the 1000 person example described three paragraphs back, potential rescuers and rescuees are temporally and physically distant from each other. Let the personal identity of the victims be unknown so that no vivid and evocative description of the plight of those who might be helped is possible. I say none of this should matter. The point I insisted on still holds.

In many situations, it will turn out to be the case that when the most efficient scheme would be for many to offer aid, but some members of this optimal group of cooperators fail to do their part, the cost-of-giving/gain-from-receiving ratio that would characterize action by the remaining faithful cooperators that takes up this slack would significantly increase. But I submit that when this is so, it is this unfavorable shift in the relevant ratio of cost to benefit that renders it at least somewhat plausible to hold that the moral obligation to sacrifice for the benefit of others gives out. What does not seem even initially plausible is that the sheer fact of noncompliance with the optimal scheme by some (yielding increased costs to be incurred by the rest) lessens the level of sacrifice that morality demands of potential cooperators.

Notice that there can be cases in which some are not cooperating with the optimal scheme for provision of aid, so the Cooperation Principle triggers an absolute permission not to aid, but a subset of potential cooperators has made significant sacrifices that will bear no fruit unless the agent cooperates. In such cases, fairness to a subset of potential cooperators who are actually cooperating is ignored by the Cooperative Principle, which registers a fairness concern only among all potential cooperators. One might also object to the Cooperative Principle on the ground that in many cases considerations of fairness to the potential beneficiaries of our cooperative acts should outweigh any countervailing considerations of fairness among the set of potential cooperators. The Cooperative Principle takes account only of a limited aspect of the fairness considerations that bear on the determination of what I ought to do and ignores the rest.

The ratio of the cost-to-the-agent-from-giving to the-net-gains-that-the-beneficiaries-obtain is an important determinant of the moral obligation to give aid on

any plausible view that denies that agents are always obligated to do whatever will bring about the best outcome at least in situations in which deontological constraints are not in play. Of course, this cost to benefit ratio cannot be the entire story about permissible departures from consequentialist beneficence. After all, this cost to benefit ratio might be extremely favorable in a case in which only piddling amounts of good are at stake. Perhaps in some set of circumstances by raising my left eyebrow slightly at absolutely trivial cost to myself I could bring about an ever so slight feeling of comfort in someone who is staring at me. The relevant cost to benefit ratio is overwhelmingly favorable but the total gain if I undertake this act of beneficence is so slight that if I am ever morally at liberty to deviate in my conduct from optimal action, I must be morally at liberty so to deviate in this sort of case.

A similar point might hold when the cost to the agent of helping becomes extremely high. Imagine a science fiction scenario in which subjecting myself to a long life of utter hell or intense continuous suffering that is not offset by any compensating goods at all, a life that is far worse than instant death, would somehow bring it about that huge numbers of people threatened with hellish lives come to have excellent lives. One might take the line that no matter how favorable the cost to benefit ratio, if the gain to others consequent on giving is too small, the agent must be permitted not to give, and if the cost to the agent of giving is too large, the agent must also be permitted not to give.

Against Murphy, I have urged that the proper limits of the moral requirements of beneficence must depend on the ratio of the cost to the agent/gains to the potential beneficiaries. The Cooperative Principle entirely ignores this factor, so it is unacceptable. Later in this essay, discussing Samuel Scheffler's agent-centered

prerogative, I deny that this cost/benefit ratio does constrain the requirements of beneficence, so it might seem my argument undermines itself. Not so. Against Murphy I hold that IF there are moral limits on beneficence, they must be set or at least significantly partly determined by the moral cost/benefit ratio. Against Scheffler I deny the antecedent of this conditional.

SELF-INTEREST AS CONSTRAINT ON PERMISSIBLE SELF-SACRIFICE

Up to this point we have been reviewing a proposal to limit the extent of morally required optimal beneficence. Jean Hampton interestingly suggests a far more radical proposal that asserts stringent limits on the extent of morally permissible beneficence (whether optimal or not).⁷ Hampton's account turns the problem that the Singer Principle raises on its head: The needs of the self and of those loved by the self take strict priority over the needs of mere strangers, so the question, what stops the needs of strangers from dominating the appropriate choices of how to live, is just a nonstarter.

Hampton imagines an agent, Terry, who sacrifices her own interests for the sake of the interest of her family members. She does so in a lopsided fashion, preferring to sacrifice her own important interests even with respect to such basics as good health in order to achieve trivial advantages for her family members. Terry believes in a rigid division of labor between men and women such that childrearing and homemaking are the responsibilities of the wife and mother, and these responsibilities must be fulfilled even at great cost to oneself. Hampton rather plausibly suggests that as described, Terry's self-abnegating behavior seems morally wrong. The moral of the story drawn by Hampton is that many views of morality construe it as concerned with requiring individuals to cater to the well-being and freedom of other persons, but this undue

emphasis on what we owe to others tends to obscure the issue of what each of us owes, as a matter of morality, to herself.

Focusing on the self-regarding aspect of morality, Hampton finds its demands to be stringent, amounting almost to morally compulsory egoism. The story she tells is complex. Its upshot is that altruistic behavior is morally wrong when it conflicts with the requirements of self-respect, which include valuing oneself as a human being of equal worth and taking care that one flourishes as a human being and as a particular human being. The latter involves self-authorship, which one achieves by developing preferences and aims that are genuinely one's own, defined by oneself, and having content that does not "conflict with what is required to meet the person's objective needs as a human being."⁸ Choosing to help others at cost to oneself is only compatible with the requirements of self-respect and hence morally permissible if it is "authentic and done out of love."⁹ An authentic choice is autonomous, and proceeds from a self-authored preference. Choices are done out of love when the doers conceive themselves as "so unified with those whom their acts are attempting to benefit that what they regard as good for themselves is what will be good for those with whom they are united."¹⁰ It might seem then that true self-sacrifice is ruled out and egoism is morally obligatory, on Hampton's analysis. Altruism that involves "self-sacrifice" and is morally permissible gives to another whose interests are fused with one's own interests, so that what might seem self-sacrifice really redounds to one's own interest. But this is not quite correct. Hampton is prepared to acknowledge that there might be authentic altruistic choices done from love that benefit the helper to some extent, because to some extent a gain for the helped is a gain for the helper, but overall the helper predictably ends up worse off from

the altruistic act, because the gain from fused interests is overbalanced by loss to the giver. Hampton gives the example of a woman who must choose either to care for her ailing alcoholic father, whom she loves, or to pursue an independent life that develops her talents, when she cannot do both. In such a case, really sacrificing oneself for the sake of another might be morally permissible, Hampton allows.

The main argument that Hampton offers to motivate this elaborate structure of moral constraints on self-sacrificing beneficence is the case of Terry. For the moment let us accept, what seems at least somewhat plausible, that Hampton is right to judge that Terry acts wrongly in sacrificing herself. What might explain this judgment? Hampton leaps to a position that is very near to moral egoism to explain the wrongness of what Terry does. There are other possible explanations that she neglects to consider:

1. In the example Terry sacrifices a lot of her own well-being or utility in order to gain a very little well-being for her close relatives. A straight utilitarian morality condemns this self-abnegating altruism, which we will construe as sacrificing one's own interests to benefit others when the benefits are worth less than the costs to self as assessed by impartial moral standards.

2. Perhaps the most salient ground for thinking that Terry acts wrongly is that she lets down the side. By accepting an arrangement with her husband in which he gets the lion's share of benefits and the squirrel's share of burdens and she gets the squirrel's share of benefits and the lion's share of burdens, she contributes to the poor bargaining position of many other women struggling for fair terms in their relationships with male husbands and boy friends. Moreover, in acting out the terms of this grotesquely lopsided bargain in daily family life, she trains her children to accept inferior caste status for

women as morally appropriate. In these ways she behaves unfairly to women who continue to suffer from social norms, sustained by widespread beliefs and expectations, that put women at a disadvantage in dealings with men.

This line of thought is especially interesting but raises issues that are tangential to the concerns of this paper, so I shall ignore it.

3. Terry's behavior would also be condemned by a moral principle that was advanced by C. D. Broad under the label "self-regarding altruism."¹¹ According to this view, one ought to act so as to maximize a function of human good that gives special weight to advancing the good of those people to whom one is connected by ties of affection and involvement. That is to say, one should give special weight to advancing the good of those who are near and dear to us. The person to whom one is most near and dear is oneself, so each of us should put ourselves first. The good of mere strangers should count in these calculations, but at a discount compared to the good of those to whom we have close ties. Self-regarding altruism would approve Terry's giving priority to helping her own family members rather than to attending to the needs of strangers but would disapprove her putting the good of family members above her own good in the manner of self-abnegating altruism.

My first claim about Hampton's argument then is that each of these three explanations of why Terry's conduct as described is morally wrong is initially more plausible than Hampton's postulated near-egoism. They at least deserve a careful hearing. So even if Hampton had guessed right in her account of what makes Terry's conduct wrong, Hampton would not be entitled to claim her account is correct, in the absence of any careful assessment of the obvious alternative accounts.

But is Terry's conduct morally wrong? If it is not, Hampton is providing an explanation of a nonfact. A plausible line of thought leads to the conclusion that the answer is No, Terry's conduct is permissible. Before tracing this line of thought, I want to register some doubts about Hampton's account of the moral priority of helping oneself before helping others. (I should note that I ultimately will hold that Terry's conduct as described is morally wrong because it violates the Singer Principle and act-consequentialism.)

Some of Hampton's restrictions on beneficence seem too stringent. Consider the requirement that altruistic self-sacrifice, if it is to be morally permissible, must proceed from a love that connects the person aiding and the person aided so that, from the standpoint of the aiding person, their interests are unified. If it is the case that most humans are capable of summoning up such love only for their kin and friends and close associates, then for most of us Hampton's rules straightaway forbid any significant self-sacrifice for the sake of distant strangers. This conclusion looks to be implausible on its face, and since she does not provide any argument that might mitigate this initial impression, the impression stands.

The requirement that the choice to sacrifice oneself for the sake of others must be authentic, meaning in part that it must proceed from a self-authored preference, seems to rule out as impermissible all altruism chosen by those who through bad luck or their own fault have failed to develop self-authored altruism. When Sidney Carton, the wastrel character in the novel by Charles Dickens, says, "It is a far, far better thing that I do, than I have ever done," he may well be speaking truly, and his act of self-sacrifice might be motivated in part by inchoate guilt and feelings of low self-esteem.¹² Yet his act might

be reasonable and morally permissible even if he cannot summon up the will to do the act just because he perceives it to be right, and even if he is motivated in part by considerations that are nonauthentic. The act might still be the one he ought to do, and a fortiori morally permissible, even if his doing it lacks moral worth.

Hampton starts from the insistence that the agent must respect himself in the sense of regarding himself as a human being and so possessing the equal worth that all humans share. But in a world where natural and human evil bring it about that many of these equally morally worthy humans will not have the opportunity to satisfy their basic needs, why must the agent who recognizes he is equal in moral worth to any other person give priority to bringing about the satisfaction of his own basic needs when he could produce a better outcome by sacrificing his needs so that others can flourish? The transition in Hampton's argument from the assertion of Kantian moral equality of persons to the presumption in favor of oneself is obscure.

Let's return to the question: Is Terry's conduct morally wrong? Suppose one holds that favoring oneself over the interests of other persons in deciding what to do is not morally required, as Hampton urges, but is at least morally permissible. Terry then is allowed to deviate from doing what would bring about best consequences as impartially assessed by giving special weight to her own personal needs and interests. She is morally permitted to pursue her own lesser good instead of the greater good of others. But if she is permitted to favor herself over others, counting a gain in her own utility (avoidance of utility loss) as more valuable than a same-sized utility gain (or avoidance of loss) for others, why should she not equally be morally at liberty to give the same greater weight to the good of any other person rather than to her own good? This is exactly what she

does. If favoring self is morally permissible, then favoring anyone else other than the self should be equally morally permissible, and then the ground for condemning Terry's conduct slips away beneath our feet.

THE AGENT-CENTERED PREROGATIVE

The idea that the moral limit of required benevolence is set by the comparison between the cost of benevolence to the agent and the gain to others it would bring about is the basis of an appealing proposal developed by Samuel Scheffler. Scheffler suggests that one is morally permitted to deviate from acting as act-consequentialism would dictate according to what he calls an "agent-centered prerogative."¹³ This allows one to give greater weight to one's own personal concerns in deciding what to do than the weight that would be accorded those concerns in an agent-neutral consequentialist calculation. Each of us has a personal perspective and cares about things in a certain way from that perspective. What makes no difference at all from the impartial perspective makes all the difference from my individual partial perspective. It is entirely natural and understandable that each of us in deciding what to do gives weight to her own perspective in assessing the consequences of actions she might take.

Scheffler asserts that an agent is always permitted to do a less than optimal act just in case the loss to him if he were to do any act available to him whose outcome is superior, as assessed by impartial calculation, instead of this less than optimal act is less than M times greater than the loss to others as assessed by impartial calculation if he were to do the impartially superior act instead.¹⁴ The multiplier M registers the disproportionate weight the agent is allowed by the agent-centered prerogative to give to her own projects and plans in deciding what to do. The agent is always permitted to do

the optimal act if she chooses, but she is also free to choose any less than optimal act whose benefit to her (compared to what would ensue for her if she did any alternative impartially superior act) when multiplied by the factor M is greater than the loss to others consequent on her doing her preferred less than optimal act.

The agent-centered prerogative provides a partial but appealing resolution of the puzzle about beneficence introduced by Singer. When the agent faces the drowning child in the pond, the impartial value of avoiding the child's premature death outweighs the value of avoidance of minor inconvenience (mud-stained pants) even when the agent's presumed personal concern for his personal convenience is given extra weight according to the prerogative. But when we imagine the agent being asked to give more and more of his monthly income to save the lives of needy persons around the globe, at some point the disvalue to the agent of his lost income outweighs the great impartial good of saving more lives in a calculation in which the agent's personal concerns are given extra weight by the prerogative. What is appealing about this sketch of a resolution is that we preserve the intuitively attractive judgments that the agent must help the child and need not help the distant strangers. (If the agent happens to have a special personal concern to help those needy strangers near at hand and vividly present to her mind, that personal concern gives extra weight to choosing to help the drowning child, so it might well be permissible to save the drowning child even when the impartially best outcome the agent could produce would be to abandon the child and use the extra resources saved to help a greater number of needy distant strangers.)

The match of common-sense moral judgment and the recommendations dictated by the agent-centered prerogative is not complete. If the impartial good the agent could

gain by ever greater contributions to the project of helping distant needy strangers is sufficiently large, then a prerogative with a reasonable multiplier would require the agent to sacrifice her personal concerns to the impartial good to an enormous extent. An alternative related possibility is that if the prerogative is set with a multiplier large enough to allow the agent to pursue her own personal concerns in the face of enormous impartial good to distant strangers that tending her own private garden forgoes, this prerogative will give so much weight to the agent's own personal concerns that it allows her to let the drowning child sink to the bottom of the pond with the result that her concern to avoid muddying her pants is satisfied. But the agent-centered prerogative makes it possible to develop moral judgments about beneficence far more in line with ordinary moral opinion than the Singer Principle and act-consequentialist moralities.

Most of the critical response to Scheffler's proposal centers on the fact that it draws no distinction between options and constraints.¹⁵ If it would be acceptable for me to act to get myself a benefit of a certain size, which imposes costs totalling C on other agents by way of neglecting to act to help them, it would be equally acceptable for me to get a benefit for myself of that same size by doing an act that violates a moral constraint against harming others and imposes in this way costs on others totalling C . The critics claim that any multiplier stipulated by the agent-centered prerogative that is large enough to give an agent an intuitively satisfying range of options would yield counterintuitive recommendations when the agent-centered prerogative is applied to choices in which an agent would suffer loss to self consequent on obeying some significant moral constraint against harming other people in some specified way. This issue is beyond the scope of this essay, and so not our concern here.

My worry is different. The Scheffler proposal yields implausible recommendations that are flawed in essentially the same way that the Hampton proposal was found to be. Scheffler's agent-centered prerogative makes the permissibility of acting to promote a concern depend on whether or not it is mine. For example, suppose the prerogative would have it that I am permitted to give medicine to my sick father rather than do any of the alternative acts available to me that would bring about a better outcome as impartially assessed. In the same circumstances the prerogative would not permit me to give the same medicine to someone else's equally sick father even though the gain to him would be identical to the benefit my father would gain and the act of helping someone else's father would produce an outcome that is just as good as assessed from an impartial perspective as the outcome that would result from the act of helping my own father. In other words, the Scheffler prerogative supposes that there is moral magic in the pronoun "my." My concerns, my projects and plans are given special extra weight in moral calculation that determines what it is permissible for me to do, whereas exactly similar concerns, plans, and projects that are other people's rather than my own do not generate special moral license to act to fulfill those equally morally worthy matters.

For another example, consider the act of buying for my own consumption a giant hot fudge sundae instead of performing some act with an outcome that is superior according to impartial assessment. The costs and benefits to self and others are such that the prerogative gives me permission to do the self-favoring act. Now suppose that on a whim I decide instead to buy for someone else's consumption the same giant hot fudge sundae instead of performing some act with an outcome that is superior according to impartial assessment. We suppose the beneficiary of my whimsical largesse would

derive just as much pleasure and nutritional gain and other benefits from eating the sundae as I would if I ate it myself. The agent-centered prerogative gives no extra weight to gains to this beneficiary, so it could well happen that doing the nonoptimal act that benefits another person but not me is forbidden by the agent-centered prerogative. This result strikes me as arbitrary, just as arbitrary as the judgment by Hampton that whereas it would be perfectly morally acceptable for Terry to act so as to favor herself over others, it would be morally unacceptable for Terry to favor others over herself.

It might be supposed that I am misinterpreting Scheffler in an obvious way. If I choose to benefit another person rather than myself, then it follows that I have some personal concern, plan or project to benefit that other person. So if the prerogative gives enough weight to my personal concerns to render it permissible for me to donate resources to my sick father rather than use them to secure some impartially better outcome, then if I choose to aid someone else's similarly situated father instead, then doing that must be my personal concern, so the prerogative kicks in, and renders the decision to favor others just as permissible as the decision to favor myself. And the same goes for the permissibility of the choices canvassed in the giant hot fudge sundae example.

This suggestion is incorrect. It is true that the personal concerns that are given extra weight in calculation of morally permissible action according to Scheffler's prerogative need not be personal concerns to benefit myself. I might have a personal concern to pursue some cause, or be nice to animals, or follow any of myriad interests. But the Scheffler construction supposes one can distinguish matters that are for an agent matters of her personal concern from matters that are for that agent not matters of

personal concern. The prerogative gives permissibility-enhancing weight only to the agent's personal concerns, the things she cares about, whatever they might be. But it is not taken to be necessarily or even just contingently true that each agent always acts only to further her own personal concerns. Scheffler's construction supposes that an agent in given circumstances might instead of acting to promote her personal concerns choose instead to do what is best as assessed from an impartial perspective. The agent can choose to act to promote her personal concerns or decline to do so. And just as the agent can choose to promote the impartial good instead of her personal concerns, the agent equally might choose to promote concerns of some other people rather than either her personal concerns or the dictates of impartial good. It is possible to do this, but Scheffler's agent-centered prerogative says it is wrong to do so.

My objection could be met by amending Scheffler's prerogative so that it would allow an agent to give disproportionate weight to the satisfaction of anybody's personal interests of her choosing beyond the value that would be accorded these interests by impartially neutral calculation. The prerogative allows the individual to count the satisfaction of these selected interests as M times more valuable than the satisfaction of other interests in deciding what to do. Severed from the idea that people should be free to pursue their own projects and plans to a greater extent than act-consequentialist calculation or the Singer Principle would allow them to do, the prerogative may now look just bizarre. But there are reasons for devising a prerogative that does not rely on the claimed moral reasonableness of deference to personal concern. Moreover, what we might call the *bare prerogative* has some appeal.

The Scheffler agent-centered prerogative attempts to blend together two different lines of objection to act-consequentialism and provide a single solution that is responsive to both objections. On the one hand, act-consequentialism is too demanding in that in many circumstances it would force an agent to sacrifice the pursuit of her own personal concerns to an excessive extent. On the other hand, act-consequentialism is too demanding in that it is too restrictive of the agent's moral freedom of choice. That these are quite different objections pulling in different directions is immediately obvious. Acceptance of act-egoism would fully satisfy the first objection but do nothing to ameliorate the second difficulty. The second difficulty could be satisfied by a satisficing consequentialism which selects among the options of acts available to an agent those that are "good enough" in their consequences as ranked by an impartial consequentialist standard. The agent is then free to do any of the acts selected as "good enough" in their consequences. Provided the bar of "good enough" consequences is set sufficiently low, satisficing consequentialism in many circumstances would not select one act as morally required to do but would rather give the agent a wide range of morally permissible choices. But all of these choices might still be disastrous for the agent's own personal concerns. Hence satisficing consequentialism might be subject to the first objection but not the second.

Once these two strains of thought are separated, each taken separately looks problematic. Take the imperative of restricting the demands of morality to allow the individual to pursue her own personal concerns even if these are not weighty when viewed from the impartial perspective. Hampton's Terry case shows that it is not plausible to hold that morality requires one to give priority to one's own personal

concerns in deciding what to do. If it would be acceptable for Terry to give special weight beyond what impartial rankings of consequences would allow to her own rational self-interest, then it should be acceptable for her to give comparable extra weight to her non-self-interested personal concerns. Morality does not command that one act egoistically. Nor does morality command that one act in a way that gives special priority to the things one happens to care about. But then if it is acceptable for Terry to give special weight to the interests of Fred, for whom she has personal concern, why would it not be equally acceptable for her to give similar priority to the interests of Samantha, a stranger for whom she has no such personal concern. After all, what is morally special about the things I happen to care about?

Of course there may be answers to this sort of question from the standpoint of impartial rankings of consequences. But these answers are irrelevant in this context. For example, it may be impartially desirable that friends be true to friends, so the world in which one friend is nice to another is a better world than another otherwise identical except that the friend is nice instead to a random stranger. Relationships of parent to child might beget similar special values that are realized when parents help their own children (and that are not realized when parents give similar help to other people's children). But these considerations so regarded will affect the impartial rankings of consequences and hence the determination of what it is best from the impartial consequentialist standpoint for an agent to do in given circumstances. Our concern is different. We are concerned with cases in which the consequentialist rankings of states and affairs and acts that might produce them brings it about that it would be optimistic for the agent to act in some way that is too demanding.

Consider this formulation of Scheffler's objection against act-consequentialism: By requiring us always to do what brings about what is impersonally best, consequentialism alienates us from our personal concerns. In this sense it threatens integrity. Is acting with integrity then what we morally may do or morally must do? Scheffler's position represents a compromise. According to his agent-centered prerogative, an agent is always morally at liberty to do either (1) the impersonally best act or (2) any act that brings about an outcome that advances one's own personal concerns (multiplied by the prerogative) by comparison with the benchmark level of advance of one's personal concerns that the impersonally best outcome would have achieved by an amount no less than the amount of loss of impersonal good that would be occasioned by doing this act rather than what is impersonally best. This means one is morally at liberty according to the prerogative to do an act that advances one's personal concerns less than one could permissibly attain according to the prerogative so long as the constraint just described is still satisfied. The agent-centered prerogative then gives one permission to stray from what is personally best only in the direction of what is better from the standpoint of one's personal concerns, but this act need not be the one that does the best one permissibly can do from the standpoint of one's personal concerns.

A limit case of acting from one's personal concerns would be acting to satisfy a momentary desire or whim. Imagine that the Schefflerian agent-centered prerogative allows me to satisfy my whim to bequeath my wealth to a random person whose physical attractiveness I fancy now in Grand Central Station rather than to bequeath to my expectable heirs, which would best satisfy my long-run overall, personal concerns, or to do what is impartially best. But since I happen not to have the momentary desire of

bequeathing my wealth to a random person whose physical attractiveness I do not fancy now in Grand Central Station, that choice would be morally impermissible. Too much here seems to hang on too little. If maximal fulfillment of one's overall personal concerns is of great moral significance, and trumps any requirement to do what is impartially best, why is acting from integrity in this sense not morally required if one is going to do what is not impartially best? On the other hand, if having wide moral liberty to do as one chooses is of great moral significance, why is this moral liberty limited at all by the shape of one's actual personal concerns?

If one responds to this last question by opting for wide moral liberty over personal concerns, one then ends up with the revision of Scheffler's proposal that I have called the bare prerogative. The agent-centered prerogative then becomes a blanket permission to give special weight to any interests, above the weight these interests would get in an impartial consequentialist ranking. These interests need not be special concerns of the agent. When we press on the idea that morality allows each of us to give special weight to his own personal concerns in deciding what to do, we end up leaving the idea of the agent's own personal concerns behind. What we end up with is essentially the claim that agents should be granted wide moral freedom, should be morally permitted to act in a wide variety of ways on most occasions of choice. The question then arises why this ideal of moral freedom should be thought to be compelling. Why is this the key determinant of the limits of what we owe to distant strangers? Moral freedom in and of itself seems a bloodless and abstract matter. Imagine cases in which the agent herself has no special concern for moral freedom and takes no interest in it. The position we are

considering says that even here preserving the agent's moral freedom is of great moral importance. Maybe so; but there is room for doubt.

Another source of doubt insinuates itself into the story when we ask how much weight the prerogative should give to the interests that are to receive extra weight above what impartial ranking would give them? What is supposed to determine a correct answer to this question? I find I cannot get a grip on these questions.

The proposed elaborations of the thought that act-consequentialism is too demanding (the Scheffler prerogative and the bare prerogative) are neither of them as plausible as the initial claim.

WHAT IS MORALLY RIGHT IS NOT MORALLY OBLIGATORY

I propose a way of dealing with the puzzles of benevolence discussed by Singer, Murphy, Hampton, and Scheffler that holds onto the act-consequentialist claim that one morally ought always to do the act of the available alternatives that would produce the best consequences as impartially assessed.¹⁶

In response to the claim that affluent people should entirely devote their lives to the relief of the suffering and crippling poverty of the worst-off persons on the globe, one's initial response is not so much that this claim is unjustified as that affluent people cannot be blamed for failing to live up to this austere ideal of self-sacrifice. Given human nature, which strongly inclines us to put ourselves and those near and dear to us first in our priorities when we decide how to act, it would be at the least extremely difficult for people to adhere to the Singer Principle, and given the great difficulty of complying with this code, it is no great sin that we do not, and priggish to act as though it were. There is the further consideration that people are not trained to accept anything

close to the demands of the Singer Principle, so their intuitive responses instilled by processes of socialization pull against acceptance of it and compliance with it. So we should sharply distinguish what it is right and wrong to do and what we should be praised and blamed for doing and not doing. It is morally right to do the optimal act and morally wrong to do anything else, but in many circumstances including the context in which issues of famine relief and relief of global poverty arise, one is not (or hardly at all) blameworthy for failing to do what is morally right.

One might resist this easy division of responses to the global poverty puzzle by asserting that doing right and wrong and being subject to praise and blame are conceptually connected, so the proposal to separate them is incoherent. John Stuart Mill clearly expressed the claim that there is a conceptual connection between an act's being morally wrong and being fit for punishment. He writes, "We do not call an act wrong unless we mean to imply that a person ought to be punished in some way or other for doing it—if not by law, by the opinion of his fellow creatures; if not by opinion, by the reproaches of his own conscience."¹⁷

I disagree. What is morally right and wrong to do is established by the decisive balance of reasons for and against various courses of action, on the assumption that moral reasons are trumping considerations or perhaps all things considered considerations, which incorporate and assign the correct weight to all nonmoral reasons that bear on the choice of action. What is morally right is what there is most reason to do, and the other act alternatives one might do instead are morally wrong. On the act-consequentialist analysis, the act that there is most reason to do is identified with the optimal act, that act the doing of which would produce best consequences as impartially assessed. On each

occasion of choice, the alternative acts one might do that are not optimal are morally wrong.¹⁸ Nothing in the story so far says or implies anything about blame or punishment.

Shouldn't we feel guilty if we do what is morally wrong? Not necessarily. From a consequentialist standpoint, whether someone should be punished for doing or refraining from doing some act depends on what the consequences would be of performing that act of punishment. Even the punishment of guilt feelings that spontaneously well up in the consciousness of the wrongdoer can be regarded as to some extent the outcome of a mechanism that is set in place by prior human acts—acts of socialization by others and acts of self-culture or character building by the agent herself. The consequentialist evaluation of this psychic mechanism that induces guilt in a person can be traced to the assessment of the acts that combine with circumstances and the person's innate dispositions to produce the mechanism.

Let us say that an act is *obligatory* just in case one should be punished (at the minimum, by guilt or self-blame) for failure to do it. If it were the case that a person should always be punished on each occasion that she fails to do whatever is optimific, then the category of *action it would be morally wrong not to do* and the category of *obligatory action* would always coincide. But it is a common observation in discussions of consequentialist morality that failure to do what is optimific often does not render one fit for punishment in the sense that punishing one would have consequences better than any alternative. Failure to distinguish what is right and wrong from what is obligatory and forbidden renders act-consequentialism formulated in terms of the former contrast a less plausible-sounding doctrine than it should seem.

The same analysis can apply to decision problems in which one can accomplish significant good for the wretched of the earth by significant self-sacrifice. In the cases of global famine and poverty relief described by Singer, it is stipulated that doing the benevolent and self-sacrificing act produces best consequences and there is nothing morally to be said against this benevolent choice other than that it involves real self-sacrifice. Is doing the benevolent act in these circumstances the morally right act, and failure to do it morally wrong? The act-consequentialist replies resoundingly in the affirmative. Is the failure to be benevolent at significant cost to oneself a violation of moral obligation, for which one should be punished? From the act-consequentialist standpoint, the answer to this question may be very uncertain, because it is unclear, human nature being what it is, whether punishing such humdrum selfish acts would have good or bad consequences. To this extent the act-consequentialist position, in this context identical to the Singer Principle, does a good job both in matching our considered moral responses and explaining them. I am envisaging that the morally perceptive affluent person, munching his croissants and fingering the keys to his fancy car, does not hold that it is morally justified that he enjoys the good life at the moral cost that others live avoidably wretched lives. His position is, "Can you blame me?"

Raising the level of moral obligation will generally tend to produce two effects that pull in opposite directions. Setting the level higher will tend to improve the outcomes of the actions of those people who conscientiously strive to comply with the higher norm. But setting the level higher, above some point,¹⁹ will also tend to reduce the numbers of people who conscientiously strive to comply with moral requirements. As the level of requirements becomes ever more demanding, more people will tend to

become alienated from the enterprise of morality and to become less disposed to carry out even its minimal requirements. At some level of obligation these two effects balance, so that any raising or lowering of the level of obligation would produce worse consequences overall; at this point the level of obligation is optimal. But there is no particular reason to think the optimal level of obligation will hold that people are always morally obligated to do what according to act-consequentialism is morally right.

Distinguishing what is morally right and wrong from what is morally obligatory and forbidden helps make sense of intuitive puzzles about the morality of beneficence and the difficulty of coming to a stable view about the moral limits on the extent to which beneficence, when optimal, is required. If the agent is not a saint, she will be strongly motivated to give special priority to promoting her personal projects and concerns by her actions, and choosing acts that cater to these concerns that are subjectively important to her rather than choosing the available act whose consequences are ranked best from the impartial standpoint. If the agent chooses acts that cater to her personal concerns, there will virtually always be a strong excuse for her choosing to do that rather than what is impartially best. Hence blame for such cases would often be misplaced. This across the board presumption that the agent should be excused for failing to do what is optimal because she is subject to strong motives to cater to her personal concerns does not come into play if the agent acts in a way that is nonoptimal and that also does not further her own personal concerns but rather any concerns selected randomly or capriciously at the moment of choice.

I have urged that there is no more justification for acting to further one's own personal concerns than acting to further anyone else's concerns instead of doing what

would produce best consequences impartially assessed. So if acting nonoptimally to further personal concerns is justified, acting nonoptimally to advance any comparable concerns not one's own is also justified. But it is odd that personal concerns and what matters from the personal point of view drop out of consideration in this way. But if we switch from thinking in terms of justification to thinking in terms of excuses and the conditions under which blame is inapplicable, the banished but seemingly sensible idea of deferring to the agent's personal concern returns and falls into place naturally. In the types of case in which we are intuitively inclined to say that act-consequentialism along with its cousin the Singer principle are too demanding in the requirements of beneficence they impose, we should instead say that these are cases where, although one should always do what would produce the best consequences, when doing so pulls sharply against one's personal concerns, the agent is not to be blamed for acting nonoptimally.

One difficulty with my proposal is that it might seem to leave unanswered the concern that act-consequentialist morality is too demanding in the sense that it dictates leave the individual choosing what to do too little moral freedom to choose among a wide range of alternative courses of action and ways of life.

My response is that once the idea of what is morally right and wrong is separated from the idea of what is morally obligatory and forbidden, with phenomena of guilt, blame, and punishment linked to violation of what is morally obligatory, the moral freedom concern becomes less urgent. The mere fact that I am not morally free to do anything except the available act on each occasion of choice that would produce consequences no worse than the consequences of any alternative I might choose is simply not in and of itself a severe restriction on my effective freedom to live as I choose. My

effective freedom in a situation of choice is constituted by the options such that if I were to choose them, I would get them. Lack of opportunities and resources and lack of abilities restrict my effective freedom. If I am not morally free to choose X, that only means that I morally ought not to choose it. This is fully consistent with its being the case that I am effectively free to X--if I choose it, my choice will be effective.

More should be said to clarify the idea of an act's being fit for punishment that determines the idea of an obligation in my proposed usage. We can distinguish being obligated and feeling oneself to be obligated. Whether or not one is obligated depends on whether or not one ought to be punished for omitting to do the obligatory thing. Whether or not one feels obligated or experiences a sense of obligation depends on how one has been socialized, what norms are current in one's society, one's innate dispositions, and the current state of one's moral thinking.

We can also distinguish being ideally and pragmatically fit for punishment. The distinction turns on the degree to which the context in which the action occurs is taken for granted. In the ideal sense, an act is fit for punishment just in case a society like the actual society in which the act occurs except that it is ideally regulated by act-consequentialist principle would establish practices of socialization, social norms of conduct, and institutions of punishment that would bring it about that acts of this type would incur punishment. In the pragmatic sense, an act is fit for punishment just in case it would be optimal (produce best consequences) in the circumstances in which the act occurs that it incur punishment. Corresponding to these two senses of being fit for punishment we can identify two different notions of being morally obligatory.

The ideal sense of obligation gives us a way of characterizing what people morally ought to do and are obligated to do under conditions that are as favorable as we could hope for the production of impartially good consequences. In the context of considering moral duties and obligations to contribute to the alleviation of avoidable global poverty and misery, the relevant society to consider would be a single world society ideally regulated by consequentialist principle. This regulatory regime will compromise with human nature rather than wage war against it. Suppose that you and I are impoverished, disease-plagued, jobless and homeless members of this society, and that many of its members are affluent to the degree that they could contribute significantly to the improvement of the quality of life of people such as you and I at a cost of self-sacrifice that is such that act-consequentialism (the Singer Principle) would endorse these possible acts of beneficence by the affluent. But in fact the affluent do not make these contributions and the social norms, socialization practices, and punishment practices of the society do not sanction their failure to help. What gives? Should not society be reformed so the affluent well off people are made to help us? If the society is ideally regulated by consequentialist principle, the following answer can be made: It is true that the affluent act morally wrongly in failing to give to help alleviate your plight. They do what is wrong. But they do not violate any moral obligation in the ideal sense. Since the order of society as it is is optimific, any attempt to sanction the affluent or alter institutions or practices in any way in order to bring it about that your and my lives are improved would bring it about that some people with (in the aggregate) greater moral claims to improved quality of life than you and I now have would be made worse off.

For all that has been said, it might well be that the widespread failure of well-off persons in affluent societies to contribute enormously more than they do at present to the relief of global poverty and misery and premature death is not only morally wrong by act-consequentialist standards but also a violation of obligation in the pragmatic sense as well as in the ideal sense. In this case, even if we do not feel guilty and blame ourselves for failure to contribute at a decent level, we should. My task in this essay is not to insulate the comfortable from blame but to indicate that an act-consequentialist morality can generate a nuanced analysis of this situation that is not obviously at variance with our reflective judgments after critical scrutiny. Moreover, achieving a reflective equilibrium among our moral judgments in this domain is not easy. If act-consequentialism and the Singer principle do poorly on this test, Murphy's Cooperative Principle, Hampton's espousal of strict requirements for permissible altruism, and Scheffler's agent-centered prerogative all score worse.

¹ . Peter Singer, "Famine, Affluence, and Morality," *Philosophy and Public Affairs*, vol. 1 (Spring, 1972), pp. 229-243. For criticism from a deontological perspective, see Frances M. Kamm, "Famine Ethics: The Problem of Distance in Morality and Singer's Ethical Theory," in Dale Jamieson, ed., *Singer and His Critics* Oxford: Blackwell, 1999), pp. 162-208. For elaboration of a morality of beneficence (inter alia) from a virtue ethics perspective, see Michael Slote, "The Justice of Caring," *Social Philosophy and Policy*, vol. 15, no. 1 (Winter, 1998), pp.

171-195; also Slote, *From Morality to Virtue* (Oxford: Oxford University Press, 1995).

² . Liam Murphy, "The Demands of Beneficence," *Philosophy and Public Affairs*, vol.22, no. 4 (Fall, 1993): 267-292; see p. 278. For further discussion, see Murphy, *Moral Demands in Nonideal Theory* (Oxford: Oxford University Press, 2000).

³ . Murphy, p. 280.

⁴ . Here and subsequently in this essay I use "optimal" as Murphy does: An optimal act is one that would bring about the best consequences as impartially assessed. This usage should not be confused with the entirely different idea that a situation is Pareto optimal if any alteration from it to make someone better off would make someone worse off.

⁵ . One may wonder about how to interpret the Cooperative Principle. Are its requirements timeless, so that one is required by morality to do what would produce optimal results if everyone at all times had acted optimally? Or do we interpret the Principle as limiting the requirements of beneficence only by the level of activity that would produce optimal results if everyone now fully complied with the principle? If we accepted the Cooperative Principle, we would have to solve these problems about its interpretation, but I shall argue the Principle is unacceptable for reasons that are independent of the answers to the questions just raised. For further discussion of the Cooperative Principle see Tim Mulgan, "Two Conceptions of Benevolence," *Philosophy and Public Affairs*, vol. 26, no.1 (Winter, 1997) , pp. 62-79; also Liam B. Murphy, "A Relatively Plausible Principle

of Beneficence: Reply to Mulgan," *Philosophy and Public Affairs*, vol. 26, no. 1 (Winter, 1997), pp. 80-86.

⁶ . See Murphy, *Moral Demands in Nonideal theory*, pp. 127-133; also Peter Unger, *Living High and Letting Die* Oxford and New York, Oxford University Press, 1996).

⁷ . Jean Hampton, "Selflessness and the Loss of Self," *Social Philosophy and Policy* (December, 1993): pp. 135-165.

⁸ . Hampton, "Selflessness and the Loss of Self," p. 152.

⁹ . Hampton, "Selflessness and the Loss of Self," p. 157.

¹⁰ . Hampton, "Selflessness and the Loss of Self," p. 158.

¹¹ . C. D. Broad, "Self and Others," in *Broad's Critical Essays in Moral Philosophy*, ed. D. Cheney (London: George Allen and Unwin, 1971). I owe this reference to David Brink.

¹² . Charles Dickens, *A Tale of Two Cities*, in *The Oxford Illustrated Dickens* (Oxford: Oxford University Press, 1949), p. 358. The novel was first published in 1859. Actually the quoted words are Carton's unspoken thoughts just before he is guillotined. I have a memory of Carton uttering the words, but this must be derived from seeing a movie version of the book or reading the *Classics Comics* summary. My suggestion in the text probably twists the meaning of the novel. Dickens, not averse to bathos, paints Carton's act of self-sacrifice as self-constituting, self-redemptive, done from love, and thoroughly authentic.

¹³ . Samuel Scheffler, *The Rejection of Consequentialism*, revised edition (Oxford: Clarendon Press, 1994), p. viii. This book was originally published in 1982.

¹⁴ . Here I am interpreting Scheffler, p. 20. Shelly Kagan finds the quoted lines from Scheffler hard to construe, and I agree. See Kagan, "Does Consequentialism Demand Too Much? Recent Work on the Limits of Obligation," *Philosophy and Public Affairs*, vol. 13, no. 3 (Summer, 1984), pp. 239-254. I am not entirely sure that the interpretation I give in the text, following Kagan, fully captures what Scheffler has in mind.

¹⁵ . See for example, Shelly Kagan, "Does Consequentialism Demand Too Much? Recent Work on the Limits of Obligation"; also Frances Kamm, "Supererogation and Obligation," *Journal of Philosophy*, vol. 82, no. (1985), pp. 118-138. Scheffler replies in "Prerogatives without Restrictions, reprinted in the revised edition of *The Rejection of Consequentialism*, pp. 167-192.

¹⁶ . The two-level approach that I sketch here owes a lot to the two-level theory of morality proposed by R. M. Hare in *Moral Thinking: Its Method, Levels, and Point* (Oxford: Oxford University Press, 1980).

¹⁷ . John Stuart Mill, *Utilitarianism*, ed. by George Sher (Indianapolis: Hackett Publishing, 1979), chap. 5, p. 47. (*Utilitarianism* was originally published in 1861.)

¹⁸ . One might, following Mill's suggestion (*Utilitarianism*, chap. 2, p. 7), soften the rigor of this stark division of acts by holding that the more an act

diverges in its consequences from what is optimal, the more wrong it is. All the nonoptimal acts are wrong, but some are “wronger” than others.

¹⁹ . One might speculate that below some point, lowering the standard of obligation further would tend to increase alienation from morality among people who become disgusted at what morality countenances. The broad-stroke empirical speculation in the text just indicates the character of the cost-benefit calculation a social planner would need to undertake in order to determine how altering the current standard of obligation might promote or dampen the production of best outcomes.