

SHELLY KAGAN

Does Consequentialism
Demand Too Much?
Recent Work on the
Limits of Obligation¹

Consequentialism claims that an act is morally permissible if and only if it has better consequences than those of any available alternative act. This means that agents are morally required to make their largest possible contribution to the overall good—no matter what the sacrifice to themselves might involve (remembering only that their own well-being counts too). There is no limit to the sacrifices that morality can require; and agents are never permitted to favor their own interests at the expense of the greater good.

Our ordinary moral intuitions rebel at this picture. We want to claim that there is a limit to what morality can require of us. Some sacrifices for the sake of others are meritorious, but not required; they are supererogatory. Common morality grants the agent some room to pursue his own projects, even though other actions might have better consequences: we are permitted to promote the good, but we are not required to do so.

The objection that consequentialism demands too much is accepted uncritically by almost all of us; most moral philosophers introduce permission to perform nonoptimal acts without even a word in its defense. But the mere fact that our intuitions support some moral feature hardly constitutes in itself adequate philosophical justification. If we are to go beyond mere intuition mongering, we must search for deeper foundations. We must display the *reasons* for limiting the requirement to pursue the good.

1. The following works are reviewed in this essay: David Heyd, *Supererogation: Its Status in Ethical Theory* (Cambridge: Cambridge University Press, 1982); Thomas Nagel, "The Limits of Objectivity," in *The Tanner Lectures on Human Values*, vol. I, ed. Sterling McMurrin (Cambridge: Cambridge University Press, 1980), pp. 77–139; and Samuel Scheffler, *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions* (Oxford: Clarendon Press, 1982).

Nor is it sufficient to consider this issue in isolation: the arguments offered must cohere with the rest of what we want to defend about morality. The challenge is to provide a plausible defense that preserves the features which make the ordinary moral view a moderate one. A defense of the common view, obviously enough, must explain why it is sometimes permissible to refuse to perform some optimal act. But all those unwilling to embrace egoism must at the same time avoid arguments that rule out the possibility of there being any moral requirements at all. Thus the explanation must also account for the fact that sometimes a given optimal act *is* required by morality. Similarly, the account must capture the view that even when a given sacrifice is not required, it is nonetheless permissible—indeed, meritorious. Most arguments offered seem to be incompatible with these or other features of the common moral view, and therefore need to be rejected by those who wish to defend anything like that view.

Furthermore, discussions of the claim that consequentialism demands too much are often undermined by failure to distinguish this claim from the widely discussed objection that consequentialism permits too much—improperly permitting sacrifices to be *imposed* on some for the sake of others. Some theories include deontological restrictions, forbidding certain kinds of acts even when the consequences would be good. I will not consider here the merits of such restrictions. It is important to note, however, that even a theory which included such restrictions might still lack more general permission to act nonoptimally—requiring agents to promote the good within the permissible means. It is only the grounds for rejecting such a general requirement to promote the overall good that we will examine here.

I

David Heyd's *Supererogation* provides an instructive example of some of the problems I have just described. After presenting a historical survey of the place of supererogation in some major ethical theories, Heyd offers his own favored analysis of the concept. The details of Heyd's discussion, however, need not concern us, except for this: as the first systematic study of the subject of supererogation, one would expect Heyd to give particular care to justifying his view that some acts *are* supererogatory—optimal but not required. But *Supererogation* is a disappointing book.

When it finally addresses this task in its concluding chapter, instead of sustained argument we find a bewildering grab bag of incomplete thoughts. Like most who defend the existence of supererogatory acts, the intuitive appeal of this position seems to blind Heyd to the obvious inadequacies of his own discussion.

According to Heyd, the justification of supererogation has two “aspects”: the negative aspect shows the justification for limiting what is morally required; the positive aspect shows “the value of non-obligatory well-doing as such” (p. 166). Heyd’s summary of his argument suggests that each aspect is in turn supported by several considerations. Unfortunately, the actual presentation of the justifications makes no attempt to demarcate the boundaries of individual arguments, making the whole discussion difficult to keep straight. By my rough count, there are eleven arguments, offered in no particular order.² Several are actually no more than hints of arguments, and some may not even be meant as arguments at all. Furthermore, even the best of them have a half-baked quality about them, as though Heyd had desperately grabbed at anything at all that might support his case. Although there is no room here to review all of Heyd’s arguments, I have reconstructed the more interesting ones, and given them names.

The Incommensurability of Reasons (pp. 170–71). Heyd claims that there are two types of reasons: reasons to promote overall human welfare, and reasons of “autonomy,” which support the pursuit of one’s own ends. Although the first sort of reason is “*morally* superior” to the second kind, the two types of reasons “cannot be compared in terms of strength.” Thus it simply isn’t generally true that promoting the overall good is backed by stronger or conclusive reasons—that is, that an agent “ought” to promote the good. If the agent chooses instead to act on the reasons that support promoting his own ends, this cannot be faulted from the standpoint of practical reasons; he hasn’t failed to do what he *ought* to do.

By viewing the two types of reasons as incommensurable, Heyd would be able to explain why an agent might be free to perform some optimal act at great cost to himself, or to refrain, as he chooses. But such an

2. A guide to the perplexed (using my own labels): *The Incommensurability of Reasons* (pp. 170–71); *The Good-Ought Gap* (pp. 171–72); *The Basicness of Rights* (pp. 172–73); *Coherence* (pp. 173–74); *Integrity* (p. 174); *The Minimalist Model of Morality* (p. 174); *Justice* (pp. 174–75); *The Intrinsic Value of Supererogatory Acts* (pp. 175–76); *Mill* (p. 176); *Good Samaritan Legislation* (pp. 176–78); and *Evidence of Concern* (p. 179). Perhaps others would carve up the discussion differently.

account would make it impossible to explain how an agent could ever be *required* to promote the good—even at negligible cost to himself. Presumably in cases of this kind we want to say that the two sorts of reasons *can* be compared in terms of strength, and that reasons of autonomy have become so weak that they are conclusively outweighed. Yet how can supposedly incommensurable reasons become commensurable simply through variations in their magnitude? It just won't do for Heyd to support one part of his moral theory (supererogation) with a thesis that rules out another part of his moral theory (the existence of some requirements to aid; see, e.g., p. 90).

The Minimalist Model of Morality (p. 174). Morality, says Heyd, is not aimed at “the maximization of general good or happiness” but is rather “a means of securing some minimal conditions of cooperation and justice.” Presumably only these minimal constraints on the autonomy of the individual are justified; beyond that, individuals have “the right to pursue their own ends” (p. 172). Heyd does not explicitly state how the minimal level is to be set, but he later seems to endorse the view that moral requirements “constitute the minimum required for the preservation of society” (p. 181; cf. p. 177).

Such a standard, however, would set morality unacceptably low: if I can save someone's life at no cost to myself, surely I am required to do so; but the preservation of society would not be threatened if no one ever fulfilled such obligations. Thus all but an egoist will recognize requirements to aid that go beyond those grounded by Heyd's minimalist model. Heyd may very well want to include a modest principle of aid as part of his model—but to do so is *ad hoc*. Once this is seen, and one realizes just how minimal a genuinely minimalist model of morality would be, the view should lose its appeal altogether. But at any rate, the minimalist conception cannot be used by one (such as Heyd) who wants to claim that the requirement to aid should be limited—but not nonexistent.

These first two arguments bring out nicely the importance of bearing in mind that ordinary morality lies precariously perched between *two* extremes. Many attacks on consequentialism can only be achieved by opening the door to egoism (or other overly minimal views); most of us, therefore, cannot consistently accept such arguments.

Justice (pp. 174–75). Heyd claims that “the non-utilitarian concept of justice serves both to counter utilitarian arguments for the punishment

of an innocent individual and to support the distinction between duty and supererogation.” Just as “individual persons should not be sacrificed for the promotion of overall happiness (as in the case of punishing an innocent man in order to save the lives of many others),” similarly “considerations of justice make it unacceptable to require any individual to work ceaselessly for the welfare of others.”

This is an obscure argument. Who or what is supposed to violate justice when the individual is required to promote the overall good? Heyd does not say, and neither morality nor the individual himself seems an especially promising candidate for blame. Is it then *society* that acts unjustly—by requiring the promotion of the good, backing its requirements with social sanctions? This seems more promising: we have a straightforward case of imposing sacrifices on someone for the sake of others (*exactly* like punishing the innocent man). Other passages support this interpretation as well (e.g., pp. 166 and 176–77); but if this is indeed what Heyd has in mind it is simply beside the point. The immorality of *imposing* a sacrifice upon an individual is completely compatible with that individual’s nonetheless being morally *required* to take the sacrifice upon himself. In effect, this third argument suffers from a confusion I mentioned earlier: the impermissibility of society imposing sacrifices might well be relevant to demonstrating that consequentialism *permits* too much; but it seems irrelevant to the question of whether consequentialism *demands* too much. Thus, even if it could be shown that it is unjust for society to coerce individuals into promoting the good, this would do nothing at all to support Heyd’s view that the individuals themselves are free of such a moral requirement. (Heyd himself levels a similar objection against others; e.g., pp. 109–10.)

The Intrinsic Value of Supererogatory Acts (pp. 175–76). Since supererogatory acts go beyond what is required, their performance is “purely optional.” Heyd believes that this makes such acts especially valuable. If he is right, this would give us a reason to reject a general requirement to promote the good. But what is this valuable feature? Heyd informs us that “some types of virtuous behaviour can be realized only under conditions of complete freedom and would be stifled under a more totalitarian concept of duty.” This claim is never substantiated, however, for Heyd gives no clear statement of which virtues he has in mind. He does note that supererogatory acts display “individual preferences and virtues”; but

surely doing one's duty can do the same. He also notes that the optionality of supererogation allows favoritism and partiality—but if these are “virtues” it is not clear to me that we should mourn their passing.

A more interesting possibility may be hinted at in Heyd's repeated suggestion that the special value of supererogatory acts consists in “their being totally optional and voluntary” (pp. 9; cf. pp. 18–19, 41–2, 53, 133). He never explains why this gives these acts special worth (other than the unhelpful allusions to virtue), but one possibility seems close to hand: since the supererogatory act goes beyond what is required, if an agent makes a sacrifice for another, he is doing it not because he must, but because he *wants to*. That is, his act is not done from duty—but out of concern for the one he aids. Thus Heyd may believe that supererogatory acts are (typically) done from an especially valuable kind of *motive* (cf. p. 177), which would be lost if duty were more encompassing.

Despite the initial appeal of this line of thought, however, it essentially depends on the assumption that if promotion of the overall good were morally required it would be impossible (or harder) to act out of this higher kind of motive. But there is no reason to believe such a claim. Being morally required to aid, after all, is perfectly compatible with being motivated by direct concern for those in need.

It seems, then, that Heyd's arguments repeatedly miss their mark: of the four we've examined, two are incompatible with Heyd's own views, one is irrelevant to his conclusion, and the last fails to deliver on its promise. In the end, I believe that all that can be salvaged from Heyd's discussion is the general suggestion that an adequate defense might contain both positive and negative elements: e.g., negatively, some considerations might point to the necessity of limiting moral requirements; positively, other considerations might indicate the desirability of doing so. Obviously, however, this suggestion does not itself indicate how this necessary defense might be provided.

In his closing paragraph, Heyd expresses doubts about whether basic disagreements in moral outlook can be “resolved by rational argument”; and he asserts that he has not attempted to “prove” his “picture of man and of the nature of morality” (p. 183). Surely, however, we can at least ask that Heyd paint a coherent picture, instead of a series of inadequate sketches. It should be noted, furthermore, that the faults of Heyd's discussion are typical of most writing in this area. Although few philosophers

offer as *many* arguments as Heyd does, generally the particular arguments are just as weak. In the following sections, however, I will consider two works which are notable exceptions: each offers a sustained argument for the rejection of a general requirement to promote the good.

II

Thomas Nagel's "The Limits of Objectivity" is a rich and suggestive work in metaphysics and ethics. A series of three lectures, the first treats subjectivity and objectivity in the philosophy of mind, the second discusses the objectivity of values, and the third argues for the inadequacy of consequentialism.

In the second lecture, Nagel makes a key distinction between *agent-relative* and *agent-neutral* reasons (pp. 101–3). Roughly, an agent-neutral reason "is a reason for *anyone* to do or want something"—that is, it is a reason that applies to everybody, regardless of their particular circumstances or interests. If we judge that some state of affairs has agent-neutral value, "that means that *anyone* has reason to want it to happen." Agent-relative reasons, however, apply only to particular individuals; if we judge that something has agent-relative value, then we only believe "that someone has reason to want and pursue it if it is related to him in the right way (being in *his* interest, for example)." Much of the second lecture is devoted to arguing that there are at least *some* agent-neutral reasons; Nagel believes, for instance, that "*anyone* has a reason to want *any* pain to stop, whether or not it is his" (p. 108).

This is, of course, a controversial claim, but it won't be examined here, for it would certainly be accepted by the consequentialist. Indeed, it seems that consequentialism can be usefully viewed as the theory that the *only* reasons for action are agent-neutral ones (p. 119). Against this position, Nagel argues in the third lecture for the existence of genuine agent-relative reasons. "Deontological reasons" fall outside our topic. The existence of "reasons of autonomy," however, is directly relevant to the charge that consequentialism demands too much, for such reasons "would limit what we are *obliged* to do in the service of agent-neutral values" (p. 120). They spring from "the desires, projects, commitments, and personal ties of the individual agent, all of which give him reasons to act in the pursuit of ends that are his own" (p. 120).

Nagel writes:

If I have a bad headache, anyone has a reason to want it to stop. But if for instance I badly want to become a first-rate pianist, not everyone has a reason to want me to practice. *I* have a reason to want to practice, and it may be just as strong as my reason for wanting my headache to go away. But other people have very little reason, if any, to care whether I become a first-rate pianist or not. Why is this?

I think it is easier to believe in this distinction than to explain it (pp. 121–22).

It is indeed easy to believe in Nagel's distinction. But the question, as Nagel seems to recognize, is whether or not he can offer an adequate explanation of the distinction, and thus manage to defend our intuitions.

Before investigating this, however, it may be worth spelling out exactly how such agent-relative reasons bode ill for consequentialism's requirement to promote the overall good (a topic Nagel rushes over, pp. 124–25). Suppose, first, that my projects and interests did generate agent-neutral reasons, and only agent-neutral reasons. Since my projects would have agent-neutral value, not only would I have reasons to promote them, *everyone* would have such reasons; the promotion of my projects would count as part of the overall good. Since my projects would generate only agent-neutral reasons, however, I would have no *more* reason to promote them than would anyone else; the mere fact that they were *my* projects would give me no extra reason to favor them.

Suppose instead that my interests also generated agent-relative reasons (in addition to the agent-neutral reasons, if any). Then *I* would have reasons to promote my projects which not everyone else possessed; and if sufficiently strong, these extra reasons to promote *my* projects would override the agent-neutral reasons impinging on me. Thus it would be permissible, in some cases, to pursue my own projects at the possible expense of the overall good.

Nagel concentrates on denying that there is any agent-neutral value to the satisfaction of preferences per se (p. 124; there may be derivative value, since the frustration of desires can be painful, p. 123). This enables him to offer an indirect argument: since it is implausible to hold that *no* genuine reasons are generated, it must be agent-relative ones that are created by preferences (p. 125). Even if sound, I do not think such a *via negativa* will be especially illuminating. It may help establish the exist-

ence of agent-relative reasons, but it will not explain them. Furthermore, if the account in my two previous paragraphs is correct, the denial that there is any agent-neutral value to preference satisfaction per se seems an unnecessarily extreme thesis for Nagel to maintain, for the permissibility of favoring my own projects will be established so long as my preferences generate agent-relative reasons. There is no need to deny that they also generate agent-neutral reasons; it is simply that these must not exhaust their reason-giving force.³ What is especially in need of explanation, then, is why agent-relative reasons are created. Unfortunately, given the needs of Nagel's indirect argument, the discussion focuses instead on explaining why agent-neutral reasons are *not* generated.

These are, perhaps, relatively minor complaints; but there is a more central difficulty with Nagel's account. Nagel's explanation of why it is sometimes permissible for an agent to promote his own interests, rather than sacrificing them for the overall good, is that agent-relative reasons will often outweigh the opposing agent-neutral reasons. Such an account, however, makes it mysterious how it could be permissible for the agent to *make* the sacrifice. Doing so, after all, would be in blatant disregard of the agent-relative reasons which, by hypothesis, outweigh the agent-neutral ones. In order to maintain the intuitive claim that the agent can make the sacrifice or refrain, as he chooses, it seems that Nagel will need a more complicated account of reasons than the one given here.

Nagel's problem is the complement of one that plagued Heyd. Heyd claimed that reasons of autonomy and reasons to promote the overall good are incommensurable. This enabled him to account for the agent's freedom to sacrifice his interests, or to refrain, as he chooses; but it ruled out the possibility of accounting for moral *requirements*. Nagel views agent-relative and agent-neutral reasons as commensurable, and thus can account for the existence of moral requirements; but he seems unable to account for the agent's ever having the *freedom* to sacrifice his interests if he chooses. It is far from clear whether a coherent account of reasons could manage to avoid both of these problems simultaneously.

Leaving this difficulty aside, let us now return to the earlier question:

3. If my projects lack agent-neutral value, then others have *no* direct reason to help promote them. This would provide a second way in which morality would be less demanding than most consequentialists believe it to be. But it would not show the inadequacy of consequentialism's requirement to promote the good; it would simply establish that it is an error to believe that preference satisfaction per se is *part* of the overall good.

Why does pain generate agent-neutral reasons, but the desire to be a pianist only agent-relative ones? Nagel distinguishes between involuntary desires, and those desires that are “adopted” or chosen (p. 122). His view is that, roughly, involuntary desires generate agent-neutral reasons, but adopted desires create only agent-relative ones. (This view was tacitly appealed to in the earlier argument for the agent-neutral value of the avoidance of pain, p. 109.)

Against those who claim that the satisfaction of an adopted desire has agent-neutral value, Nagel objects that “one would have to be deranged to think it *did* matter impersonally” (p. 123). Much of the rhetorical force of this reply, however, derives from the mistaken view that objects with agent-neutral value must be “good or bad in themselves” (p. 119; cf. p. 109). To say that an object is good in itself is presumably to say that its value does not depend upon anyone’s desiring it; and it is admittedly much more plausible to hold that objects of adopted desires (typically) have value only because someone does desire them. But Nagel apparently overlooks the possibility that although an object’s having value at all may depend upon its being desired by someone or the other, given that there *is* someone who desires the object, this generates agent-neutral reasons. This view is not at least obviously crazy, and so Nagel still needs to explain why adopted desires generate only agent-relative reasons.

In what appears to be the key passage, Nagel writes that when we look at adopted desires

objectively, from outside, we can acknowledge the validity of the reasons they give for action, without judging that there is an agent-neutral reason for any of those things *to be done*. That is because, when we move to the objective standpoint, we are not *occupying* the perspective from which these values have to be accepted. Their diversity and their dependence on the history and circumstances of the agent insures this. From a point of view outside the perspective of my ambition to become a first-rate pianist, it is possible to *recognize* and *understand* that perspective and so to *acknowledge* the reasons that arise inside it; but it is not possible to *accept* those reasons as one’s own, unless one adopts the perspective rather than merely recognizing it (pp. 122–23).

I must confess that I am unable to grasp what the argument or explanation is supposed to be here. The passage repeatedly asserts that if I do

not share an adopted desire, I do not have—and indeed cannot have—a direct reason to promote the satisfaction of that desire. But I simply cannot see how Nagel gives any *explanation* of this. The trivial reminder that I do not share desires I do not share cannot in itself explain why such desires do not generate reasons binding upon me.

One possibility is that Nagel is tacitly assuming that an individual can have a reason to promote something only if he takes a (logically prior) interest in it. This would explain why unshared desires do not generate reasons; but such a Humean view seems to be in violation of the entire spirit of Nagel's enterprise, and is explicitly rejected elsewhere (e.g., p. 110).

A second possible explanation might stress the difference between involuntary desires—which are simply “evoked” (p. 109)—and adopted desires, which are chosen. Since in some sense it is my *fault* that I have adopted desires which may go unsatisfied, it might be suggested that this explains why others have no responsibility for helping me in this regard. (Compare the view that the guilty and the lazy do not deserve help in overcoming the troubles they have brought upon themselves.) I do not know if this account can be adequately developed and defended, but it would clearly go beyond anything Nagel has suggested.

Ultimately, then, Nagel's discussion is inadequate. Yet despite the failure in *detail* of Nagel's argument, I think that its general thrust is in the right direction. Nagel is concerned about the conflict between the objective and the subjective perspectives, and he is eager to emphasize the legitimate claims of the latter. Surely if an adequate defense can be provided for rejecting a general requirement to promote the overall good, that defense will be grounded in the existence and nature of the subjective standpoint. Nonetheless, it seems to me that “The Limits of Objectivity” does *not* provide an adequate argument along these lines. The most promising development of the Nagelian strategy, I believe, can be found instead in our final work.

III

Samuel Scheffler's *The Rejection of Consequentialism* is one of the most interesting works of moral philosophy that I have read in years. Starting from the reasonable view that “the salient features of *all* moral concep-

tions stand in need of principled motivation" (p. 121), the book is, as the subtitle indicates, "a philosophical investigation of the considerations underlying rival moral conceptions."

Three particular features engage Scheffler's attention: distribution-sensitive theories of the good; deontological restrictions; and what Scheffler calls *agent-centered prerogatives* which (within limits) permit the agent to perform acts, if he so chooses, that are less than optimal from an impartial perspective. Only the last will concern us here. Scheffler believes that he can offer a plausible rationale for rejecting consequentialism's general requirement to promote the good. As I've indicated, this rationale turns on the nature of the subjective or personal point of view.

Before turning to the rationale itself, let's look at the kind of agent-centered prerogative Scheffler finds most plausible. Consequentialism insists that an agent should act in accord with an impartial standpoint. In contrast, then, an agent-centered prerogative would make it permissible for the agent "to devote energy and attention to his projects and commitments out of proportion to their weight in the impersonal calculus" (p. 14). This basic idea is straightforward enough; but the particular passage where Scheffler describes the details of his prerogative is rather garbled. He suggests that a plausible prerogative

would allow each agent to assign a certain proportionately greater weight to his own interests than to the interests of other people. It would then allow the agent to promote the non-optimal outcome of his choosing, provided only that the degree of its inferiority to each of the superior outcomes he could instead promote in no case exceeded, by more than the specified proportion, the degree of sacrifice necessary for him to promote the superior outcome (p. 20).

Taken literally, this passage seems to say that an act is permissible when the difference between two (obscure) magnitudes is less than or equal to some specified proportion. It is not at all clear what intuitive interpretation this formula could have.

Scheffler's lapse is unfortunate, for it does not seem difficult to understand the sort of prerogative that he probably meant to be describing. Imagine that I want to perform some act, *S*, rather than an alternative, *O*, because *S* is more in my interests. Consequentialism says I must weigh the interests of others just as heavily as my own; thus I can perform *S* rather than *O* only if the loss to others doesn't objectively *outweigh* the

gain to me. Suppose, however, that I am permitted to count my own interests more heavily than others; in my calculations, let us say, I can magnify my benefits and losses, giving them up to M times their objective weight. Unlike consequentialism, where the size of the loss to others (of S rather than O) must be less than or equal to the gain to me, under such an agent-centered prerogative performing S rather than O would be permissible even in cases where the loss to others *does* outweigh the gain to me, provided that the size of the loss to others is less than or equal to M times the gain to me.

I believe that this is the sort of prerogative that Scheffler had in mind: I can justify certain nonoptimal acts, by giving my interests up to M times their objective weight; but I am not *required* to weigh my interests more heavily, and so I can still choose to sacrifice my interests and perform the optimal act. Furthermore, such a prerogative would differ from egoism, as Scheffler notes (p. 21),⁴ for I would *not* be permitted to pursue my interests when the size of the loss to others would be greater than M times the gain to me.

There is, however, an important difficulty that Scheffler overlooks (pp. 23–5): such a prerogative will not only permit agents to *allow* harm, it will also permit agents to *do* harm in the pursuit of their nonoptimal projects (Scheffler readily permits optimal harmings). For the prerogative is only sensitive to the *size* of the loss to others, and not to whether the loss is caused by the agent's act. Thus, for example, it will apparently be permissible to kill my rich uncle in order to inherit \$10,000. Lest it be suggested that a plausible M will avoid this result, bear in mind that most of us believe we would not be required to *pay* \$10,000 in order to *save* the life of some stranger; any M large enough to save such results will obviously work in the former case as well. Such permissions to do non-optimal harm could be blocked by introducing deontological restrictions against harming, but Scheffler rejects these; nor does he offer any other rationale for limiting the prerogative to cases of *allowing* harm.

Let us now consider the proposed rationale for agent-centered prerogatives (of whatever form). Scheffler's exposition proceeds dialectically (see especially pp. 56–67), but the final outline of the argument seems to be this: 1) Each person has a point of view, the nature of which is

4. Provided that M is finite. Scheffler is wrong, however, to claim that such a prerogative would place restrictions on the *kind* of projects the agent can pursue (see p. 21; cf. pp. 18–19).

such that the individual's "own projects and commitments have a distinctive claim on his attention; he cares about them out of proportion" to their weight from an impersonal standpoint. In this sense, the personal point of view is "independent" of the impersonal standpoint (pp. 56–7). 2) Since an adequate morality must take account of the nature of persons, we should reject any moral principle "which ignores the independence of the personal point of view" (pp. 57–8). 3) There are, however, at least two different rational methods for taking account of the fact of personal independence (pp. 60, 63): (a) a "maximization" strategy, which is available to the consequentialist (pp. 58–61); and (b) the "liberation" strategy, which creates an agent-centered prerogative (pp. 61–2). 4) Since the liberation strategy is at least as adequate as the maximization strategy, we have a rationale for incorporating prerogatives: doing so "embodies a rational strategy for taking account of personal independence, given one construal of the importance of that aspect of persons" (p. 67).

Scheffler's conclusion is modest: he does not argue that it's *necessary* to include such prerogatives, but only that there is a plausible rationale for those theories which do (pp. 64–7). Whether this is so, of course, depends on the details of the liberation strategy; for the mere fact that a theory can be construed as *some* sort of response to the nature of persons does not in itself show that there is a rationale for responding in that way.

Consider, first, the maximization strategy. Despite its name, the core of this response is the recognition that personal independence "fundamentally affects the character of human fulfilment and hence the constitution of the individual good" (p. 60). Since abandoning one's projects can be especially hard, it is only plausible to "count the cost of such hardships in arriving at our overall assessments of relevant outcomes, thereby acknowledging the special concern people have for their projects *as their* projects" (p. 59). Consequentialists (whether correctly or not) go on to require the promotion of the overall good; but this sophisticated account of individual good should be accepted even by nonconsequentialists (cf. pp. 63–4, 123–24). There may be additional appropriate responses to personal independence as well, but at the very least the rationale for adjusting one's theory of the good seems clear and undeniable.

Matters are less clear with the liberation strategy. Scheffler claims that an agent-centered prerogative "takes account of the natural independence of the personal point of view precisely by granting it moral inde-

pendence” (p. 62). Such a claim is misleading, however. Admittedly, any prerogative would give the personal point of view some freedom in practical deliberation from the constraints of the impersonal standpoint; nonetheless, like consequentialism, a limited prerogative is “unlikely to exhaust” (p. 61) the agent’s own feelings about his projects. Thus nothing short of egoism would actually grant genuine moral independence to the personal point of view. If, like Scheffler, we want to reject egoism, we need a rationale for prerogatives that grant only *partial* moral independence.

Furthermore, Scheffler’s central description of the rationale underlying the liberation strategy (p. 62) is obscure; and although other passages shed some light (pp. 64, 94, and 125–27), I believe the account remains fundamentally inadequate. In the liberation strategy, says Scheffler, the importance of the natural independence of the personal point of view “is conceived as stemming primarily from its impact on the character of human agency and motivation” (p. 94)—that is, people do not typically act in accordance with the impersonal standpoint. Scheffler suggests that “given *this* conception of the importance of the natural fact of personal independence, a moral view gives sufficient weight to that fact only if it *reflects* it, by freeing people from the demand that their actions and motives always be optimal from the impersonal perspective” (p. 62).

Thus, faced with the fact that people typically *don’t* promote the overall good, the liberation strategy responds that morally they’re not required to. But what is the underlying rationale for this response supposed to be? Surely Scheffler doesn’t mean to be arguing that since people are going to do something anyway we might as well say that this is morally permissible—a quick road to egoism, and implausible to boot. Personal independence may constitute an implicit appeal for agent-centered prerogatives—but what is the rationale for *granting* this appeal? (Surely not the mere fact that the appeal is made.) Doing so is a response, to be sure, but why is it a *rational* response? Scheffler never raises these questions; but without answers, I don’t see how we can accept his claim to have shown that the liberation strategy is a rational method for taking account of personal independence. Unlike the maximization strategy, Scheffler provides no genuine rationale at all for the liberation strategy.

Let me quickly suggest two possible approaches to providing such a rationale. The first would develop Scheffler’s observation about the impact of personal independence on motivation and agency: it might be sug-

gested, as a result, that a general requirement to promote the good would lack the motivational underpinning necessary for genuine moral requirements; and so moral theory must grant at least some sort of agent-centered prerogative. The second would stress the importance of personal independence for the existence of commitments and close personal relations: it might then be suggested that the value of such commitments yields a positive reason for preserving within moral theory at least some moral independence for the personal point of view. Neither of these approaches need lead to egoism; but I am not sure whether either of them can be adequately defended. Without some such approach, however, it seems that Scheffler fails to provide an underlying rationale for incorporating agent-centered prerogatives into moral theory. Like Nagel and Heyd, Scheffler's rejection of the general requirement to promote the good remains inadequately supported.

IV

In this essay I have examined three recent attempts to defend the view that (within limits) it is permissible for agents to pursue their own projects rather than the overall good. That all three attempts have proven inadequate should give pause to those who do not even feel the need to defend this common view. The shortcomings of a grab bag approach like Heyd's will not, perhaps, surprise us; but the failure of even the sustained attempts of Nagel and Scheffler should alert us to the difficulty of supporting the common belief. Indeed, if the intuition that consequentialism demands too much remains impossible to defend, we may have to face the sobering possibility that it is not consequentialism, but our intuition, that is in error.

I am indebted to Thomas Nagel and Samuel Scheffler for valuable discussions of their work, and to the Editors of *Philosophy & Public Affairs* for criticisms of an earlier version of this essay.