



OXFORD JOURNALS
OXFORD UNIVERSITY PRESS

Mind Association

Agent-Centred Restrictions, Rationality, and the Virtues

Author(s): Samuel Scheffler

Source: *Mind*, New Series, Vol. 94, No. 375 (Jul., 1985), pp. 409-419

Published by: [Oxford University Press](#) on behalf of the [Mind Association](#)

Stable URL: <http://www.jstor.org/stable/2254812>

Accessed: 05/11/2014 16:25

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at
<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Oxford University Press and *Mind Association* are collaborating with JSTOR to digitize, preserve and extend access to *Mind*.

<http://www.jstor.org>

Discussions

*Agent-Centred Restrictions, Rationality,
and the Virtues*

SAMUEL SCHEFFLER

There is no substantive moral theory that is obviously correct. All such theories stand in need of some defence. However in my book, *The Rejection of Consequentialism*,¹ I argued that the need is particularly acute in the case of typical deontological theories. For although the common-sense morality of our culture is substantially deontological in content, and although many moral philosophers find themselves drawn toward some version of deontology, I maintained that there is a distinct air of paradox surrounding such views. And this mixture of real appeal and apparent paradox—always a potent combination in philosophy—lends a special urgency to the defence of deontology.

That typical deontological views are apparently paradoxical, I argued, is to be explained by their inclusion of what I call ‘agent-centred restrictions’. An agent-centred restriction is, roughly, a restriction which it is at least sometimes impermissible to violate in circumstances where a violation would serve to minimize total overall violations of the very same restriction, and would have no other morally relevant consequences. Thus, for example, a prohibition against killing one innocent person even in order to minimize the total number of innocent people killed would ordinarily count as an agent-centred restriction. The inclusion of agent-centred restrictions gives traditional deontological views considerable anti-consequentialist force, and also considerable intuitive appeal. Despite their congeniality to moral common sense, however, agent-centred restrictions are puzzling. For how can it be rational to forbid the performance of a morally objectionable action that would have the effect of minimizing the total number of comparably objectionable actions that were performed and would have no other morally relevant consequences? How can the minimization of morally objectionable conduct be morally unacceptable?

In the two published versions of her Presidential Address to the Pacific Division of the American Philosophical Association,² Philippa Foot attempts ‘to show that there is no paradox at the heart of non-consequentialist morality’.³ Foot

¹ Oxford, Clarendon Press, 1982.

² The reference for the first version is ‘Utilitarianism and the Virtues’, *Proceedings and Addresses of the American Philosophical Association* (abbreviated hereafter as *PAAPA*) 57 (1983), pp. 273–83. The second version appears, with the same title, in *Mind* 94 (1985), pp. 196–209. Foot describes it as ‘an expanded version’. She adds: ‘Much of the text is unaltered and all the ideas are the same, but I hope to have explained myself more clearly this time around.’ See *Mind*, p. 196. When quoting, I will always indicate whether or not the passage as quoted appears in both versions of Foot’s paper. Where it does, I will give both page references. Where it does not, I will give the page reference for the version in which it does appear, and I will also compare the quoted passage with the corresponding passage in the other version, if there is one. If there is no corresponding passage, I will so indicate.

³ *PAAPA*, p. 282. These words have been eliminated from the second version, but the description they provide of the aim of the paper fits both versions equally well.

agrees that agent-centred restrictions *appear* paradoxical. And she believes that consequentialism, which first gives some principle for ranking overall states of affairs from best to worst from an impersonal or agent-neutral standpoint, and then says that the right act in a given situation is the one that will produce the best overall outcome of any act available, has a 'spellbinding force'.⁴ But she also believes that a certain kind of non-consequentialist moral view can in the end be shown to be free of paradox despite the fact that it includes agent-centred restrictions, and that the spell of consequentialism can thus be broken. The kind of moral view she has in mind is one in which a conception of the virtues plays a central role. Now many of what I have been calling traditional deontological views do not assign this kind of role to the virtues. Indeed, so-called 'virtue theories' are often thought to represent an alternative to both consequentialist and deontological moral conceptions. For the purposes of this discussion, therefore, it is important to remember that Foot's claim is, in effect, that agent-centred restrictions are not paradoxical when they are set in the context of a non-consequentialist view of a certain kind. There will be occasion later in this paper to consider the extent to which assignment of a central role to the virtues really is essential to the sort of defence of agent-centred restrictions that Foot wants to give.

Foot says that what seems compelling about consequentialism is 'the rather simple thought that it can never be right to prefer a worse state of affairs to a better'.⁵ And what seems paradoxical about those non-consequentialist views that include agent-centred restrictions is that they appear to claim that it is sometimes morally impermissible to produce the best state of affairs that one is in a position to produce. Sometimes, they seem to say, we must do less good, or prevent less evil, than we could. Perhaps, for example, we must refrain from harming one innocent person even if harming that person would result in the minimization of the total number of innocent people comparably harmed. That consequentialism should seem compelling, and that agent-centred restrictions should seem paradoxical, Foot believes, is inevitable once we grant the apparently innocent idea 'that there *are* better and worse states of affairs in the sense that consequentialism requires'.⁶ But, she maintains, this idea is really not so innocent; it can be challenged, and it is through such a challenge that she hopes to break the spell of consequentialism and dissolve the air of paradox surrounding agent-centred restrictions.

Foot does not claim, as some others have, that evaluations of states of affairs never make sense in moral contexts. On the contrary, she thinks it is important 'to see the place that there indeed is *within* morality for the idea of better and worse states of affairs'.⁷ 'That there is such a place', she adds, 'follows from the fact that the proper end of benevolence is the good of others, and that in many situations the person who has this virtue will be able to think of good and bad states of affairs in terms of the general good.'⁸ Thus, for example, if there is 'a question of riding out to rescue

⁴ *PAAPA*, p. 274; *Mind*, p. 198.

⁵ *PAAPA*, p. 275; *Mind*, p. 198.

⁶ *PAAPA*, p. 275; *Mind*, p. 199.

⁷ *PAAPA*, p. 281. In the *Mind* version, the quoted material has been slightly altered: 'to see that there is indeed a place *within* morality for the idea of better and worse states of affairs' (*Mind*, p. 206).

⁸ *PAAPA*, pp. 281–2. In *Mind*, the corresponding passage reads: 'That there is such a place is true if only because the proper end of benevolence is the good of others, and because in many situations the person who has this virtue will be able to think of good and bad states of affairs, in terms of the general good' (*Mind*, p. 206).

a small number or a large number then benevolence would urge that the larger number be saved'.⁹ What Foot wants to argue, however, is the following. Although someone who possesses the virtue of benevolence will indeed be disposed to promote good states of affairs in certain circumstances, benevolence is not the only virtue. Justice, for example, is also a virtue. And there are various rules and requirements that the person who possesses the virtue of justice must observe: rules of distributive justice, truth-telling, respect for rights, and so on. Rules and requirements such as these restrict the area 'in which benevolence is free to pursue its ends';¹⁰ for 'sometimes justice will forbid a certain action, . . . and then it will not be possible to ask whether "the state of affairs" containing the action and its result will be better or worse than the one in which the action is not done. The action is one that *cannot* be done because justice forbids it, and nothing that has this moral character comes within the scope of the kind of comparison of total outcomes that benevolence may sometimes require'.¹¹

Now by itself the claim that an unjust action falls outside 'the scope of the kind of comparison of total outcomes that benevolence may sometimes require' is not entirely unambiguous. One might wonder whether it means that no meaningful comparison of outcomes is *possible* in cases where one of the outcomes would result from an unjust action, or whether it means instead that since one must not perform the unjust action in any case, it is inappropriate actually to *carry out* the relevant comparison of overall outcomes. However, the following passage from the earlier version of Foot's paper suggests that it is the first interpretation that more nearly reflects her thinking:

When we . . . give expressions such as 'best outcome' and 'good state of affairs' no special meaning in moral contexts other than the one that the virtues give them, we shall no longer think the paradoxical thought that it is sometimes right to act in such a way that the total outcome, consisting of one's action and its results, is less good than some other accessible at the time. What the non-consequentialist should say is that 'good state of affairs' is an expression which has a very limited use in these contexts. It belongs in cases in which benevolence is free to pursue its ends, and chooses among possibilities . . . But the expression has no meaning when we try to use it to say something about a whole consisting of what we would illicitly do, allow, or wish for, together with its consequences. In the abstract a benevolent person must wish that loss and harm should be minimized. He does not, however, wish that the whole consisting of a killing to minimize killings should be actualized either by his own agency or that of anyone else. So there is no reason on this score to say that he must regard it as the 'better state of affairs'. And therefore there is no reason for the non-consequentialist, whose thought of good and bad states of affairs in moral contexts comes only from the virtues themselves, to describe the refusal as a choice of a worse state of affairs. If he does so describe it he will be giving the words the sense they have in his opponents' theories, and it is not surprising that he should find himself in their hands.¹²

⁹ *PAAPA*, p. 282. In the *Mind* version, the phrase 'we may suppose' has been inserted between 'would' and 'urge'. See *Mind*, p. 206.

¹⁰ *PAAPA*, p. 282. This phrase does not appear in the *Mind* version, but the view I am describing surely does.

¹¹ *PAAPA*, p. 282. In *Mind*, the corresponding passage reads: 'sometimes justice will forbid a certain action, . . . and then it will not be possible to ask whether "the state of affairs" containing the action and its result will be better or worse than one in which the action is not done. The action is one that *cannot* be done, because justice forbids it, and nothing that has this moral character comes within the scope of the kind of comparison of total outcomes that benevolence may sometimes require' (*Mind*, p. 206).

¹² *PAAPA*, p. 282.

The view expressed in this passage seems to be that comparisons of overall states of affairs in moral contexts can only be meaningfully made when action aimed at promoting the good of others is called for, that such action is forbidden in cases where it would transgress some rule of justice, and that in cases of this kind it is not possible meaningfully to say that the prohibited action would produce a better overall state of affairs than the alternative. Now in the later version of her paper, Foot has eliminated that portion of the passage just quoted which begins with the words ‘What the non-consequentialist should say’ and ends with the words ‘together with its consequences’, thus withdrawing the explicit claim that the expression ‘good state of affairs’ has no meaning when the outcome of an unjust act is in question. Nevertheless, the broad outlines of her position remain unchanged. She continues to maintain that while comparisons of states of affairs in moral contexts can meaningfully be made when benevolent action is called for, the claim that some unjust act would result in a better overall state of affairs than any of the available alternatives lacks any clear sense in ordinary non-consequentialist moral thought. And, she argues, while consequentialist theories may *give* it a sense, someone who has not already accepted one of those theories has no reason to believe that there *are* better and worse states of affairs in the consequentialist’s sense.¹³ Thus, Foot believes, the air of paradox surrounding non-consequentialist views that include agent-centred restrictions can be dispelled. For what seems paradoxical about those views, according to Foot, is that they appear to claim that we must sometimes produce a worse overall outcome instead of a better one. And if she is right, this appearance can be shown to be illusory. There will of course be situations in which the *consequentialist will describe the non-consequentialist* as insisting that we must produce a worse overall outcome rather than a better one. But the non-consequentialist can, if Foot is right, deny that that description has any ordinary meaning in such situations. The non-consequentialist can thus maintain that either the consequentialist is talking nonsense, or else he is supplying his words with some special meaning derived from his own theory, in which case he is begging the question against the non-consequentialist.

There seem to me, however, to be three reasons for doubting whether Foot has really succeeded in dispelling the air of paradox surrounding agent-centred restrictions. First, I am sceptical of the idea that, in ordinary non-consequentialist moral discourse, evaluations of overall states of affairs are meaningful when benevolent action is called for, but meaningless when the outcome of an unjust action is in question. People who deny that such evaluations *ever* make sense typically do so because they do not believe that the benefits and harms of different human beings can be meaningfully summed. But this worry about aggregation does not seem to be what concerns Foot, since she is happy to speak, for example, of ‘the important place that the idea of maximum welfare has in morality’.¹⁴ And, as we have seen, she wants to claim, not that evaluations of states of affairs never make sense in moral contexts, but only that they may lose their sense whenever the candidate for assessment is the outcome of an unjust act. But do we really cease to understand what is meant by ‘a better state of affairs’ if the question is raised whether infringing a right

¹³ The expression ‘good state of affairs from a moral point of view’, she writes in the second version of her paper, ‘may mean nothing; and it may lack a reference when a special consequentialist theory has given it a sense’ (*Mind*, p. 204).

¹⁴ *PAAPA*, p. 282; *Mind*, p. 206.

or telling a lie or treating a particular individual unfairly might perhaps produce a better state of affairs than failing to do so? I do not think so. Many moral dilemmas take the form of conflicts between considerations of justice, rights, or fairness on the one hand, and considerations of aggregate well-being on the other. And it seems to me quite natural to characterize the dilemmatic feature of a situation of this kind by saying, for example, that one is faced with a problem because violating someone's rights would in this case produce better results on the whole than would respecting them. I do not think that it is only consequentialists who think of matters in these terms, and unless it can be shown that there is something incoherent about *any* interpersonal aggregation of benefits and burdens, I see no reason to deny us this way of speaking and conceiving of the matter.

Second, in order for Foot's attempt to dissolve the apparent paradox surrounding agent-centred restrictions to be successful, it must be the case that the alleged paradox cannot be formulated without using the idea of one overall state of affairs being better than another. But, as my initial characterization of the paradox at the beginning of this paper was meant to suggest, it can in fact be formulated without using the notion of an 'overall state of affairs' at all. How, I asked, can it be rational to forbid the performance of a morally objectionable action that will have the effect of minimizing the total number of comparably objectionable actions that are performed and will have no other morally relevant consequences? How can the minimization of morally objectionable conduct itself be morally unacceptable? Even if, for the sake of argument, we grant Foot's claim that the idea of one overall state of affairs being better than another lacks any clear non-consequentialist sense in cases of injustice, *these* questions can still be formulated and understood, and the answers to them still do not seem obvious. Even if Foot's claim is granted, the defender of agent-centred restrictions can hardly say that it is meaningless to assert that circumstances can arise in which a certain moral rule will be violated several times unless I violate it once. And while he can, if Foot's claim is granted, deny that it is meaningful to say that the state of affairs containing several violations is worse than the state of affairs containing just one violation, I do not believe that we need the latter claim in order to see the agent-centred prohibition as puzzling. All we need is the recognition that fewer violations will occur if I act one way rather than another, together with the idea that such violations are morally objectionable, in the rather unambitious sense that it is morally preferable that no such violations should occur than that any should. And while Foot may in fact want to reject even this weaker idea, I believe, as I shall argue in a few pages, that the costs of doing so are prohibitive.

Third, although Foot begins her paper by acknowledging that 'utilitarianism tends to haunt even those of us who will not believe in it',¹⁵ and although her paper is meant as an 'exorcism',¹⁶ an attempt to rid consequentialism of its 'spellbinding force', the way in which she ultimately tries to do this is such as to make it seem mysterious how consequentialism was ever taken seriously in the first place, let alone viewed as spellbinding. For if, in asking how it can ever be right 'to prefer a worse state of affairs to a better',¹⁷ the consequentialist is either talking nonsense or else using the language of his own theory instead of the language that the rest of us

¹⁵ *PAAPA*, p. 273; *Mind*, p. 196.

¹⁶ *PAAPA*, p. 273; *Mind*, p. 196.

¹⁷ *PAAPA*, p. 275; *Mind*, p. 198.

speak, how is it that we find his question troubling, haunting? After all, if Foot is right, it is not clear that we even understand the question. So wherein lies its power to haunt us? I do not believe that Foot's view allows any adequate answer to this question, and for this reason if for no other her position seems to me worrisome.

Although I do not agree with the idea that attempts to make agent-centred restrictions seem paradoxical are question-begging, or with the idea that we will find views that include such restrictions paradoxical only if we have already conceded the truth of consequentialism in accepting the description of those views that is supposed to generate the difficulty, I think I understand one reason why these ideas seem tempting. Moreover, although I do not agree with them, I think that there is *something* right about them, and that in an appreciation of what is right about them lies the key to any adequate defence of agent-centred restrictions. These ideas seem tempting partly, I believe, because we have the sense that in finding the restrictions paradoxical, we are relying on a conception of rationality that seems to lie at the heart of consequentialism, and that if accepted seems inevitably to make the restrictions look problematic. And there is a way in which this is right. The reason that it is nevertheless not question-begging to say that the restrictions seem paradoxical is that although the conception of rationality that generates the appearance of paradox lies at the heart of consequentialism, it is not peculiar to consequentialism. On the contrary, it is a fundamental and familiar conception of rationality that we accept and operate with in a very wide and varied range of contexts. The fact that this powerful conception of rationality seems both to lie at the heart of consequentialism and to generate the sense that agent-centred restrictions are paradoxical does not show that the restrictions will only seem paradoxical to us if we have already, wittingly or unwittingly, accepted consequentialism. It shows rather that the 'spellbinding force' of consequentialism, its capacity to haunt even those who do not accept it, derives from the fact that it appears to embody a notion of rationality which we recognize from myriad diverse contexts, and whose power we have good independent reason to respect. It also shows that the seeming paradox of agent-centred restrictions goes deep; no questions need be begged to find the apparent clash between the morality of common sense and the rationality of common sense troubling, haunting, difficult to ignore or dismiss. At the same time it suggests that a fully satisfying defence of agent-centred restrictions could take one of two forms. It might, first, consist in showing that the conflict between such restrictions and the kind of rationality they seem to defy is only apparent: that, appearances to the contrary notwithstanding, the restrictions can be reconciled with that familiar form of rationality. Or it might, alternatively, consist in showing that the restrictions embody a limitation on the scope of that form of rationality, and give expression to a different form of rationality which we also recognize and which also has its place in our lives.

The kind of rationality that consequentialism seems so clearly to embody, and which makes so much trouble for views that incorporate agent-centred restrictions, is what we may call *maximizing* rationality. The core of this conception of rationality is the idea that if one accepts the desirability of a certain goal being achieved, and if one has a choice between two options, one of which is certain to accomplish the goal better than the other, then it is, *ceteris paribus*, rational to choose the former over the latter. Consequentialism seems to embody this kind of rationality because it starts from a conception of what is desirable (the overall good) and then tells us always to

promote as much of it as we can. Views that incorporate agent-centred restrictions, by contrast, seem troubling, relative to this notion of rationality. For they appear to identify certain kinds of actions as morally objectionable or undesirable, in the sense that it is morally preferable that no such actions should occur than that any should, but then tell us that there are situations in which we must act in such a way that a greater rather than a lesser number of these actions are actually performed.

There is, of course, nothing within maximizing rationality itself that requires us to accept the consequentialist's choice of goals, and so although consequentialism embodies that form of rationality, it is not the only normative theory of action that does so. For example egoism, construed here as the view that one ought always to pursue one's *own* greatest advantage, also embodies maximizing rationality. Indeed common-sense deontological morality, standing between egoism and consequentialism, sometimes seems to be caught in a kind of normative squeeze, with its rationality challenged in parallel ways by (as it were) the maximizers of the right and of the left: those who think that one ought always to pursue one's own good, and those who are convinced that one should promote the good of all.

I said a moment ago that a satisfying defence of agent-centred restrictions could take one of two forms. The first would be to show that, appearances to the contrary notwithstanding, there really is no conflict between such restrictions and maximizing rationality. Thus it might be denied, to start with, that views incorporating agent-centred restrictions actually do present as desirable any goal whose maximum accomplishment they then prohibit. They assign each person the agent-relative goal of not violating any restrictions himself, it might be said, but they do not present the overall non-occurrence of such violations as desirable. Thus in forbidding the minimization of overall violations, they are not in fact thwarting the achievement of any goal whose desirability they recognize.

Now I do not believe that defenders of standard *deontological* views are really in a position to make these claims. The difficulty is that such views do, as I have suggested, seem committed to the idea that violations of the restrictions are morally objectionable or undesirable, in the sense that there is a moral point of view from which it is preferable that no violations should occur than that any should. Defenders of deontological views are typically happy to say things like this, and with good reason. For on standard deontological views, morality evaluates actions from a vantage point which is concerned with more than just the interests of the individual agent. In other words, an action will be right or wrong, on such a view, relative to a standard of assessment that takes into account a number of factors quite independent of the interests of the agent. And defenders of such views are unlikely to claim that the relevant standard of assessment includes agent-centred restrictions, but that it is a matter of indifference, from the vantage point represented by that standard, whether or not those restrictions are violated. For if it is not the case that it is preferable, from *that* vantage point, that no violations should occur than that any should, it is hard to see how individual agents could possibly be thought to have reason to observe the restrictions when doing so did not happen to coincide with their own interests or the interests of those they cared about. In other words, deontological views need the idea that violations of the restrictions are morally objectionable or undesirable if the claim that people ought not to commit such violations when doing so would be in their own interests is to be plausible. Yet if such views do regard violations as morally objectionable or undesirable, in the sense

that it is morally preferable that none should occur than that any should, it does then seem paradoxical that they tell us there are times when we must act in such a way that a larger rather than a smaller number of violations actually takes place. Notice that egoism, by contrast, *does* seem committed exclusively to agent-relative goals. It assigns each person the agent-relative goal of maximizing his own advantage. And since it does not purport to assess actions from a point of view which is concerned with more than just the interests of the individual agent, it is not committed in the way deontology is to presenting as desirable any non-relative goal whose maximum accomplishment it then prohibits. That is why *it does not for a moment seem paradoxical for the egoist to say that one ought to maximize one's own advantage even if that means that fewer people overall will be able to maximize theirs.*

Thus defenders of standard deontological views do not appear to be in a position to make the claim that, in forbidding us to minimize the violation of those restrictions they insist on, they are not thwarting the achievement of any goal whose desirability they recognize. The situation may be different, however, with other kinds of non-consequentialist views. In particular, someone who accepts a view like Foot's may be in a position to make this claim more plausibly. For if agent-centred restrictions are seen as restrictions that those who possess certain virtues will be disposed to observe, and if these virtues are thought of as traits of character whose possession enables a person to live the kind of life that is good for him,¹⁸ then it may perhaps be denied that the commitment to agent-centred restrictions involves any commitment to assessing actions from a 'moral point of view' which is concerned with something more than just the interests of the individual agent.¹⁹ Such a denial would reveal a significant difference between this kind of view and standard deontological views, and it would make the assignment of a central role to the virtues essential to the defence of agent-centred restrictions; but it would also carry with it a commitment to the idea that actions are right or wrong—if at all—relative to a standard of assessment that does not ultimately take anything but the well-being of the agent into account. Thus, perhaps, what would be wrong with injustice, lying, and the like would be, roughly, that the disposition to engage in such activities does not contribute to a good life for the agent, and that the disposition not to does. But this, it seems to me, rather glaringly fails to capture our actual sense of what is ordinarily wrong with these things. Even if we agree that the disposition to behave unjustly does not in fact contribute to the agent's ability to live a life that is good for him, we are unlikely to agree that that is the only reason injustice is wrong. It may be objected that the kind of view under discussion is best understood as claiming, not that certain kinds of actions are wrong because the disposition to perform them does not contribute to the living of a good life by the agent, but rather that the disposition

¹⁸ I am not in fact sure that Foot herself would be prepared to say this. (See footnote 20 below.) But the argument I am imagining in defence of agent-centred restrictions depends on a willingness to say it. And as indicated in footnote 19 below, that argument seems in obvious respects to be rather in the spirit of Foot's overall position.

¹⁹ Such a denial would of course be entirely consistent with Foot's general scepticism, expressed in both versions of her paper, about the phrase 'the moral point of view'. It would, I think, also be in keeping with the spirit of the following passage from the second version: 'Perhaps no . . . *shared end* appears in the foundations of ethics, where we may rather find individual ends and rational compromises between those who have them. Or perhaps at the most basic level lie facts about the way individual human beings can find the greatest goods which they are capable of possessing' (*Mind*, p. 209).

to perform them does not contribute to the living of a good life by the agent because they are wrong (by some independent standard).²⁰ Understood in this way, however, the view loses its ability to avoid the deontologist's predicament. For it no longer claims that the standard relative to which actions are right or wrong is one that takes nothing but the well-being of those who perform them into account. It thus loses its ability to disclaim any commitment to the idea of assessing actions from a point of view which is concerned with more than just the interests of the individual agent, and hence to the idea that there is a moral point of view from which it is preferable that no violations of the restrictions should occur than that any should. And so it loses its ability to make the claim that, in forbidding the minimization of overall violations, it is not thwarting the achievement of any goal whose desirability it recognizes.

As an alternative to trying to make that claim, someone who wanted to show that there was no conflict between agent-centred restrictions and maximizing rationality might point out that, if the *ceteris paribus* clause in the formulation of maximizing rationality were fully cashed out, one of its main features would be a provision to the effect that it can sometimes be rational to act in such a way as to worse achieve one goal if that will make it possible to better achieve another. Since that is so, it might be said, views that include agent-centred restrictions need not come into conflict with maximizing rationality when they tell us to further the agent-relative goal of not violating the restrictions ourselves at the expense of the non-relative goal of minimizing violations of the restrictions. By itself, however, this claim is not fully persuasive. The problem is that the agent-relative goal and the non-relative goal appear to be related to each other in such a way as to make the insistence on giving priority to the relative goal puzzling, from the standpoint of maximizing rationality. Since, as our earlier discussion suggested, the fact that violations of the restrictions are objectionable from a moral point of view constitutes at least part of the basis for claiming that individual agents ought not ordinarily to commit such violations, the agent-relative goal looks as if it is derivative from, and given life by, the non-relative objection, and does not appear to represent something independently desirable. Rather, the desirability of achieving the agent-relative goal seems contingent on its serving to advance the non-relative goal of minimizing the morally objectionable. And if that is so, then the insistence that one must satisfy the agent-relative goal even when doing so will inhibit achievement of the non-relative goal is incompatible with considerations of maximization.

The project of reconciling agent-centred restrictions with maximizing rationality thus faces the following difficulty. On the one hand, as I have already argued, the compatibility of such restrictions with that form of rationality cannot be satisfactorily established by dispensing altogether with the idea of a moral point of view. For if one dispenses with that idea, one cannot do justice to our sense of what is

²⁰ Judging from various of Foot's other published works, it is not clear that *she* would want to make *either* of these claims. For in some of her more recent writings she has expressed increasing doubts about the closeness of the connection between one's possession of the virtues and one's good. (See, for example, her introduction to *Virtues and Vices* (Berkeley, University of California Press, 1978), her paper 'Morality as a System of Hypothetical Imperatives' (reprinted in *Virtues and Vices*), and the final footnote in the version of 'Moral Beliefs' that appears in the same volume.) Without something like the first claim, however, the defence of agent-centred restrictions that I have been sketching does not get off the ground. And as I indicated in the preceding note, Foot does seem to have more than a little sympathy for at least some elements of that defence.

ordinarily wrong with the conduct that the restrictions prohibit. Nor, as I have also argued, can the compatibility of the restrictions with maximizing rationality be established by accepting the notion that morality evaluates actions from a point of view that is concerned with more than just the interests of the individual agent, but denying that violations of the restrictions are objectionable or undesirable from that point of view. For if one accepts the former notion, then one needs the claim that violations are morally objectionable or undesirable in order to explain why individuals ought not to commit such violations when doing so would be in their own interests. On the other hand, however, the argument of the preceding paragraph suggests that, if the compatibility of agent-centred restrictions and maximizing rationality is to be established, neither can it be conceded that the *entire basis* for the restrictions is the objectionableness from the moral point of view of the behaviour they prohibit. For if one makes that concession, then the requirement that the agent-relative goal be given priority over the non-relative goal cannot be reconciled with considerations of maximization. To show that agent-centred restrictions are compatible with maximizing rationality, therefore, one must agree that the behaviour they rule out is morally objectionable or undesirable, but deny that that very objectionableness constitutes the entire rationale for the restrictions. And then, of course, one must supply the remainder of the rationale.

One idea, along these lines, would be to argue that agent-centred restrictions serve some independent maximizing purpose. Thus it might be said, for example, that the inclusion of such restrictions enables a moral conception to give more weight than consequentialism does to some important fact or consideration: some natural feature of persons, perhaps. In *The Rejection of Consequentialism* I tried to use a strategy of roughly this kind to motivate an 'agent-centred prerogative', a prerogative allowing each agent to devote energy and attention to his own projects and commitments out of proportion to the weight in any impersonal calculus of his doing so. If my argument there was correct, such a strategy can thus be used to explain why one is not always *required* to give the non-relative goal of minimizing overall violations priority over the agent-relative goal of avoiding violations oneself. At the same time, I indicated that I myself do not see how, specifically, to deploy such a strategy in defence of agent-centred restrictions: in defence of the view that one is not always *permitted* to give the non-relative goal priority over the relative one. I do not, in other words, see how to make a convincing case that there is some particular important fact or consideration to which a moral theory gives sufficient weight only if it includes agent-centred restrictions. Obviously, however, that is hardly conclusive, and this strategy continues to represent a means by which it might be possible to reconcile agent-centred restrictions and maximizing rationality, thereby dispelling the apparent paradox attached to the restrictions.

Of course, even if no reconciliation were possible, that would not show that agent-centred restrictions are indefensible. As I said earlier, a satisfying defence of the restrictions could take either of two forms. Reconciliation with maximizing rationality would be one sort of defence. But it is, after all, not obvious that maximizing rationality constitutes the whole of rationality. And if in fact there were no way to defend agent-centred restrictions while remaining within the framework of maximizing rationality, then the alternative for a defender of the restrictions would be to try to show that they embody a departure from maximization which is licensed by the more comprehensive tapestry of full human rationality. In other

words, the task would be to try to set the restrictions convincingly within the broad contours of practical rationality as we understand it.

Now it might be thought that this task could be easily dispatched. After all, if it really is true that, as I said earlier, agent-centred restrictions are congenial to the common-sense morality of our culture, and if the restrictions thus embody constraints on practical reasoning that seem to us natural and intuitively appealing, then that might be thought sufficient to show that they do in fact have their place within what we are prepared to recognize as human practical rationality, even if they represent a departure from maximization. This idea may not in fact be so very different from what Foot wishes to maintain. The difficulty with this quick solution is that the appearance that the restrictions are *irrational* is generated by an apparently appropriate application of a very powerful form of thought which itself occupies a central place within what we recognize as human practical rationality. The seeming paradox arises out of a process of reasoning that itself seems natural and intuitively compelling, and not through the introduction of some theoretically attractive but humanly unrecognizable model of rationality. Thus to dispel the paradox and give a satisfying account of the place of the restrictions within full human rationality, more must be done than simply to call attention to their naturalness and appeal. For to do no more than that is to leave in place all of those elements which combine to create the impression that, in so far as it is drawn to agent-centred restrictions, human practical reason may be at war with itself.

Viewed from one perspective, it may seem odd that agent-centred restrictions should be thought to have a specially insecure relationship to considerations of practical rationality. For such restrictions are often thought of as broadly Kantian in spirit, and it is Kant, along with Aristotle, who is most closely associated with the idea that moral norms are rooted in the structure of practical reason. The oddity may be lessened somewhat if we remember that the normative view whose rationality is in question, although standardly referred to as Kantian, represents at most one aspect of Kant's own view. Roughly speaking, we can distinguish the following elements, among others, in Kant's moral thought: a view about the nature of moral motivation (an act done purely from inclination lacks any genuine moral worth), a view about the constraints imposed by reason on the maxim of an action (the categorical imperative procedure), and a view about the substantive moral norms derivable from the categorical imperative. If there is a genuinely Kantian view being challenged here, it is this: that it is possible to interpret the categorical imperative in such a way that it is plausibly thought of both as a requirement of practical reason and as supporting agent-centred restrictions in particular. This leaves much of what Kant thought about the relation of morality and rationality untouched. At the same time, the question it does raise is one to which the answer, I think, is not at all clear.

Department of Philosophy
University of California, Berkeley
Berkeley, California 94720
U.S.A.

SAMUEL SCHEFFLER