Morality as a Rational Requirement

JULIAN BAGGINI

One of the most long-running disputes in moral philosophy concerns the relationship between ethics and rationality. On one side, there are those who, in agreement with Kant, believe that the requirement to behave according to moral principles can be established by reason alone. As Kant himself put it, 'The ground of obligation here must be sought not in the nature of the human being or in the circumstances of the world in which he is placed, but *a priori* simply in concepts of pure reason.'[1]

In the opposing camp are those who, following Hume, argue that reason can never provide the basis for morality and that at best it can serve the desires and impulses, moral or otherwise, that arise in us for non-rational reasons. On this view, as Hume put it, 'Reason is and ought only to be the slave of the passions.'[2]

John Searle has recently produced an argument in the Kantian tradition in which he attempts to show that there are 'rationally binding *desire-independent* altruistic reasons for action.'[3] These reasons are sufficient, Searle argues, to provide a rational basis for strong altruism. In this paper, I intend to show why Searle's argument fails and what its failure tells us about the wider debate between the successors of Hume and Kant.

1. Strong and weak altruism

Searle distinguishes between strong and weak altruism. Weak altruism is where a person is 'naturally inclined to care about the interests of others'.[4] This is termed 'weak' altruism because it is

[1] Immanuel Kant, *Groundwork of the Metaphysics of Morals*, in *Practical Philosophy*, Mary J. Gregor (ed. and trans.) (Cambridge: Cambridge University Press, [1797] 1996) 45 [4:389]

[2] David Hume, *A Treatise on Human Nature* (Second Edition), ed. L. A. Selby-Bigge (Oxford: Clarendon, [1739–40] 1978) 415 [Book 2, Part III, Section III].

[3] John Searle, *Rationality in Action*, (Cambridge, Massachusetts: The MIT Press, 2001) 158.

[4] Ibid., 158.

doi:10.1017/S0031819102000384 ©2002 The Royal Institute of Philosophy

https://doi.org/10.1017/S0031819102000384 Published online by Cambridge University Press

grounded in nothing more than a disposition or preference. As long as altruism is merely a matter of personal preference, we can no more hope to give any rationally compelling reasons why someone should be altruistic than we can give rationally compelling reasons why they should like the taste of beer.

A strong altruist, on the other hand, is a person who 'recognizes the interests of others as a valid reason for acting *even in cases where he has no such inclination*.'[5] (My italics.) In such cases, there are what Searle terms 'rationally binding *desire-independent* altruistic reasons for action.' The reasons are rationally binding because to disregard them requires one to act inconsistently and hence contrary to reason. They are desire-independent because they exist whether or not one wants them to exist. So, for instance, to recognize that one will contradict oneself if one does not give some of one's income to charity regardless of whether one is disposed to do so or not is to recognize that there are rationally binding desire-independent altruistic reasons to do so. More generally, Searle argues that to recognize that one has rationally binding reasons to help others whether one is disposed to do so or not is to accept the case for strong altruism.

The important thing to note here is that strong altruism requires that one accepts reasons for action which are *both* rationally binding and desire-independent.

2. Searle's argument

Searle's argument requires three steps and is illustrated with the example of being in pain. Central to the argument—and something I will not question here—is what Searle calls the 'generality requirement'. Searle defines this as 'when you make an assertion of the form *a* is *F*, rationality requires that you be able to will that everyone in a similar situation should assert that *a* is *F*.'[6] As it concerns language, the point is a broadly Wittgensteinian one: in order for my utterances to have any meaning, they must be bound by public rules of meaning, and this means that in using a word I have to accept that other people may also use that word with the same meaning. As it concerns ethics, the point is a broadly Kantian one, though it is equally accepted by anyone who sees universalisability as being an essential feature of ethics. The basic principle of universalisability

[5] Ibid., 158.

[6] Ibid., 159. Searle does not explicitly define what the generality requirement is, but this sentence seems to function as a defintion.

https://doi.org/10.1017/S0031819102000384 Published online by Cambridge University Press

is that if I say 'X ought to do Y', I am committed to the view that anyone ought to do Y in relevantly similar circumstances.

The first step in Searle's argument requires us to accept that when we say, 'I am in pain', we are 'committed by the generality requirement to recognize that in a similar situation you would be in pain.'[7] In order for me to use a word like 'pain' I have to accept the public rules for its use, and that means accepting that when someone else says they are in pain, they may mean just what I do.

The second step is to see that my pain creates a need. I am in pain so I need help. As with the first step, the generality requirement again means that I have to accept that if anyone else is in pain, their pain too creates a need. I cannot make a special case of myself: if my pain creates a need for me, then the pain of others creates a need for them.

The third and final step is to see that my need for help is a reason for other people to help me. Once more, the generality requirement comes into play and so I have to accept that to be rational, if my pain creates a reason for other people to help me, then the pain of others creates a reason for me to help them.

This is enough, believes Searle, to establish strong altruism. The reason for this is simple: independently of what I desire or not, I have to recognize that there exist reasons why I should help others. These reasons are rationally binding in that they exist merely as a consequence of me recognizing what rationally follows from the fact that I myself say things like 'I am in pain'. Therefore, Searle concludes, we have rationally binding desire-independent altruistic reasons for action and the existence of strong altruism has been established.

3. Why the argument fails

I want to make the case that Searle's argument fails. But I want to do so without questioning its most basic assumption: the truth of the generality requirement, both with regard to language and ethics. In my view, the problem with the argument is much more specific. The error comes in the third and final step. As we have seen, Searle says that 'I believe my need for help is a reason for you to help me'. Let us call this belief R. The generality requirement means that, on pain of contradiction, I also have to accept that your need for help would be a reason for me to help you. That too can be accepted as a simple consequence of accepting the generality requirement. But this is not enough to show that there are rationally binding desire-

[7] Ibid., 161.

https://doi.org/10.1017/S0031819102000384 Published online by Cambridge University Press

independent reasons for action. I can accept that your need for help is a reason for me to help you but still reject strong altruism. The reason for this is that Searle has glossed over an important distinction in formulating R. What we need to do is state *for whom the reason is operative*. Depending on how we interpret this, R can have two possible interpretations:

R1: I believe my need for help is a reason *for me* for you to help me.
R2: I believe my need for help is a reason *for you* for you to help me.

A simple example shows why the difference is real and important. I apply for a job, you are the person who decides who gets the job and I would be greatly helped if I got this job. This fact—that the job would help me enormously—provides *me* with a reason for you to give me the job. My need, in other words, is a reason *for me* for you to help me. But it is not a reason *for you* for you to help me. The reasons which are operative for you are ones concerning the need to give the job to the person best suited to do it well and so on.

As this is an important point, I will give a second example. I have entered a painting competition and I know that if I win, I will appreciate the prize much more than my rivals. That gives a reason *for me* for the judge to choose me as the winner. But it is not a reason *for the judge* for the judge to pick me: the only reasons operative for him are those concerning artistic merit.

In order for Searle's argument to work, R must mean R2. Only if I accept that my need for help is a reason *for you* for you to help me must I, because of the generality requirement, accept that your need for help is a reason *for me* for me to help you. If R just means R1, then all the generality requirement demands is that I believe your need for help is a reason *for you* for me to help you. That is insufficient to establish that there are rationally binding desire-independent altruistic reasons for action.

But Searle's argument is only enough to establish R1. In recognizing that I am in pain and need help, I am only rationally compelled to accept that this provides a reason *for me* for you to help me. There is nothing incoherent or irrational in believing that although I have reasons for you to help me, these reasons are not operative for you. And even if I *think* these reasons are operative for you, I might be wrong, and no false belief can provide rationally-binding reasons for action. Therefore in accepting R1 I am not rationally compelled to believe that I ought to help other people.

https://doi.org/10.1017/S0031819102000384 Published online by Cambridge University Press

4. A second flaw

I have tried to show why Searle fails to establish that there are rationally binding desire-independent altruistic reasons for action. I want now to show why his argument, even if right, cannot establish strong altruism. This will also provide a reply to a possible counter argument to my objection.

On Searle's view, accepting that one has a need requires accepting that the need is a reason for others to help me. The problem with this is that the notion of 'a reason for others to help me' is too thin. At any one time I have a large number of needs: I need to lose weight, get to the train station on time, have some lunch, take my medicine and so on. On Searle's view, all these needs must create reasons for others to help me. But this cannot mean that other people are somehow morally *obliged* to help me. If it did, we would have the absurdity that at any one time, we would all be morally obliged to help other people meet any need they have.

On the only plausible reading, therefore, 'a reason for others to help me' has to be something which does not necessarily carry with it any moral obligation. If Searle's argument is valid, it cannot mean that anyone is morally obliged in anyway to help me just because there is a reason for them to do so. In other words, the fact that there are rationally binding desire-independent altruistic reasons for action does not in anyway place an obligation on me to act on those reasons. But if Searle has not shown that there is any *obligation* to act altruistically, he cannot have established the existence of strong altruism.

So Searle's argument fails because there is a gap between any credible definition of strong altruism—which must include the idea of obligation—and what Searle claims to have established: namely that there are rationally binding desire-independent altruistic reasons for action. And if my argument of section 3 is correct, he has not established this anyway.

A possible response to my argument in section 3 is that there is no real gap between R1 and R2. It is wrong to insist that a reason *for me* for you to help me is different from a reason *for you* for you to help me. If there is a reason to help me, this reason must be capable of being recognized by any rational agent.

My argument in this section provides the response to this. The fact that there exists a reason to help someone is not sufficient to establish that someone ought to act on that reason. We need to go further and say why that reason should, in the language I have been using, be operative for any particular individual. Recognizing that

https://doi.org/10.1017/S0031819102000384 Published online by Cambridge University Press

reasons can be operative for different people and that not all reasons can be equally operative for all people is sufficient to show that Searle's argument doesn't work. This is the case whether we accept there is a real and important difference between R1 and R2 or if we prefer to say that a need creates a reason *for anyone* to help someone else, but that such a reason does not necessarily create an obligation to help that person. Either way—and these two are the only credible options—Searle fails to show there are rationally binding desire-independent reasons for strong altruism.

5. The logical gap

There is an important lesson to learned from Searle's mistake. The general line of argument he follows has a noble tradition, going back to Kant. The general strategy is to attempt to show that reasons for acting morally can be derived purely from considerations of what is required by rationality. This project can only work if there are reasons to help others or consider their interests which are operative for everyone.

The problem is that the logical gap between R1 and R2 is unbridgeable. It can never follow as a matter of logic that because there is a reason *for me* for you to help me that there is a reason *for you* for you to help me; nor that because there is a reason *for you* for me to help you that there is a reason *for me* for me to help you. Even if we reject the distinction between R1 and R2, the same problem arises: it can never follow as a matter of logic that because there is a reason for you to help me which is operative for me or places an obligation on me, that there is a reason for you to help me which is operative for you or places an obligation on you.

The same point can be extended to cover reasons for moral action more generally. Instead of R, consider the general moral principle S: 'There is a reason for you to behave morally.' Just as R admitted of two interpretations, so does S:

S1: There is a reason *for me* for you to behave morally.
S2: There is a reason *for you* for you to behave morally.

The same logical gap exists between S1 and S2 as exists between R1 and R2. That is to say, it can never follow as a matter of logical necessity from the fact that there is a reason *for me* for you to act morally that there is a reason *for you* for you to act morally; nor can it follow from the fact that there is a reason *for you* for me to act morally that there is a reason *for me* for me to act morally.

https://doi.org/10.1017/S0031819102000384 Published online by Cambridge University Press

The only way to bridge this gap is to *start* with a reason why *everyone* should act morally. A reason which is both a reason *for me* for you to act morally and a reason *for you* for you to act morally must already be a reason why *anyone* in relevantly similar circumstances should act morally. But, of course, this is just the kind of reason it has proven so difficult to find. This is why moral philosophers in the Kantian tradition have tried to arrive at a reason for anyone to act morally by starting from a reason for *some* people that they should act morally and trying to establish that there is a rational requirement to generalize from that. Such attempts impale themselves on the fork of a dilemma: either the reason is only genuinely operative for certain people, in which case, it cannot be generalized from; or it is already a reason operative for everyone, in which case the argument is not really starting from a reason which is operative for only some people at all. In other words, there is no short cut or back-door route to morality. If we want to show that altruism, for example, is demanded by rationality, we have to show that there are reasons *for me* for me to help others *and* reasons *for you* for you to help others; not just that there are reasons for me or for you to help others. Nor can it ever be enough to show that there are reasons to help others unless we can also show that these reasons are operative for everyone.

Salamanca, Spain

https://doi.org/10.1017/S0031819102000384 Published online by Cambridge University Press