

[This is a chapter from my doctoral thesis, which is a defence of the ethical theory of R.M. Hare. It centres on the main argument presented at the Oxford Moral Philosophy seminar on 5<sup>th</sup> May 2008, with some additional material. There are a few references to other chapters, but these can be safely ignored.]

## Humean and Kantian approaches to normative justification

Daniel Y. Elstein\*

### 3.0 Introduction

In this chapter I will consider the prospects for obtaining an acceptable view of moral objectivity given quasi-realism. Taking it that quasi-realism has been satisfactorily established, I turn to some difficulties which Blackburn's position faces which are not (perhaps surprisingly) consequences of his quasi-realism, but rather of his other commitments, and in so doing I hope to sketch out a distinct position that is more in line with Hare's. I will argue that Blackburn's Humean approach to normative justification is inadequate, and that a Kantian approach is superior. I will then show how to reconstruct Kant's arguments in *Groundwork* §II-III within the quasi-realist framework, and thus how his Categorical Imperative and Formula of Universal Law can be established.

There is a strategic reason for pursuing this line, in my overall project of rehabilitating Hare's version of expressivism. The central distinction between Hare and other expressivists, including Blackburn, is Hare's insistence on universalizability, not as one normative principle amongst many, but as an indispensable part of anything comprehensible as morality. Unfortunately, Hare's arguments for this position tend to be linguistic: he tries to elicit linguistic intuitions to show that moral terms had to be used in a universalizable way.<sup>1</sup> The problem with this strategy is that even if we do use moral terms in the way that Hare claims, that does not in itself provide us with a reason not to change our usage; thus the danger is that even if Hare succeeds in showing that utilitarianism can be derived from universalizability, he will not have shown that we should be utilitarians, because we will still be able to ask why we should not start thinking in non-universalizing terms. And if Hare insists that the use of the word 'should' commits us to universalizing, we can ask why we cannot use it differently.<sup>2</sup> He could reply that when we are doing moral philosophy we are in the business of answering questions of the form, 'What ought I to do?', and since this question makes use of the word 'ought', it is appropriate for us to consider the meaning of that word when answering it. But this seems only to push the worry further up the line: we could be interested in different questions, so what Hare's story leaves out is an explanation of the importance and inescapability of *that* question. Why are the questions we ask when moralising better or more important questions than those we would ask if we were smoralising instead (where by 'smoralising' we intend a practice similar to moralising but not implying universalizability)? Indeed, there is a risk that Hare's theory will end up being as implausible as the analytical descriptivism that was dismissed in Chapter 1, in that they both seem to make the ultimate normative questions into questions about meaning.<sup>3</sup> What we need is an argument that we

\* University of Leeds; D.Y.Elstein@leeds.ac.uk

1 See e.g. Hare (1952: 129-30).

2 This objection is considered by Singer (1990: 156-9) and Blackburn (1998: 227).

3 Hare (1981: 218) worries about this problem, but whilst I do not disagree with what he says in response, it does not seem to answer the objection; Hare seems more concerned about the classificatory question of whether he counts as a descriptivist than in the more pressing problem of avoiding the implausibility of descriptivism.

must engage in a form of thinking that involves the requirement of universalizability; an argument that normativity itself presupposes universalizability. And that is what this chapter aims to provide. But enough with motivating what follows; whatever my particular reasons for pursuing it, a defence of the Kantian project in metaethics is surely of independent interest.

### 3.1 Fundamental error

We begin our consideration of Blackburn with the defence of quasi-realism which he offers in his 1999 article ‘Is Objective Moral Justification Possible on a Quasi-realist Foundation?’. Here Blackburn rightly insists that the quasi-realist is not committed to relativism simply in virtue of not being a realist. It is true that the main difference between us and those we think in grave moral error (his example is the Taliban) is that we and they have different attitudes. But there is no need to hold that there is nothing to choose between our attitudes and theirs. The following passage is crucial:

And why does that [i.e. expressivism] not imply that divergent moral opinions are on all fours? Well, all I can hear that as meaning is that they are all *equally good*. And that is just not true. The Taliban’s opinion on the education of women is not as good as mine. In fact, it is diametrically wrong, wrong root and branch. And notice that this would be true even if we were less minimalist than I have been about facts. Suppose a substantive or robust theory of truth were developed, giving us some notion of correspondence. Suppose it proceeds by isolating some metaphysical category of Facts (note the upper-case). And suppose finally that for the kinds of reason I have outlined, there are no normative or ethical Facts (all these doctrines belong to the *Tractatus Logico-Philosophicus*). This would be a metaphysical result. So it clearly could not imply that all moral opinions are on all fours. It could not imply, for instance, that it is permissible to hold that women should not be educated. It could at best imply that, in holding this, you do not trespass against the upper-case Facts. But that is all right. It was not *that* (or, not simply that) that is wrong with the Taliban view. The main thing that is wrong with the view is that it is inhumane, cruel, arbitrary, and so on. The metaphysics cannot imply that it is all right to be like that! [Blackburn 1999: 217]

Blackburn rightly insists that quasi-realists are perfectly entitled to talk of the correctness and incorrectness of moral opinions, and of facts as justifying some opinions rather than others. In saying all of these things they are (on the quasi-realist view) expressing further attitudes (higher-order ones). The only reason to doubt that having such attitudes amounts to holding that those opinions are really correct or justified is the assumption that normative judgement could not be a matter of attitude, which would beg the question against quasi-realism in the worst way. So there is nothing relativist about quasi-realism. Indeed there is a hint towards the end of the passage that it is realism which is confused on this point: if realism takes the justification of a moral opinion to consist ultimately in correspondence with facts, then this risks a disconnect with the things of genuine moral importance which figure in genuine moral justification. This relates back to the criticism of realism which I made in Chapter 1: that to think of a fact as being relevant to moral justification presupposes a normative judgement concerning the justificatory power of facts of that kind.

What Blackburn establishes is the right of expressivists (or at least quasi-realist expressivists) to talk of moral justification; there is nothing in expressivism which entails that moral judgements cannot be justified. But this still falls short of an account of how moral judgements are justified. There are two powerful arguments which can be levelled against Blackburn’s position: the argument from fundamental moral error, and the normative regress

argument. The upshot of these arguments (and a common worry lies behind both) is that expressivism is compatible with objective moral justification, but only as Kantian expressivism; what gets Blackburn into trouble is not expressivism but Humeanism about practical reason. There is an immediate reply that may be tempting to Humeans like Blackburn even before hearing the arguments, which is that quasi-realism is not in the business of saying what justifies moral judgements, since that is a normative question rather than a metaethical one. For instance, opposition to torture is justified *inter alia* by the pain which torture causes, but this is no part of quasi-realism. The reason that the reply is unsatisfactory is that the objections we will consider challenge the right of quasi-realists to say anything of that form at all.

The objection from fundamental moral error is formulated by Egan (2007), and I follow his account of the challenge below. In order to explain the problem with fundamental moral error, Egan first recounts Blackburn's explanation of how quasi-realists can talk about the possibility of ordinary moral error. It initially seems difficult for quasi-realists to account for claims like "I believe that stealing is wrong, but I could be mistaken." If the first part expresses disapproval of stealing, what does the second part express (it had better not simply be the opposite attitude)? Blackburn has a response to this worry:

The problem comes with thinking of myself... that I may be mistaken. How can I make sense of my own fears of fallibility? Well, there are a number of things that I admire: for instance, information, sensitivity, maturity, imagination, coherence. I know that other people show defects in these respects, and that these defects lead to bad opinions. But can I exempt myself from the same possibility? Of course not (that would be unpardonably smug). So I can think that perhaps some of my opinions are due to defects of information, sensitivity, maturity, imagination, and coherence. If I really set out to investigate whether this is true, I stand on one part of the (Neurath) boat and inspect the others. [Blackburn 1998: 318]

The thought is that I might in future change my mind about the wrongness of stealing, and this change might be a result of applying my other norms to my moral reasoning, in which case I would count the change as an improvement. So the "I could be mistaken," expresses indecision about whether to count a putative change in my opinion as an improvement; in other words (applying an expressivist account of improvement) indecision about whether to approve of such a putative change. So, as Egan concedes, the quasi-realist has a plausible account of moral error in such cases. But isn't it also possible to worry that I am more fundamentally mistaken: that although there is no change in my moral judgement of a case that I would recognise as an improvement, this is only because I am also wrong about what counts as an improvement? The quasi-realist might be tempted to simply deny the coherence of this thought. To guard against this, however, Egan strengthens the objection by pointing out that it seems possible for moral disagreement to be fundamental: two people may hold conflicting moral opinions which are stable, in the sense that they hold background norms such that neither would count a change as an improvement.<sup>4</sup> If we hold, as a quasi-realist must, that when there is moral disagreement at least one side is wrong, we must also acknowledge that a case of stable disagreement is a case of fundamental error on at least one side. So if I am party to such a stable disagreement, it seems that it would be 'unpardonably smug' (Blackburn's phrase above) to discount the possibility that the fundamental error is on my side. And yet the quasi-realist has no way of interpreting the thought that I am

---

4 Note that in this sense of 'fundamental', the disagreement may not be over what either side thinks of as most important. There might be fundamental disagreements between two social democrats, as well as between a social democrat and a Taliban. And similarly, fundamental error might not be error that has wide-ranging practical consequences, but simply error that is stable in light of the agent's norms of improvement.

fundamentally morally mistaken, because Blackburn's proposal for dealing with moral error is that I am morally mistaken when I could recognise a change in the putatively mistaken belief as an improvement; and this is precisely what I cannot do in the case of fundamental error.

### 3.2 Normative regress

Egan takes the foregoing argument to be decisive against quasi-realism, because he cannot see an alternative quasi-realist account of moral error that copes with fundamental error (and neither can I). I want to draw a more subtle moral: Blackburn's position stands refuted, but not quasi-realism, because Kantian quasi-realists do not have to count the rejection of fundamental error in one's own case as unpardonably smug. To understand the difference between the Kantian and Humean views we need to investigate the second argument against Blackburn: the normative regress argument. This is an argument we met before in Chapter 1 as an argument against realism and for expressivism. It is not clear exactly who first came up with it; as I tried to show earlier, its roots lie with Socrates and Cudworth, though it is given more explicitly by Copp (1995: 37-49) and Cohen (2003); it is likely one of those arguments that periodically gets re-invented. I believe that Kant had something like this argument in mind when he claimed that morality depended on there being a categorical imperative. The sceptical form of the argument goes as follows: suppose that some norm N is justified; we can ask what justifies it; but whatever the answer, it will presuppose that the fact cited justifies N, which is a further normative claim (since claims about justification are normative); so every justified norm depends on some further norm for its justification; thus there is a vicious regress of normative justification, and so no norms are justified.<sup>5</sup>

The challenge for Blackburn and other Humeans is to resist this regress. There are a number of options to consider. We might reject the idea that our norms need to be justified, perhaps by equating it with the Kantian idea (which Blackburn 1999: 226 rejects) that in order to be justified in applying a norm we must be able to justify it to those against whom it is applied. But this would be a conflation, because it is not justification to others which is at issue here. The regress challenges our ability to justify norms to ourselves. Alternatively we might hold out for some kind of coherentist solution to the regress, as is suggested by Blackburn's use of the Neurath boat analogy. The thought would be that we justify each of our norms by relying on some of the others, but that it is acceptable for chains of justification to be circular, so that at least one norm depends for its justification on a norm that it is ultimately used to justify. The problem now is very similar to the earlier difficulty with fundamental error: it is possible for there to be two conflicting sets of norms, each of which is coherent (i.e. it is possible within each set to justify each of the norms by relying on the others). If there are two such conflicting but coherent sets, then only one of them can be justified, which means that coherence is not sufficient for justification. Moreover, it is then necessary to say what makes one of the sets justified and the other not. Blackburn entertains the possibility that the Taliban have such a coherent set of norms, but that our norms are better than theirs.<sup>6</sup> He implicitly

---

5 A point about terminology: where I use the term 'norm', Copp uses 'standard' and Cohen uses 'principle'. I am confident that we are all talking about the same thing, and that the difference is merely verbal.

6 It is tempting to try to avoid the problem by denying the coherence of the Taliban's norms. But nothing here hangs on whether the Taliban's norms in particular are coherent; all that is required for the argument to go through is that there be at least two coherent, but mutually inconsistent, sets of norms. And it can easily be seen that this is possible: let Set 1 consist of the two norms A (Drainpipes should be green and norm B is justified) and B (Norm A is justified), and let Set 2 consist of the two norms C (drainpipes should be red and norm D is justified) and D (Norm C is justified). Sets 1 and 2 are both coherent but that does not make it plausible that either is justified.

holds that the Taliban's norms are worse than ours because they discriminate against women. This is doubtless correct, but it cannot be an ultimate explanation of the superiority of our norms, because it presupposes that norms ought to be neutral between the sexes, which is one of our norms. It would only be sufficient as a justification if coherence were enough; but if coherence were enough the Taliban's norms would also be justified, which is impossible. It seems that Blackburn's position only makes sense if one switches back and forth between accepting and rejecting a coherentist theory of justification. (On coherentism, see also Copp's 1995: 41-2 discussion).<sup>7</sup>

One alternative worth considering for Blackburn is the one that Copp endorses. Copp (1995: 42-4) holds that there is a regress, but that it is not vicious. His view is that all the regress argument shows is that if there are any justified norms, then there are infinitely many justified norms (so that when we go up the chain of justification we never reach a final norm). It is only if we assume that there cannot be infinitely many justified norms that the argument can lead to scepticism. Copp notes that he does not mean that all infinite chains of norms (with a hierarchy of justification) are justified; that would be objectionable for the same reason that coherentism is: "But the point I am insisting on is that the fact that a standard has a place in an infinite hierarchy of standards is *not* sufficient to show it *not* to be justified" (Copp 1995: 43). It is worth noting that such a view has uncomfortable epistemic consequences: we surely do not grasp an infinite hierarchy of norms, and thus we never grasp the full justification for any of our norms. But such epistemological worries are secondary here, for there are more basic problems. Suppose that we are in possession of norms which do indeed form part of an infinite hierarchy of justified norms. There will presumably be other infinite hierarchies of unjustified norms; suppose that the norms of the Taliban form part of such an unjustified hierarchy. In virtue of what could our hierarchy be the justified one (and the Taliban one unjustified)? What reason could we have for thinking that our norms are justified? If some infinite hierarchies have a special ingredient that others lack which makes them justified, as it seems they must, then this ingredient must enter into the story at some point, and it seems needlessly obtuse and of doubtful coherence to defer this point to infinity.<sup>8</sup>

---

7 Talk of coherence here may seem to strengthen an analogy between the normative regress argument and the famous regress argument concerning epistemic justification which is used to justify foundationalism. That argument is roughly that in order to be justified in believing a proposition on the basis of some evidence one has also to be justified in believing that that evidence adequately supports that proposition; so either some beliefs are foundationally (rather than inferentially) justified, or else there is a vicious regress and no beliefs are inferentially justified. Although this analogy is strong, with the Kantian position I defend being an unorthodox kind of foundationalism, and the solutions to the infinite regress I reject being similar to epistemic coherentism and Klein's (1998) infinitism, there are two ways in which the strength of the analogy might be overestimated. The first would be to simply assimilate justification as a status of a norm with justification as a status of a belief. These are not at all the same: it is (at least conceptually) possible that a norm is justified though no-one is justified in believing that it is justified. The question of whether a norm is justified is not an epistemic question, and thus the normative regress argument is not itself an epistemic regress argument. The second mistake would be to think that the solutions to the two regress arguments must be the same, or that if the normative regress argument demands a Kantian solution the epistemic regress must do too. I do not rule out a Kantian approach to the epistemic regress, but it may be unnecessary; that is because the epistemic regress is really a restricted version of the normative regress, in that it focuses on epistemic norms, which are a subset of norms in general. Because the question of epistemic justification is in this sense local, it might be settled by reference to other norms. Another reason why the epistemic regress might be more tractable is that coherentism is more of an option, because it is not obviously wrong to say that to conflicting (coherent) sets of beliefs are both justified, whereas conflicting sets of norms cannot both be (all things considered) justified.

8 Again, it matter not at all whether there could be an infinite hierarchy including the Taliban norms; what matters is that there can be two infinite hierarchies which conflict. Let Set 1 consist of the norms A (Drainpipes should be green), B (Norm A is justified), C (Norm B is justified) and so on; and Set 2 be structurally similar but with the first norm stating that drainpipes should be green. Sets 1 and 2 are both

The regress can only be defeated if there is some way for norms to get justified which does not presuppose prior norms. Copp's solution needs this to distinguish between justified and unjustified infinite hierarchies; but once it is in play there is no need to grant the regress at all. The point against Copp can be restated slightly differently: suppose that there are two infinite hierarchies of norms, of which one is justified and the other not; now there must be a reason why the justified one is justified (something which differentiates the two hierarchies); and so there must be a further norm to explain why that reason is a reason; that norm cannot be justified simply by its place in the justified hierarchy, because the unjustified hierarchy may contain a contrary norm, according to which it is justified, and thus justificatory symmetry between the hierarchies would be preserved; and so the further norm must require further justification, and the regress continues. The upshot is that even if Copp is right that there is nothing special to prevent infinite hierarchies from being justified, they are of no use for stopping the regress.<sup>9</sup>

Radzik (2000) suggests a way of solving the regress that is vulnerable to a similar criticism: she holds that although some norms cannot justify themselves, some can, and thus stop the regress. To be clear, there are many norms which purport to justify themselves, in that they are justified according to themselves (e.g. 'This norm is justified'); call such norms 'self-supporting'. What Radzik is after is a norm that genuinely provides itself with justification, so that it really is justified in virtue of being self-supporting. But just as we ask what the difference is between the infinite justified hierarchy and the infinite unjustified hierarchy, so we must ask what differentiates the justified self-supporting norm from the unjustified self-supporting norm, and there will be no satisfactory answer to this question. Radzik could hardly say that it is just a basic fact that particular self-supporting norms are self-justifying; that would be simply dodging the normative question, rather than answering it.

There are three further possible responses to the normative regress argument which should be ruled out. One is to say that some norms are simply obviously justified, and so need no further justification. But it is unclear how this avoids the central problem avoiding all other accounts: it seems quite possible that two conflicting norms will both seem obvious, and thus seeming obviousness cannot be sufficient for justification. If the reply to this is that we are not talking of apparent obviousness but rather genuine obviousness, then the account becomes question-begging, because a norm cannot be genuinely obvious without being justified, and we are entitled to ask for an account of genuine obviousness, which will be no easier to come by than an account of justification itself.

A second account seems related: it is a kind of primitivism about justification, according to which it is simply a brute fact whether a norm is justified or not. The crucial point concerning

---

infinite hierarchies, but not plausibly justified.

9 Radzik (1999) agrees that Copp's infinitary solution to the regress does not work, but for rather different reasons, which I disagree with. She holds, against Copp, that the regress argument relies on a first-person conception of justification, according to which agents must have access to their reasons in order to be justified. She claims that an infinite chain of justified norms would flout this requirement, because agents cannot use infinite chains of reasoning when they deliberate. I agree with Copp, against Radzik, that the regress does not depend on the first-person conception. The argument is that in order for a norm N to be justified, there must be some reason why it is justified; but whatever that reason is, it seems that there must be further norm to explain why the cited reason is a justification for N. That argument concerns what reasons there are, not what reasons are accessible to an agent. My objection to Copp is thus dialectically stronger than Radzik's, because I do not rely on her dubious assumption of a first-person conception of justification, and so grant more to Copp. My version of the regress argument is effective whether or not we accept this assumption.

such a view is that it is not compatible with expressivism. The business of this chapter is to show that Blackburn's version of expressivism is inadequate, and that a different, more Kantian, version of expressivism is needed. It is only to be expected that some will take the argument as showing what is wrong with expressivism itself, and so retreat to some variety of realism. I have tried to explain in Chapter 1 why realism is so unfruitful as an answer to the explanatory question. What we have here is the attempt to use realism to answer the normative question; but note that a realist answer to the normative question cannot make sense unless realism can also answer the explanatory question.<sup>10</sup> It is no good to say that the fact that some norms are justified is simply bedrock in answering the normative question unless you also have some account of what you are talking and thinking about when saying and judging that norms are justified. So unless realism can be defended against the objections of Chapter 1 (which I shall not repeat), it is out of order to attempt to stop the normative regress by realist foot-stamping.

A final attempt is to claim that "the normative question presupposes, incorrectly, that it makes sense to morally evaluate moral values" (Curry 2005: 169).<sup>11</sup> Certainly the normative question does presuppose this, but why think it incorrect? One reason would be despair at the regress. Another would be the thought that since expressivists have given a full explanation of human moral judgement, there is no further question to ask about what really is good or right. According to Curry, a separate argument is required to show that there is anything to get right in the realm of normative judgement. One point against this which can be made immediately is that we actually do make moral evaluations of moral values, so to say that this does not is to put forward an error theory of our moral thinking, and to give a sceptical answer about at least second-order normative questions. And it seems that such scepticism is bound to seep downwards: if it makes no sense for me to think that there is anything wrong with moral values which permit torture, then it is unclear how I can think that there is anything wrong with torture itself. Part of Kant's argument, which I will explain in more detail in §3.5, is that this kind of normative scepticism is seriously unattractive, indeed untenable. Finally, the point of Kant's transcendental argument is precisely to show how there can be an objective answer to normative questions in a way that fits in with expressivism; and that is a task for the remainder of this chapter.

At this point it is worth noting a difference between Copp's version of the regress argument and Cohen's. For Cohen, the argument is meant to show that fundamental norms cannot be fact-sensitive. The structure of Cohen's regress is an alternation between norms and facts: facts justify norms, and higher-level norms justify the ability of facts to justify lower-level norms, but these higher-level norms depend on further facts etc. That is why Cohen can say that if norms can be justified in some other way then the regress ends, and normative scepticism does not result. But for Copp the regress is meant to cover all possible ways in which norms might be justified. His thought is that simply by considering whether a norm N is justified we invoke a further norm to settle the question; that further norm might make whether N is justified depend on the facts, or it might not; that is not Copp's concern. There is

---

10 The idea that there are explanatory and normative questions in metaethics is from Korsgaard (1996). The idea is that the explanatory question concerns the nature of normative judgement and its place within a causal/explanatory scientific world-view, whereas the normative question concerns the justification of norms. For more on this see §3.4.

11 Note that Blackburn does not take this line – insofar as Blackburn and Curry are both broadly Humean, there is a split in the Humean camp on whether answering explanatory questions makes it unnecessary or senseless to try to answer normative questions. As we have seen, what Blackburn objects to is not normative questions, not even normative questions about moral values, but Korsgaard's idea that there is a single normative question.

a simple way to extend Cohen's argument so that it is closer to Copp's: substitute 'reason' for 'fact'. Sometimes the reason why a norm N is justified is some fact, but the next step in the argument does not depend on that; because whatever the reason is, there is always a further question of why it is a reason for N, and that presupposes a further norm. So if Cohen's argument establishes that fundamental norms must be fact-insensitive, it also seems to show that they are reason-insensitive. But it is harder to make sense of the idea of reason-insensitive norms, and thus normative scepticism again looms large. We seem to need a reason-insensitive norm to stop the regress, and that is of dubious intelligibility.

### 3.3 A Kantian solution?

As I have already intimated, my view is that the normative regress can only be stopped in a Kantian fashion: the thought is that the idea of a reason-insensitive norm is the same as Kant's idea of a categorical imperative. Before I go further with my exegesis of Kant, it is worth briefly stating what I aim to achieve therein. The primary goal here is to argue for Kant's Formula of Universal Law (FUL); insofar as I present arguments for it which I attribute to Kant, it is more important that those arguments are successful than that they are really Kant's. I do not have space here to present the detailed textual evidence which would be necessary to demonstrate the superiority of my interpretation to that of contemporary Kantians. The main evidence that I present here in favour of my interpretation is simply the reasoning for the success of the arguments I attribute to Kant; it is better, other things being equal, to interpret Kant as giving good arguments rather than bad ones (as we shall see, leading Kant scholars admit that Kant's arguments are fallacious on their interpretations). Why then, if scholarship is not my primary goal, do I bother to say that these are Kant's arguments, rather than presenting them as my own? The reason is that these are the arguments which I find in the Groundwork; it would be dishonest to claim to have invented them myself.

With these disclaimers, I will begin by sketching the stages of Kant's main argument in the Groundwork. I take it that in Section I Kant is trying to show how his theory fits in with common moral concepts, and that the argument only really gets started in Section II (first a description of what we implicitly think, and then an explanation of why it makes sense to think that). There are then four important parts to the argument. Firstly, Kant argues that there must be a categorical imperative (4:406-417). Secondly, he argues that the content of a categorical imperative must be deducible from its form, and thus that there is a single categorical imperative which can be stated as FUL (4:417-421). Thirdly, he explains how to apply FUL to derive directly action-guiding principles (4:422-424). Fourthly (this is now Section III), he argues that the Categorical Imperative (CI) is binding on all rational beings because rational beings have to take themselves to be free (4:446-463). There is, of course, substantial material at the end of Section II, dealing with the Formula of Humanity (FH) and autonomy; much of what is said there is relevant to the rest of the argument, but these passages are largely inessential to the central thrust of the argument. The current fashion for emphasising FH over FUL seems to me to make little sense of the text, and is perhaps motivated by the mistaken view that Kant's argument for FUL is fallacious.

Of these four elements in Kant's argument, I will defend the first, second and fourth. I postpone discussion of the third until the next chapter, but I will reject it. My rejection of Kant's application of FUL is not intended to rely on the arguments or interpretation which I give in this chapter; indeed there is no reason why those who disagree with me about the argument for FUL should not agree with me about its application, and vice-versa. On the

other hand, the arguments of this chapter are not independent of the arguments given in the previous chapters for expressivism. I thus perceive a kind of independence between two parts of an overall argumentative strategy to establish an ethical theory; but it is not the traditional gap between metaethics and normative ethics, unless FUL is counted as part of metaethics. Rather the distinction is between the formal and substantive parts of ethics. The standard view that metaethics and normative ethics are independent is embraced by Blackburn, who holds that quasi-realism does not commit one to any view in normative ethics. This contrasts with the position which I argue for (and which Hare held) that universalizability is not an optional extra for expressivists; my arguments in this chapter are meant to show that expressivism is only defensible if one also accepts FUL.<sup>12</sup>

I return now to the puzzle presented to us by the normative regress argument; it seems that the regress can only be stopped if there is a norm (at least one) which is not sensitive to reasons. I now wish to point out that Kant is worried by exactly the same problem. First, he accepts what looks very much like the normative regress argument, or at least its crucial step<sup>13</sup>:

Nor could one give worse advice to morality than by wanting to derive it from examples. For, every example of it represented to me must itself first be appraised in accordance with principles of morality, as to whether it is also worthy to serve as an original example, that is, as a model; it can by no means authoritatively provide the concept of morality. [Kant G4:408]

Notice here the point that the use of an example would presuppose a prior norm. And what he says follows is that “there is, then, no genuine supreme basic principle of morality that does not have to rest on pure reason independently of experience.” (4:409) Later we are told that this is what is meant by a ‘categorical’ imperative: “if the action is represented as *in itself* good, hence as necessary in a will in itself conforming to reason, as its principle, *then it is categorical.*” (4:414) And Kant is clearly worried by normative scepticism:

On the other hand, the question of how the imperative of *morality* is possible is undoubtedly the only one needing a solution, since it is in no way hypothetical and the objectively represented necessity can therefore not be based on any presupposition, as in the case of hypothetical imperatives. Only we must never leave out of the account, here, that it cannot be made out *by means of any example*, and so empirically, whether there is any such imperative at all, but it is rather to be feared that all imperatives which seem to be categorical may yet in some hidden way be hypothetical. [Kant G4:419]

Suppose that all imperatives were hypothetical, or in other words (going by the definition of ‘categorical’ which Kant gives above) that no actions were good in themselves, what would follow? As I interpret these passages, that would mean that the question of the goodness of an

---

12 Note that it is the combination of expressivism and genuine objectivity which entails universalizability for ethics. Since humour and aesthetics, where expressivism is also correct, are not objective, a non-universalizable discourse for those areas might make sense. That is not to deny that actual aesthetic discourse involves universalizability; the point is that even if aesthetics is universalizable, shmaesthetics might not be, whereas even shmethics would have to be universalizable if it was still meant to be objective.

13 It is not entirely clear from the next two passages that Kant endorses the whole of the normative regress argument. But he is committed to the two premises of that argument: that the justification of a norm by a fact or a reason presupposes a further norm, and that fundamental norms (such as ‘the imperative of morality’) cannot be based on further normative presuppositions. Since putting these two premises together gives us an argument to the conclusion that either there is an imperative which is somehow valid without there being any reason why (i.e. a categorical imperative), or else no norms are valid, and since this is the very conclusion which Kant wants, it is more than tempting to attribute the whole argument to him.

action could never be settled by just looking at the action; there would always have to be some reason why that action was good, which would depend on it belonging to a class of actions which are good. But that class of actions would also (*ex hypothesi*) not be good in themselves: their goodness would depend on their belonging to some more general class of good actions, and so *ad infinitum*. (Why is  $\phi$ -ing good? Because  $\phi$ -ing is  $\psi$ -ing and  $\psi$ -ing is good. But why is  $\psi$ -ing good? Etc.) This is the very same normative regress as before. So it looks very much like Kant rejects the idea that all imperatives are hypothetical because he sees it as leading to a regress and thus normative scepticism, and this motivates him to give an account of how a categorical imperative is possible. Whether this is the correct interpretation of Kant is, however interesting, not the most important question for current purposes; for Kant seems to have a solution to the normative regress very different from the ones canvassed and rejected above. What we need to investigate is whether Kant can come to the rescue of the expressivists, and in particular Hare.

### 3.4 The content of the Categorical Imperative

Once we see the need for a categorical principle, two questions arise. First, what its content is (what it tells us to do), and second, whether it really exists (whether we are really bound by it, and how we are if we are). The latter task is, as I have already indicated, left until Section III of the Groundwork, but the former is to be dealt with immediately: “In this task we want first to inquire whether the mere concept of a categorical imperative may not also provide its formula containing the proposition which alone can be a categorical imperative. For, how such an absolute command is possible, even if we know its tenor, will still require special and difficult toil, which however, we postpone to the last section.” (4:420) Kant’s central argument over the content of the CI is very brief, and is not generally thought to be valid. The crucial idea is that the content of the CI must be derivable from its form alone, because there is nothing else to give it content. Here is the crucial passage:

When I think of a *hypothetical* imperative in general I do not know beforehand what it will contain; I do not know this until I am given the condition. But when I think of a *categorical* imperative I know at once what it contains. For, since the imperative contains, beyond the law, only the necessity that the maxim be in conformity with this law, while the law contains no condition to which it would be limited, nothing is left with which the maxim of action is to conform but the universality of a law as such; and this conformity alone is what the imperative properly represents as necessary. There is, therefore, only a single categorical imperative and it is this: *act only in accordance with that maxim through which you can at the same time will that it become a universal law*. [Kant G4:420-1]

The move which has seemed problematic here is from the requirement to conform your actions to universal law as such (CI), to the requirement to conform them to maxims which you can will as universal law (FUL).<sup>14</sup> Aune (1979: 29) complains that “[FUL] differs significantly from [CI] in practical import, at least, and the line of thought leading naturally to [CI] does not seem adequate to render [FUL] credible.” And Wood (1999: 81) claims that the “fallacy in Kant’s deduction” is that “it does not follow from the mere concept of a categorical imperative that *the will of a rational being* – what a rational being wills or can consistently will – has any role to play in determining the content of universal laws.”<sup>15</sup> This

<sup>14</sup> The objection can be found as far back as Schopenhauer (1965/1841: §II.7), who wonders what bearing what we will could have on the CI.

<sup>15</sup> Orthodoxy holds that Kant’s real (i.e. valid) argument for the content of the CI comes later, when he discusses autonomy; I have even heard it said that Kant does not mean to be providing a derivation of FUL from CI at all. This seems to me to be a case of wishful thinking: if you think that Kant’s argument for FUL

criticism of Kant is not always taken to be decisive; Kitcher (2004) produces a convoluted explanation of how to bridge the gap, reconstructing a detailed argument from other bits of Kant. But my view is that the move is just as simple as Kant represents it as being, with no missing steps. The key is expressivism. For expressivists, normative judgement, including moral judgement, just is a kind of willing. So when I judge that my maxim of action conforms to universal law (i.e. I judge that it conforms to a norm which is universally valid), I must will that that maxim be followed universally.

Let us expand on how expressivism helps Kant out here. For expressivists, all normative judgements are attitudes (i.e. willings). But we are here considering a special kind of normative judgement: one that is not meant to depend on anything else. This means that such judgements have to be universal, since there is nothing to restrict their scope: a non-universal judgement would be one that was restricted in some way, and would thus presuppose the normative relevance of that restriction. So to judge that a maxim is categorically (i.e. non-derivatively) correct is to will it universally. Now the final step to FUL is trivial, for we can safely assume that we are obliged to act in accordance with maxims which we judge to be correct. Of course, we may make incorrect moral judgements, but this is no objection to the deliberative principle which tells us to choose actions which conform to our best judgement. The point is that to make a normative practical judgement is already to judge that one should act in accordance with it; in willing that everyone follow a maxim, *a fortiori* one wills that one follow the maxim oneself.

Wood is mystified by the idea that what we will could help to determine the content of the CI; but it is not mysterious how the content of the CI can be connected to what correct normative judgement amounts to, and it is not mysterious (given expressivism) how normative judgement can be a matter of willing. Given the simplicity of this reconstruction of Kant's argument, an argument that is usually considered to be obviously invalid or even non-existent, it is surprising that the view that Kant is an expressivist is not more mainstream. I am afraid to say that the likeliest explanation is to be found in the misunderstanding of expressivism amongst Kant scholars. It cannot be said that the expressivist interpretation is ignored because it has not been prominently suggested: Hare repeatedly insisted that Kant was a prescriptivist, and thus in contemporary terms an expressivist.<sup>16</sup> Indeed Hare (1963: 34) hints at the very same explanation of FUL as that given above. And indeed Wood for one is not ignorant of Hare's position, though he does not acknowledge, or does not realise, that it provides a solution to the problem which he sees for Kant. Unfortunately, Wood is dismissive of Hare's expressivist interpretation of Kant:

Since Kant holds that moral truth is irreducible either to what people think or to the results of any verification procedures, he is a moral *realist* in the most agreed-upon sense that term has in contemporary metaphysics and metaethics. Kant is a moral realist because realism is the only way of preserving the *critical* stance necessary to all moral thinking, the open-endedness of moral inquiry. To say that the moral law rests on an *idea* is to say that it is always in principle possible for us to be mistaken about what we think is right, no matter who we are, how many of us there are, or what decision procedures we may have applied in arriving at our moral beliefs. [Wood 1999: 157-8]

---

is embarrassingly bad (as most Kantians do), it seems charitable to interpret him as not making that argument at all, and to ignore his talk of whether "the mere concept of a categorical imperative may not also provide its formula".

<sup>16</sup> Hare makes this point from the beginning of his writing on prescriptivism (e.g. 1952: 16), and insists on it even more later on (e.g. 1993: 12).

This makes it especially wrongheaded for R. M. Hare to associate Kantian autonomy with moral antirealism and the “is-ought gap”. [Wood 1999: 375, footnote to the previous passage]

The problem with what Wood says is that he does not appear to know much about contemporary metaethics. Expressivists do not hold that moral truth is reducible to what people think or to the results of any verification procedures, or that we are immune to moral error, or that morality is mind-dependent in any other relevant way. It is not clear that even the earliest expressivists, like Ayer, held any of these views; certainly Hare did not, and, as discussed in Chapter 1, Blackburn has demonstrated how expressivists have a right to this kind of mind-independence. Regardless of whether Wood is correct to define realism with reference to these ideas of mind-independence (I think not), it is clear that he takes expressivism to be an anti-realist view, and so implicitly holds that expressivism is committed to one or other of the sins which he contrasts to realism. So Wood relies on an understanding of expressivism which had been debunked by the time he was writing.<sup>17</sup> This means that his objections to Hare’s interpretation have no force at all.

It is worth tracing the issue between Hare and Wood a little further, to discuss the passage from Hare which Wood quotes in the footnote cited above:

The reason why heteronomous principles of morality are spurious is that from a series of indicative sentences about ‘the character of any of its objects’ no imperative sentence about what is to be done can be derived, and therefore no moral judgement can be derived from it either. [Hare 1952: 30]

Hare is not exactly trying to put forward an expressivist interpretation of Kant here. Rather, it seems to be the normative regress argument, or at least its crucial step, which is doing the work. The idea is that a heteronomous judgement is one that reads off a normative conclusion from some fact about ‘the character of any of its objects’ (Hare is referring here to Kant’s definition of ‘heteronomy’ at G4:441). But as was argued above, drawing a normative conclusion from some fact presupposes a further norm. This is why Kant says in that passage that “this relation [of the object to the will], whether it rests upon inclination or upon representations of reason, lets only hypothetical imperatives become possible: I ought to do something *because I will something else*.” Hare says (to Wood’s consternation) that Kant is here relying on the Humean doctrine that an ‘ought’ cannot be derived from an ‘is’; and this should make us reflect again on the connection between expressivism and the normative regress argument. What I want to emphasise is that we should not see the is-ought gap as a consequence of expressivism<sup>18</sup>; rather it is a datum of reflection on our normative concepts that norms are not derivable from facts alone. That is because any purported derivation would assume that the relevant facts really did justify the derived norm, and that would be a normative presupposition. As was argued in Chapter 1, expressivism is the view best able to account for this datum. And yet Wood apparently simply rejects the is-ought gap; I confess that this leaves me mystified as to how he can account for the various passages, such as the one which Hare notes, where Kant seems to rely on it.

---

17 This is one reason why I am prepared to commit myself to unorthodox readings of Kant: although the Kant scholars with whom I disagree have a vastly greater knowledge of Kant’s writings, it is unlikely that this knowledge will give them a secure grasp of Kant’s metaethical views unless they also understand what the metaethical options are. Wood’s inaccurate pronouncements on metaethics do not inspire much confidence on that score. It is not clear to me how deep these misunderstandings run within Kant scholarship, though I fear that Korsgaard is an honourable exception and that Wood is more typical.

18 As it would be if we took motivational internalism to be the basic argument for internalism, and deduced the is-ought gap subsequently, as Hume is often supposed to have done.

Before leaving the topic of Kant and expressivism, I wish to draw attention to an obscure passage from Korsgaard:

Expressivism, I believe, is like realism also true after all, and also in way that makes it boring. From the descriptive and explanatory perspective that is appropriate to scientific or perhaps in this case social-scientific inquiry, those who use normative language will appear to be simply expressing their values. When you are not in the grip of practical problems that provide standards for their own solutions, the truth and falsehood of statements employing concepts that embody those problems must be elusive. The trouble with expressivism is that it describes moral language from the outside, as if we were not ourselves the creatures who face practical problems, but only someone else making anthropological observations about them. Behind that stance is the idea that so long as we are reasoning we must remain at this anthropological level, and behind that view is the same error that animates moral realism – the view that the business of cognition is describing the world. [Korsgaard 2003: 122 n.49]

There are a number of points to note here. Korsgaard holds that expressivism is true but boring. Thus she need not (unlike most contemporary Kantians) dissent from my view that Kant is an expressivist; but she surely would disagree with my claim that expressivism is crucial to the derivation of FUL (that would hardly be boring). She also compares her attitude to expressivism to her attitude to realism: they both end up being true, without being the right places to start (2003: 118). Notice that as an attitude to realism, this is quite consistent with Blackburn's take on the matter. He too holds that moral claims end up being true or false, but that what is important is how they end up being so: we have to earn the right to talk of truth. What Korsgaard is adding here is that realism and expressivism are both partial views: each is a correct description of the way things appear from one vantage point, whereas the correct account, constructivism, is able to explain both of them. Whilst Blackburn could agree that expressivism is only a partial view (hence quasi-realism), he does seem committed to assigning a kind of priority to expressivism over realism. And this is the basis for Korsgaard's accusation that expressivists are biased in favour of the anthropological, naturalistic stance from which expressivism is true, and thus assimilate reasoning and cognition to describing the world.

As I will argue in §3.6, there is some justice to this charge. But perhaps the focus of the disagreement can best be seen by considering the two questions which Korsgaard (1996) discusses: the explanatory question, and the normative question. Whilst we are asking the explanatory question, the anthropological standpoint is bound to come to the fore. Korsgaard thinks, and I agree, that asking the normative question forces us towards a Kantian or constructivist view which operates mainly from the practical stance. The crucial difference in Blackburn's view is that he rejects the idea of a normative question as on a par with the explanatory question (see §3.5 below); for him, there are individual normative questions, but no single big question (whereas he does seem to assume a single explanatory question). This explains why the anthropological stance looms larger for Blackburn, even though at least officially he would not hold that reasoning can only go on at that level: it is the one where the question he is interested in answering arises. Whether there is any real asymmetry is a more difficult question (I discuss this further in note 24 below). The point to take from Korsgaard here is that whilst Kant is an expressivist, he is not only an expressivist: he is mostly interested in the normative question, and that is how Wood can be misled into thinking him a realist. In the end, I mean to vindicate the approach to metaethics which lies in Kant, and which Korsgaard expresses in the passage above. But, as I shall discuss in §3.7, Korsgaard's transcendental argument fails where Kant's succeeds.

### 3.5 Agency, normativity and scepticism

Having given these unorthodox interpretations both of Kant's argument for the necessity of a categorical imperative (as relying on the normative regress argument), and his argument for FUL (as relying on expressivism), we now face what appears to be an even harder challenge: explaining what Kant's argument in §III of the Groundwork is all about. The challenge is hard because the task which Kant sets himself seems impossible on the interpretations given above. What §III purports to show is how the CI is binding on us; but this sounds very much like an attempt to show that the CI is justified, which seems impossible given the normative regress argument. Recall that if we try to give a reason in support of a norm, we presuppose some other norm; but then we can give no reasons in support of the CI, since the whole point of the CI is that it does not depend on any other norm. Once we start to think in this way it may appear that the normative regress argument has led us down a blind alley: it told us that we needed a categorical norm, but it also prevents that norm from being justified, and thus seems to entail normative scepticism. This I think is a clue to what Kant is up to in §III: he is trying to show that even if we cannot justify the CI, we can show that it has a status, other than being justified, which can allow us to retain it in a pivotal role.

The key to Kant's argument is the rejection of normative scepticism. Expressivism leaves us with three candidate accounts of normativity: Kantian, Humean and sceptical. The Kantian demands an answer to the normative question, and insists that it compels us to accept a norm which is the source of justification, but not itself justified (the CI). The Humean denies that there is any such single question:

But why suppose that there is then such a thing as *the* normative question? People ask why they must do some particular thing in all kinds of circumstance, and their concerns can only be addressed in appropriately different ways. There is no more a single normative question than, for instance, a single emotional question, 'What am I to feel?'. [Blackburn 1998: 258]

The sceptic on the other hand simply concedes that no norms are justified, and thus tries to do without normative judgements (it would be Moore-paradoxical to accept a norm and also deny that that norm was justified). The normative regress argument presents us with a dilemma between the CI and normative scepticism, because it rules out the Humean position. Kantians need not assume that there is a single normative question.<sup>19</sup> Instead, we can see the normative regress argument as showing how by asking individual, particular normative questions, we can be brought to see that there is in fact one big question lying behind them all: how does normative justification get started in the first place? Given that the Humean answer is off the table, if normative scepticism can also somehow be ruled out, then we would have to accept the CI. But how could normative scepticism be ruled out without giving a justification for the CI (which would involve us in the regress again)?

The answer is meant to be provided by a transcendental argument: normative anti-scepticism,

---

<sup>19</sup> To be fair to Blackburn, in the above passage he is responding to Korsgaard (1996), and whilst it is reasonably clear that she has the normative regress argument at least partially in view (e.g. 1996: 30), she introduces the idea of a single normative question first, and does not use the regress argument to support it. The more significant difficulty with Korsgaard's approach, which is closely related to my own, is that she seems not to be fully convinced of the Kantian project, and ends up with a kind of compromise between Hume and Kant that leaves a crucial role to reflective endorsement in the fundamental justification of moral principles. But such a position is still vulnerable to the normative regress argument, which is a further reason for thinking that she does not fully grasp its importance.

and thus the CI, is a necessary condition for freedom. This is connected to Kant's emphasis of the fact that we are rational beings, and that the CI is binding for all rational beings. That does not mean that some other norm is binding on non-rational beings, and that thus the status of the CI is the same as that of the pure intuitions of space and time in the Critique of Pure Reason. Such an interpretation would be misguided, because what it is to be a rational being (at least for Kant) is to be subject to norms, to be interested in acting for good reasons. Thus no norms could apply to non-rational beings.<sup>20</sup> This then is the core of my interpretation of §III of the Groundwork: Kant holds that, *qua* rational beings, we are inescapably faced with a normative question whenever we deliberate – what do I have reason to do (or believe, in the case of theoretical deliberation)? If this is right, then what immediately follows is that we cannot accept normative scepticism.

This interpretation allows us to make sense of Kant's thought that there is a distinction between freedom and autonomy, in that freedom is possible for those who are not fully autonomous, but that those who are free (i.e. all rational beings) are still committed to autonomy. If to be free is to take oneself to act on the basis of reasons, then anyone making a practical normative judgement and acting on its basis is (negatively) free. But merely making a normative judgement without accepting the CI will end up being inconsistent (as the normative regress argument is meant to show). Since being autonomous simply involves deriving the norms on which one acts from the CI, this threat of inconsistency faced by those who are merely free can be understood as a pressure towards autonomy. When Kant argues for this transition from (in his terminology) negative to positive concepts of freedom (4:446), he claims that "freedom, although it is not a property of the will in accordance with natural laws, is not for that reason lawless but must instead be a causality in accordance with immutable laws but of a special kind; for otherwise a free will would be an absurdity." The thought is surely that in order for action to be intelligible it must be in accordance with principles, because otherwise it could not be explained, and thus that for free action norms must do the work of laws of nature. Now Kant moves immediately to saying that the law involved with freedom must be autonomy, and thus the CI (4:446-7), but we should not be concerned at this argument seeming to move too quickly, because Kant already argued in §II that the idea of normative guidance presupposes the CI.

It may seem that the alleged impossibility of normative scepticism is a mirage. After all, it is not as if a being who did not accept norms could not act; such a being could still have beliefs and desires to determine its behaviour. Some philosophers would deny this, saying that for a judgement to count as a belief it would have to be normatively constrained, so that a normative sceptic could not have beliefs (and a similar story might apply to desires as well).<sup>21</sup> But we need not take that line here, because we can allow that it is possible to have beliefs and desires without norms, whilst still maintaining that this is not possible for us, since the relevant question is what it would be like from the inside to be like that. One would perhaps be aware of one's beliefs and desires, and aware that one's body was responsive to them. But it would not be possible to see oneself as an agent, because the idea of agency presupposes deliberation, which in turn involves normative judgement, and thus the movement of one's body could not be understood as action. So the argument that normative scepticism is impossible for us is that we are essentially agents, so that non-agency is impossible for us,

---

20 I have not seen the interpretation I criticise here defended in print, but I have heard Thomas Pogge put it forward in a lecture entitled 'Two Queries for Cohen', given on 21/04/07 at the University of Reading.

21 The idea might be that it is constitutive of belief that one ought to believe that *p* if the evidence supports *p*. For somewhat different reasons Kant also thinks that normative judgement is necessary to having genuine beliefs, see e.g. G4:452.

and that normative scepticism is impossible for an agent.

I will now attempt to show where Kant propounds the view about freedom which I attribute to him above. One crucial passage is as follows:

I say now: every being that cannot act otherwise than *under the idea of freedom* is just because of that really free in a practical respect, that is, all laws that are inseparably bound up with freedom hold for him just as if his will had been validly pronounced free also in itself and in theoretical philosophy. Now I assert that to every rational being having a will we must necessarily lend the idea of freedom also, under which alone he acts. For in such a being we think of a reason that is practical, that is, has causality with respect to its objects. Now, one cannot possibly think of a reason that would consciously receive direction from any other quarter with respect to its judgements, since the subject would then attribute the determination of his judgement not to his reason but to an impulse. Reason must regard itself as the author of its principles independently of alien influences; consequently, as practical reason or as the will of a rational being it must be regarded of itself as free, that is, the will of such a being cannot be a will of his own except under the idea of freedom, and such a will must in a practical respect thus be attributed to every rational being. [Kant G4:448]

Here is my paraphrase of that passage: *Every being which takes itself to be acting freely is thereby bound by the laws which are presupposed by free action (i.e. the CI). Every agent must take itself to be acting freely, because agents take their judgements and actions to proceed from deliberation, which is to say that they take themselves to be guided by norms, and so not determined by anything else that bypasses deliberation. And to take one's deliberation to be short-circuited by external forces is what we mean by unfreedom, so to take oneself to act freely is just to take oneself to be guiding one's actions by deliberation.* What is established here is the link between agency (being rational) and freedom in the guise of apparently action-guiding normative deliberation.

### 3.6 The Unity of Reason

Given the above understanding of what Kant says about rational beings, Kant's persistent talk of reason should sound less suspicious. There is a perception amongst those with more Humean inclinations that Kant's project is to show that morality can be derived from reason, and that he sets up acting from reason in opposition to acting from desire.<sup>22</sup> With regard to the second charge, I concur with O'Neill's (1989) compatibilist interpretation of what Kant says about the two standpoints.<sup>23</sup> For Kant, acting on the basis of desires and beliefs is the same as

---

22 This line is pushed persistently by Blackburn (1998: 214-224, 243-256). According to Blackburn, "the fundamental mistake about deliberation" which Kantians make is to think of it as a way of standing back from one's desires and selecting between them. But it is a mystery to me where Blackburn finds this mistake in Kant: surely Kant's view is that when we deliberate we do not think about our desires.

23 It is not, alas, standard to say straight out that Kant is a compatibilist. Indeed Schneewind (1998) says that Kant is an incompatibilist without explanation. Part of the problem may be terminological; compatibilism is strongly associated for many with the views of Hobbes, Hume and Ayer, and it is true that Kant's version of compatibilism is different from and more subtle than theirs. This should not blind us, however, to the fact that Kant says quite clearly that although freedom and determinism ('natural necessity') seem to be contradictory they are not: the apparent contradiction is an 'illusion' (4:456). Someone who thinks that freedom and determinism are not contradictory must think that they are compatible. The modern version of compatibilism most similar to Kant's is perhaps Strawson's (1962). Strawson reverses the traditional linkage between freedom and morality: whereas it was common to worry that if it were shown that freedom was an illusion then we would have to abandon the moral appraisal of others, Strawson maintains that the reactive attitudes (attitudes of moral appraisal) are inescapable, and thus we must see ourselves and others as free. This is, I think, the same move that Kant makes when talks about the two standpoints: we have to adopt the

acting on the basis of reasons, but looked at from two different standpoints, external and internal respectively (or, as Kant puts it, empirical/phenomenal and practical/noumenal). Crucially, “the concept of a world of understanding is thus only a *standpoint* that reason sees itself constrained to take outside appearances *in order to think of itself as practical*” (G4:458).<sup>24</sup> So there is no metaphysical problem of explaining how reason itself can motivate or anything of the kind; when we are investigating how agents are motivated we are interested only in explaining their actions in terms of beliefs and desires. Kant (1996/1783) makes this point decisively:

In fact, the practical concept of freedom has nothing to do with the speculative concept, which is abandoned entirely to the metaphysicians. For I can be quite indifferent as to the origin of my state in which I am now to act; I ask only what I now have to do, and then freedom is a necessary practical presupposition and an idea under which alone I can regard commands of reason as valid. [Kant 8:14]

The first charge, however, seems to rest on an oversight on the part of Hume and Humeans. Hume’s view is that theoretical questions are subject to reason in a way that practical questions are not. Insofar as he is an expressivist, his expressivism is confined more associated with practicality than normativity. Why else would he so obviously fail to be an expressivist concerning logic and theoretical reason in general, as he does by contrasting ethics, imbued with passion, with (theoretical) reason, which is not?<sup>25</sup> But as I argued in Chapter 1, it is normativity, not just practicality, which entails expressivism, so a consistent expressivism must extend to all normative judgements.<sup>26</sup> Hume’s mistake then is either to link expressivism with the practical rather than the normative, or else to fail to notice that there are theoretical normative judgements (such as judgements concerning what is reasonable). If

---

noumenal standpoint in order to think practically, and that is all it takes for us to be free.

- 24 Here there seems to be a point of commonality between Blackburn and Kant, since they both have a use for the idea of an internal/external distinction. This is not altogether surprising, since Blackburn gets the distinction from Carnap, who is pursuing a Kantian project of sorts. (The difference between Carnap’s project and Kant’s is, roughly, that Carnap and the other positivists wanted to replace Kant’s idea of pure intuitions with that of language. Thus whilst for Kant the world of experience is structured by our pure intuitions of space and time, for Wittgenstein “*the limits of my language mean the limits of my world*” (*Tractatus* 5.6).) But there is a puzzle about the way which Kant uses the distinction, which has made it tricky to reconcile his positions in practical and theoretical philosophy: when we move between them the internal and the external seem to switch places. In the *Critique of Pure Reason*, the noumenal is the way the world is in itself, stepping outside our own ways of perceiving it, whereas the phenomenal is conditioned by our pure intuitions. In the *Groundwork*, the practical is the standpoint we take from within our deliberation, and the empirical is the view we have of actions when we step outside and consider how we fit into a causal order. So it looks as though the practical should be as inner as you can get. But Kant wants to identify the noumenal with the practical, and the phenomenal with the empirical. That is why he says that the practical standpoint is “outside appearances”, when it would seem more natural for him to talk of it being internal to deliberation. For Blackburn things are straightforward: the normative, practical perspective is the internal one. So what to make of Kant’s view? Is it really plausible that when we deliberate we are somehow thinking about how things are in themselves? And doesn’t Kant’s humility about the noumenal mean that normative knowledge would be impossible? I do not have a firm view on this vexed subject; but here is a suggestion. What if the practical standpoint does indeed involve a projection, when seen from the empirical standpoint (as Blackburn requires), but when the practical standpoint is considered from within itself, it does seem as though we are going beyond the merely empirical, so that each standpoint takes itself to be the external one? Then there is no real answer to the question of which standpoint is really external, and to call the position quasi-realism is just to adopt the empirical standpoint of contemporary naturalism.
- 25 Plausibly Hume was also an expressivist about causation and induction, but these are less obvious cases for expressivism than theoretical reason; his motivation seems to have been metaphysical scepticism rather than the thought that judgements in those areas are normative.
- 26 And thus the whole framing of the debate in terms of a contrast between reason and passion is a mistake too; Kant’s view, that reason and passion are two sides of the same coin, is much more in line with expressivism.

Hume had seen that theoretical reason is normative, and that if expressivism is true of any normative judgements then it must be the true of all normative judgements, then he would have avoided overplaying opposition between reason and passion.<sup>27</sup> Once a more general expressivism is accepted, one reason for suspicion towards Kant's project is removed: connecting morality with reason will not involve an attempt to derive an ought from an is, or a confusion between belief and desire. But perhaps there is a slightly different suspicion about Kant's project, that does not depend on an illicit limitation on the scope of expressivism: that Kant is attempting to derive practical reason from theoretical reason.<sup>28</sup> The response to this kind of thought is simply that this is not what at all Kant is doing. His view is that theoretical and practical reason are not really distinct, and that neither is more fundamental; this doctrine is known as 'The Unity of Reason'.

We know that Kant believes in the Unity of Reason, because he says so explicitly – "I require that the critique of pure practical reason, if it be carried through completely, be able at the same time to present the unity of practical with speculative reason in a common principle, since there can, in the end, be only one and the same reason, which must be distinguished merely in its application" (4:391) – and because when he talks about the principle of theoretical reason, he gives FUL:

To make use of one's own reason [publicly] means no more than to ask oneself, whenever one is supposed to assume something, whether one could find it feasible to make the ground of or the rule on which one assumes it into a universal principle for the use of reason. This test is one that everyone can apply to himself. [Kant, *What does it mean to orient oneself in thinking?*, 18]

It is no surprise that Kant believes in the Unity of Reason, because the normative regress argument and his deduction of FUL entail it. Firstly recall that the normative regress argument makes no specific reference to practical norms, so if it succeeds it shows that all norms must stem from a categorical imperative. But why must it be a single categorical imperative for both practical and theoretical reason? Recall that Kant argued that the content of a categorical imperative must be deduced from its form; and since there is only one form, there can be only one content, and thus the single CI, guiding both practical and theoretical reason. This does not mean that the CI has to govern all norms. Kant's (2002/1790) view of aesthetic norms is that they are not all they seem to be: because no justification can be given for aesthetic norms which links them back to the CI, no aesthetic norms are really justified, and thus aesthetic value is not objective, even though it purports to be. Note that this aesthetic scepticism is not impossible in the way that general normative scepticism is.

The Unity of Reason is big news if it is correct. The idea that morality and reason are on all

---

27 A note of caution here: part of the dispute between Hume and Kant seems to concern what mental faculties there are. We could plausibly interpret Kant as saying not just that theoretical and practical reason have a single principle (see below), but also that they involve a single faculty; whereas one implication of Hume's reason/passion dichotomy is that they involve separate faculties. It is not clear to me that the question of what faculties there are is a good question, and thus I abstain from taking sides. But if Hume meant to deduce separate faculties from a difference in the nature of the relevant judgements, then he was mistaken.

28 This seems to be what Blackburn (1998: 223) has in mind when he criticises Kant for trying to give us something extra to say against the sensible knave – namely the accusation of irrationality. What we want to know, however, is whether the sensible knave is a normative sceptic, or just a sceptic about practical reason. If the former, then Kant would say that the sensible knave is not possible, because agents cannot be like that. But if the latter, as seems more likely, the question is how the knave can think of his norms as justified. We can think of rationality as the aspiration we have as agents towards having justified norms; but as the normative regress argument is meant to show, this presupposes the CI.

fours runs against the grain of a great deal of philosophy. Many moral philosophers seem to think that they can take rationality for granted when doing metaethics and ethical theory. For example, Scanlon (1998: 4, 153, 195-7) claims that morally acceptable principles are those which no-one can reasonably reject. But it is hardly obvious that we will be able to settle which rejections are reasonable independently of the very moral questions at issue. Notoriously, Scanlon thinks that it is reasonable to reject a principle which demands that we always save the greater number: someone who will die if such a principle is followed has a strong reason for rejecting that principle, and this rejection is made reasonable (according to Scanlon) by the fact that there are alternative principles which impose costs on others (taken individually) which are not much greater. But this depends on the idea that it is reasonable to ignore the question of how many people alternative principles impose costs on, which is surely just as contentious as the idea that it is moral to ignore the numbers, the conclusion that Scanlon aims to establish. This is not to say that Scanlon's criterion is incorrect, just that it fails to get us anywhere. The point is that Scanlon's methodology only appears plausible because we commonly assume that being reasonable is more basic than being moral, and if Kant's argument for the Unity of Reason goes through then that assumption is wrong. Even more startling, I suspect, are the implications which the Unity of Reason has for the study of theoretical reason. Consider, for instance, the quarrel between classicists and intuitionists over the law of excluded middle; if Kant is right we should at least contemplate resolving it by applying the CI (perhaps the CI will not cut any ice, but that is not obvious). I doubt that many logicians would be sympathetic to that kind of Kantian approach to logic, yet that is what the Unity of Reason suggests.

The Unity of Reason may not, at first sight, appear very plausible. Is the suggestion really that we can somehow derive all our epistemic norms, and even the norms of logic from the CI? Before tackling this difficult question full on, it is worth explaining what the CI does not have to do concerning logic, by looking at a certain kind of apparent exception to the normative regress argument. Lewis Carroll (1995/1895) famously challenges us to say why we are forced to follow *modus ponens*. Suppose I accept *p*, and *If p then q*, do I have to accept *q* as well? Can I not ask why those two premises entail the conclusion? The answer to this question must not be to supply an extra premise to the argument: *If p, and if p then q, then q*. For if the first question was in order, we can ask a similar question again, which will require a yet further premise, and we are embarked on a regress, and one seems a lot like the normative regress. There are thus two worries: first, a bad company objection to the normative regress argument, and second the thought that the CI must somehow provide the answer to the sceptical worry. The answer to the puzzle is surely to say that in accepting *If p then q* we already commit ourselves to accepting *q* when we accept *p*; indeed the conditional must be interpreted as expressing precisely this conditional commitment.<sup>29</sup> Then once we accept the premises, there is no further question as to whether to accept the conclusion. The point here is that the normative question of why we should accept *q* is answered in part by a norm that we already accept: the conditional expresses a normative judgement. Now a normative regress worry can be raised concerning the justification of the conditional, but this is not a concern about the validity of the argument, but rather about its soundness; it is thus quite different from the kind of baseless doubt which Carroll warns us against. Since all deductive arguments can be put in *modus ponens* form, we need not see the normative regress as threatening logical validity, or call on the CI to justify it.

This is by no means all there is to say about the status of logic on the Kantian view; but we

---

<sup>29</sup> Note the connection to Blackburn's commitment semantics solution to the Frege-Geach problem discussed in Chapter 2.

can get more of a handle on the more straightforward case of epistemic norms. I do not mean to attempt any derivation of those norms here; that would be beyond the scope of the current work. But perhaps by resolving some puzzles which might be thought to beset any such derivation, we can dispel some of our doubts about whether it is possible. One worry is that by being expressivists about such norms, we undermine the fact/value distinction upon which expressivism depends. For if our beliefs depend on normative judgements for their justification, are they not thereby themselves normatively loaded?<sup>30</sup> I am out of sympathy with this question: why ever should we think that the nature of a mental state is determined by what it depends on, rather than its subject matter? However much beliefs depend on norms, they are not about norms. So there is no immediate threat that Kantian quasi-realism collapses on itself. A more serious worry, however, is that we have no place to start when working from the CI. It seems hopeless to try to derive anything from the CI without relying on auxiliary beliefs as well; but if we need epistemic norms to justify beliefs in the first place, then how can the derivation ever get going? The reply we would like to make is that we start off with beliefs anyway, before we ever get into the business of making normative judgements, and that thus we can think of our practice with regard to our beliefs as more along the lines of the Neurath boat analogy that Blackburn wants to use in the case of norms.<sup>31</sup> But this observation may make us have second thoughts about the whole Kantian enterprise: if we have to rely on some kind of reflective equilibrium or coherentist approach anyway, even if we buy the argument for the CI, then Blackburn's suggestion that we do so from the start looks more attractive. Why should we be allowed to treat norms and beliefs so differently?

Of course, there have also been attempts to provide transcendental arguments in the case of beliefs too, to show that some beliefs are genuinely indispensable.<sup>32</sup> But even without considering whether any such arguments can succeed, there is a way of seeing an asymmetry between beliefs and normative judgements that can resolve our worries. Although expressivists hold that normative judgements are desires, this does not entail that all desires are normative judgements.<sup>33</sup> There is no such gap for beliefs. So when the normative question arises (which happens whenever a particular normative question first arises) we start off with beliefs and desires, but no normative judgements. It is impossible to start without any beliefs, so it is pointless to try. But we do seem to be able to step outside of our normative judgements, precisely because of the availability of the transcendental argument exhibited here, and this need not involve (impossibly) doing without any desires at all. Thus, whether or not there are transcendental arguments available for fundamental beliefs, the project of

---

30 This point is pressed by Putnam (2002).

31 Note that the traditional regress for epistemic justification is not the same as the normative regress, and is more easily answered. Since epistemic norms need not be the most fundamental norms, we can give reasons for coherentism or (non-transcendental) foundationalism within epistemology. But as the normative regress argument covers all norms, there is no room for this kind of justification from the outside.

32 As was already suggested, it seems that such arguments must be limited in comparison to the kind of transcendental argument I have been considering. The argument of the Groundwork (as I interpret it) is meant to establish the CI for all agents, and so there is no possibility of any other fundamental norm. But the transcendentially justified beliefs will presumably be restricted to beings like us, rather than any beings capable of having beliefs.

33 What is the difference then? I would suggest that normative judgements must be universalizable, partly because of the arguments in this chapter, but there is no need to assume this here. Blackburn (1998: 9) sees it as a matter of emotional ascent, which means that when we have pro-attitude we also have a pro-attitude to the presence of the first attitude, and so on. The problem with this is that it seems to make no sense of two-level moral theories (like Hare's – see Chapter 5) where one can think that it is better not to have the correct attitude. That is, one judges (at the critical level) that maximising utility is best, but one still thinks that the best attitudes to have (at the intuitive level) are not utilitarian attitudes.

starting from the CI in epistemic practice (which is an implication of the Unity of Reason) is not obviously incoherent. This is a rather tentative resting point. But remember that the normative regress argument seems to point inexorably to the CI and the Unity of Reason. So whatever doubts I feel about the derivability of epistemic norms, I am convinced that Kant's story about normativity is the best that we have.

### 3.7 Transcendental arguments

The argument which I attribute to Kant in §III of the Groundwork, and which I defend here, is a transcendental argument; thus it is important to consider how it faces up to general objections which beset such arguments, and how it compares with other transcendental arguments for normative principles. What I will attempt to show is that the transcendental argument under consideration is able to meet the objections to such arguments, whereas alternatives, notably that of Korsgaard (1996) are not.

Perhaps the most influential objection against transcendental arguments is that levelled by Stroud (1968). Some statements are such that, although they are contingent, they must be true whenever they are uttered by someone, e.g. the statement "I am here"; Stroud calls the class of such statements the "privileged class".<sup>34</sup> Stroud suggests that the idea of transcendental arguments is to defeat scepticism by showing that the statement which the sceptic doubts is a member of the privileged class. But if this is the strategy of transcendental arguments, there seems to be a problem:

In particular, for any candidate S, proposed as a member of the privileged class, the skeptic can always very plausibly insist that it is enough to make language possible if we *believe* that S is true, or if it looks for all the world as if it is, but that S needn't actually be true. Our having this belief would enable us to give sense to what we say, but some additional justification would still have to be given for our claim to *know* that S is true. [Stroud 1968: 255]

Whether or not Stroud is right that the sceptic can insist that all that has been shown is that we must believe S, he is clearly correct that there is a gap between showing that we have to believe S and showing that it is true. And it seems that this criticism will be applicable to the argument of this chapter, because what I have tried to show is precisely that we must accept the CI, and this does not seem in itself to dispel doubts about its truth or validity. Yet Skidmore (2002) has attempted to show that Stroud's objection is ineffective against transcendental arguments for principles of reason (i.e. norms). Skidmore's idea is that for norms, as opposed to claims about the world, there is no gap between our having to accept them and their being valid:

The fact that we must believe external objects exist does not entail that they *do*, and this leaves room for Stroud's skeptic. However, in the case at hand regarding NC [a principle of non-contradiction], this distinction dissolves. To demonstrate that anyone capable of belief must accept NC *just is* to demonstrate that NC is a valid principle of theoretical reasoning. [Skidmore 2002: 127]

This is a tempting line to take, but I fear that it cannot be used to defeat scepticism concerning fundamental norms. Or rather, we need a more complicated set of reflections to

---

<sup>34</sup> These are epistemic or primary necessities, in the terminology of two-dimensional modal semantics. See Chalmers (2005).

get us to Skidmore's position.<sup>35</sup> The problem is that the principle P, that if a norm N is such that the sceptic must accept it then N is valid, seems itself to be a normative principle: it is just a way of saying that the fact that the sceptic must accept a norm justifies it. It is thus in order for the sceptic to doubt P too. (This is just another instance of the normative regress that concerned us earlier.) Nor does it seem fair to stipulate that the validity of norms consists in the necessity of their acceptance; that seems no better than the Platonist move of stipulating that validity consists in correspondence with the Form of the Good. Or at least this is how things look from a fully internal perspective. Why could we not simply be mistaken even about a norm which we must accept, given that the fact that we must accept it is just that: a fact which can only have normative relevance if a further norm is presupposed? From the internal perspective normative facts are just as robust as facts about the external world, and there seems just as much room for error, even necessary error. To get the closure of the gap between appearance and reality which Skidmore wants, we have to move external: when we see normative judgements as attitudes, we must see necessity (of the kind possessed by the CI) as the limit of anything that could be thought of as validity for them. The justification possessed by the CI is thus of a radically different kind from that of the norms it justifies (or helps to justify): it has an external justification, not an internal one.

Although there is much in this chapter which is in agreement with Korsgaard (1996), it is important to note the difference between the transcendental argument for morality which she offers, and the one I find in Kant and present here. Korsgaard's argument also takes normative scepticism to be impossible, in roughly the same way that I do, but rather than taking this to be a (more or less) direct grounding for the CI, she takes it to license us to value our humanity, *qua* practical identity. The problem with this is that a separate argument is now required against egoism: Korsgaard claims that because private reasons are ruled out (by a Wittgensteinian argument), all reasons must be agent-neutral. But as Skidmore (2002) argues, this seems to equivocate between different senses in which a reason might be said to be private: a reason that is comprehensible to others is in one sense (Wittgenstein's) not private, but there is another sense in which a reason is private unless others have the same reason. What Korsgaard argues for is that reasons must be comprehensible to others, but what she needs as a conclusion is that they must be shared by others, and it is simply unclear how the two can be convincingly linked. But that is not to say that I disagree with Korsgaard's conclusion. Fundamental reasons must indeed be agent-neutral (i.e. shared by all agents), but the reason for that is the normative regress argument: to think that the facts which distinguish between people have normative relevance is to presuppose some norm concerning those facts, and ultimately such fundamental norms cannot depend on identifying facts. This is really to put in slightly different terms the argument of this chapter. Since Korsgaard comes so close to stating the normative regress argument, and connecting it to Kant's thinking, it is slightly puzzling that she does not employ it more directly to justify the rejection of fundamental agent-relative (i.e. non-agent-neutral) reasons.

I conclude by returning to Egan's argument from fundamental error. Recall that the great difficulty for Blackburn was that he seemed committed to immunity to fundamental error, but that on his account a claim to such immunity would be unpardonably smug. What the Kantian

---

<sup>35</sup> Skidmore (2002: 127) argues that it is the limited scope of the norms we are trying to justify which "removes the gap between appearance and reality on which Stroud's skeptic depends". I confess that I do not understand how this is meant to help; and anyway, it will not help *me*, because the principle I wish to defend (the CI) is not limited in scope. Indeed for Kant the position with respect to limitation of scope is opposite to that suggested by Skidmore: transcendental arguments for external objects are limited in scope because they apply only to beings like us, whereas transcendental arguments for norms are unlimited in scope because they apply to all beings for whom normative questions arise. See also §3.5 above.

considerations above are meant to have shown is that reliance on the CI is not unpardonably smug. Indeed it may help us to get clear on exactly what the transcendental argument is meant to show if we think of it as a way of persuading ourselves that it is not unpardonably smug to take ourselves to be immune from error concerning the CI.<sup>36</sup> This is how it transpires that the Kantian project and the quasi-realist project are not opposed, but rather inseparable. Hare's Kantian expressivism, with its deep commitment to universalizability, stands vindicated.

---

36 Note that we do not have to take ourselves to be completely immune from error, even about the CI. It may be that the argument of this chapter is flawed; but it does not seem right to see the judgement that the argument is correct as just another attitude to be traded off with other attitudes as part of Neurath's boat. For the argument takes a sufficiently Archimedean stance towards the boat that we can reasonably think of it as standing outside. What it tells us is that the CI is indispensable to any boat at all; without the CI as glue, any boat of norms will collapse.

## References

- Aune, B. (1979) *Kant's Theory of Morals* (Princeton University Press).
- Blackburn, S. (1998) *Ruling Passions* (Oxford University Press).
- Blackburn, S. (1999) 'Is Objective Moral Justification Possible on a Quasi-realist Foundation?', *Inquiry* 42: 213-28.
- Carroll, L. (1995/1895) 'What the Tortoise Said to Achilles', *Mind* 104(416): 691-3.
- Chalmers, D. J. (2005) 'The foundations of two-dimensional semantics', in M. Garcia-Carpintero & J. Macia, eds., *Two-Dimensional Semantics: Foundations and Applications* (Oxford University Press). Also available online at: <http://consc.net/papers/foundations.html>
- Cohen, G. A. (2003) 'Facts and Principles', *Philosophy and Public Affairs* 31: 211-45.
- Copp, D. (1995) *Morality, Normativity, and Society* (Oxford University Press).
- Curry, O. (2005) *Morality as natural history*. PhD thesis, London: LSE Research Online. Available at: <http://eprints.lse.ac.uk/archive/00000441>
- Egan, A. (2007) 'Quasi-Realism and Fundamental Moral Error', *Australasian Journal of Philosophy* 85: 205-19.
- Hare, R. M. (1952) *The Language of Morals* (Oxford University Press).
- Hare, R. M. (1963) *Freedom and Reason* (Oxford University Press).
- Hare, R. M. (1993) 'Objective Prescriptions' in E. Villanueva, ed., *Naturalism and Normativity* (Atascadero CA: Ridgeview), reprinted in R. M. Hare, *Objective Prescriptions* (Oxford University Press, 1999), references to reprint.
- Kant, I. (1996/1783) *Review of Schulz's Attempt at an introduction to a doctrine of morals for all human beings regardless of different religions*, in M. J. Gregor, ed. & trans., *Practical Philosophy*, the Cambridge Edition of the Works of Immanuel Kant (Cambridge University Press).
- Kant, I. (1996/1785) *Groundwork of The metaphysics of morals* in M. J. Gregor, ed. & trans., *Practical Philosophy*, the Cambridge Edition of the Works of Immanuel Kant (Cambridge University Press).
- Kant, I. (2001/1786) *What does it mean to orient oneself in thinking?* in A. W. Wood, ed. & trans., *Religion and Rational Theology*, the Cambridge Edition of the Works of Immanuel Kant (Cambridge University Press).
- Kant, I. (2002/1790) *Critique of the Power of Judgement*, ed. & trans. A. W. Wood, trans. E. Matthews, the Cambridge Edition of the Works of Immanuel Kant (Cambridge University Press).
- Kitcher, P. (2004) 'Kant's Argument for the Categorical Imperative', *Noûs* 38(4): 555-584.
- Klein, P. (1998) 'Foundationalism and the Infinite Regress of Reasons,' *Philosophy and Phenomenological Research* LVIII: 919-26.
- Korsgaard, C. (1996) *Sources of Normativity* (Cambridge University Press).
- Korsgaard C. (2003) 'Realism and Constructivism in Twentieth-Century Moral Philosophy', in *Philosophy in America at the Turn of the Century*, supplementary volume (APA Centennial Supplement) of the *Journal of Philosophical Research*.
- O'Neill, O. (1989) *Construction of Reason* (Cambridge University Press).
- Powell, B. (2006) 'Kant and Kantians on "the Normative Question"', *Ethical Theory and Moral Practice* 9(5): 535-44.
- Putnam, H. (2002) 'The Entanglement of Fact and Value' in his *The Collapse of the Fact/Value Dichotomy and other essays* (Cambridge MA: Harvard University Press).
- Radzik, L. (1999) 'A Normative Regress Problem', *American Philosophical Quarterly* 36(1): 35-47.
- Radzik, L. (2000) 'Justification and the Authority of Norms', *Journal of Value Inquiry* 34(4):

451-61.

- Scanlon, T. (1998) *What We Owe to Each Other* (Cambridge MA: Harvard University Press).
- Schneewind, J. B. (1998) *The Invention of Autonomy* (Cambridge University Press).
- Schopenhauer, A. (1965/1841) *On the Basis of Morality*, trans. E. F. J. Payne (Indianapolis: Bobbs-Merrill).
- Singer, P. (1990) 'Reasoning towards Utilitarianism' in D. Seanor and N. Fotion, eds., *Hare and Critics* (Oxford University Press).
- Skidmore, J. (2002) 'Skepticism about Practical Reason: Transcendental Arguments and Their Limits', *Philosophical Studies* 109: 121-41.
- Strawson, P. F. (1962) 'Freedom and Resentment', *Proceedings of the British Academy* 48: 1-25.
- Stroud, B. (1968) 'Transcendental Arguments', *Journal of Philosophy* 65: 241-56.
- Wittgenstein, L. (1922) *Tractatus Logico-Philosophicus*, trans. C. K. Ogden (London: Routledge and Kegan Paul).
- Wood, A. W. (1999) *Kant's Ethical Thought* (Cambridge University Press).