

REALISM AND CONSTRUCTIVISM IN TWENTIETH-CENTURY MORAL PHILOSOPHY

CHRISTINE M. KORSGAARD

HARVARD UNIVERSITY

ABSTRACT: In this paper I trace the development of one of the central debates of late twentieth-century moral philosophy—the debate between realism and what Rawls called “constructivism.” Realism, I argue, is a reactive position that arises in response to almost every attempt to give a substantive explanation of morality. It results from the realist’s belief that such explanations inevitably reduce moral phenomena to natural phenomena. I trace this belief, and the essence of realism, to a view about the nature of concepts—that it is the function of all concepts to describe reality. Constructivism may be understood as the alternative view that a normative concept refers schematically to the solution to a practical problem. A constructivist account of a concept, unlike a traditional analysis, is an attempt to work out the solution to that problem. I explain how the philosophies of Kant and Rawls can be understood on this model.

I. THE ORIGINS OF MORAL REALISM

T

he story I’m going to tell is a story about the moral and political philosophy of the twentieth century, but it begins a little earlier. I date it from 1706, the year that Samuel Clarke published his *Discourse Concerning the Unchangeable Obligations of Natural Religion*.¹ This work contains the first clear statement of the position we have come to know as moral realism. I start there, because the vicissitudes of moral realism in the twentieth century, and the need to put it behind us in the twenty-first, is what I want to talk about today.

By moral realism, I do not mean the view that propositions employing moral concepts may have truth values. That, as I will argue in due course, is a point on which realists and constructivists can agree.² Moral realism, rather, is a view about *why* propositions employing moral concepts may have truth values. Since articulating this view will be a large part of the work of this essay, I will only try to gesture at it in a rough and ready way here. Moral realism, then, is the view that propositions employing moral concepts may have truth values because moral concepts describe or refer to normative entities or facts that exist independently of those concepts themselves.³ We have the concepts in order to describe or refer to those facts. Seen this way, moral realism is a view about why we have moral concepts, one that I will argue is mistaken.

Moral realism was first articulated as a reaction against another theory, the theory of Thomas Hobbes.⁴ Hobbes had argued that morality is the solution to a problem. To see what the problem is, we need only consider what human life would be like if there were no morality; and to consider that, we need only reflect on the facts of human psychology. According to Hobbes, human beings are driven by an insatiable appetite for power, which is in turn fueled by a kind of bottomless insecurity. Whatever resources we may obtain in order to satisfy our appetites and protect ourselves from others, it turns out that we need more resources still in order to protect the first ones. Hobbes envisioned a “state of nature,” a sort of pre-moral condition in which the consequences of our psychology are allowed to work themselves out freely. In a state of nature, each person would try to dominate and if possible to enslave everyone else around him, because this would be the only security for his position in a world where everyone else would be doing the same thing. This would lead to that miserable condition that Hobbes called a “war of all against all,” in which, as he famously tells us, the life of a human being would be “solitary, poor, nasty, brutish, and short.”⁵

In order to alleviate this condition, Hobbes thought, human beings would be motivated to make a contract with one another, in which we give up our power and freedom to a sovereign who would be able to enforce laws that would make life tolerable. When the sovereign is placed in authority and the laws are enforced, not only the political state, but also moral obligation, comes into existence. Reason itself tells us *which* laws we must follow in order to bring about conditions of peace and security among human beings. But Hobbes thought that until the sovereign can enforce these laws and make sure that nearly everyone obeys them, no one can be *obligated* to obey them. To obey moral laws all by yourself would just make you the victim of your more ruthless neighbors. In the state of nature, Hobbes therefore said, the notions of right and wrong have no place.⁶

One might take Hobbes to be pointing out just how deeply rooted in human nature morality must be. But Samuel Clarke took Hobbes to be attacking what Clarke called “the real difference” between good and evil.⁷ Given the

relations between things, Clarke tells us, certain actions are by their very nature fit or unfit to be done. Given that God is superior to us, for instance, it is fit that we should worship Him. And given the relations between people, it is fit that we should promote the good and be faithful to our contracts, that we should treat one another justly, that we should rescue the endangered, and so on.⁸ All these things are fit “in themselves.” How do we know this? Clarke tells us that:

These things are so notoriously plain and self-evident, that nothing but the extremist stupidity of mind, corruption of manners, or perverseness of spirit, can possibly make any man entertain the least doubt of them.⁹

In fact, Clarke continues:

it might . . . seem altogether a needless undertaking, to attempt to prove and establish the eternal difference of good and evil; had there not appeared certain men, as Mr. Hobbes and some . . . others, who have presumed . . . to assert . . . that there is no such real difference originally, necessarily, and absolutely in the nature of things.¹⁰

And so begins the most frequently recurring theory of the modern period. In the eighteenth century, moral realist positions were defended by John Balguy against the moral sense theory of Francis Hutcheson,¹¹ by Richard Price against the sentimentalist views of both Hume and Hutcheson,¹² and by Thomas Reid against Hume.¹³ In the nineteenth century realism was defended by William Whewell against the utilitarianism of Paley and Bentham.¹⁴ In the twentieth century realism was first defended by G. E. Moore against the utilitarianism of Mill and Sidgwick as well as against what Moore called the “metaphysical ethics” of Kant.¹⁵ Moore took up a realist position about the Good but not, at least as he was understood by his successors, about the Right, so his admirer W. D. Ross then took up the realist cudgels against Moore himself.¹⁶ H. A. Prichard, with even more temerity than the rest of this crowd, leveled volleys of realist ammunition against Plato and Aristotle.¹⁷ Later in the century, naturalistic realists like Peter Railton and David Brink took on the emotivists and other so-called “non-cognitivists,” while Thomas Nagel took on John Mackie and Gilbert Harman.¹⁸ And in a book that was already in progress as the century turned, Derek Parfit is defending a realist position against, among other people, Christine Korsgaard.¹⁹

Perhaps it would be best for the honor of our subject if we could conclude that such an army of defenders is a sign that moral realism is the correct position. But I want you to notice all of the “againsts” in the list I just presented, because I think they are no accident. Moral realism seems to owe its perennial reappearance to its reactive character rather than to its truth. It reappears in the wake of almost every attempt to give a substantive account of our moral nature. This might make us suspect that many realists share Samuel Clarke’s view that it would not be necessary to defend ethical reality, if only

philosophers would stop trying to explain it. But what, more precisely, is the realist reacting against?

Clarke thought that Hobbes's theory could not really explain obligation. For Hobbes says that obligation springs from the social contract, and if that is right, how are we to explain the obligation to be faithful to the social contract itself? *That* obligation cannot come *from* the contract. So according to Clarke, Hobbes faces a dilemma. On the one hand, Hobbes could admit that being faithful to the social contract is fit and reasonable in itself. But in that case he may as well admit that other morally required actions are fit and reasonable in themselves too, and therefore do not depend for their obligatory character on the social contract. On the other hand, Hobbes could insist that being faithful to the contract, and the other things the sovereign compels us to do, are not fit and reasonable in themselves. In that case what we call "moral obligation" is really just the exercise of arbitrary power, on the part of either God or a sovereign.

Clarke's eighteenth-century successor, Richard Price, found the exact same difficulty in moral sense theory. Hutcheson had argued that human beings are equipped with a "sense" that gives rise to responses of approval and disapproval. We have a natural tendency to take a particular kind of pleasure in the contemplation of benevolent action, and as a result of that particular pleasure we deem benevolent action to be virtuous. The moral sense does not discern a moral quality that is already there; rather, it confers a moral quality, a virtuousness, on the action it approves of, just as the Hobbesian sovereign confers an obligatory character on the actions required by his laws. So the moral sense is after all only a sort of internalized sovereign, and can no more obligate us, according to Price, than a real external sovereign can. The moral sense, like a sovereign, can bully us, but nothing outside of the character of the action itself can make it obligatory. If an action is not intrinsically *right*, Price insists, then it cannot be obligatory. Right and obligatory *mean* the same thing, and since rightness cannot coherently be conferred on an action from outside, neither can obligatoriness. Price thinks this shows "that virtue, as such, has a real obligatory power antecedently to all positive laws, and independently of all will."²⁰ And this has a further implication. Price says:

As to the schemes which found morality on self-love, on positive laws and compacts, or the divine will; they must either mean, that moral good and evil are only other words for *advantageous* and *disadvantageous*, *willed* and *forbidden*. Or they relate to a very different question; that is, not to the question, what is the nature and true account of virtue; but, what is the subject-matter of it.²¹

Things should be starting to sound familiar now, for the argument I have just quoted from Richard Price's 1758 *Review of the Principal Questions and Difficulties in Morals* is exactly the same argument that we find in G. E. Moore's *Principia Ethica* a century and a half later. Non-normative words, Price insists, cannot just be "other words" for normative ones. It may turn out that

what is advantageous, or what is willed by God or commanded by a sovereign, gives us the content or “subject matter” of ethics, but, Price tells us, it cannot be what “virtuous” or “obligatory” *means*.

II. REALISM IN THE TWENTIETH CENTURY

The predominant ethical theory when Moore wrote was of course utilitarianism. This was Moore’s target, and he took aim at it by arguing that “good” does not *mean* “pleasant.” In fact, he argued, any attempt to define “good” falls afoul of the (misnamed) “naturalistic fallacy”—the fallacy of believing that “good” can be defined in any naturalistic or indeed any non-normative terms.²² When we say that something is good we are saying something normative—we are implying that it ought to be brought about or pursued. But Moore argued that for any natural or metaphysical quality put forward as the essential characteristic of the good, it is an open question—an intelligible question, worth asking—whether things with that characteristic really ought to be brought about or pursued. “Good” is therefore indefinable, Moore argued, and, as a corollary to that, value must be understood as an intrinsic rather than a relational property. That is, because the good cannot be defined relationally as that which someone desires or that which someone enjoys or as that which someone wills, it must be intrinsic to its object. Of course it might be *true* that the good is pleasure, or the desirable, or what someone wills. But—and now Moore might have put his point in Price’s exact words—that is a question of the subject matter of ethics, not a question of what Price called the “nature and true account” of the good. Although utilitarianism and hedonism were the most famous targets of his argument, Moore himself thought the argument just as effective against other theories. Egoism, evolutionary ethics, and the so-called “metaphysical ethics” of Kant are all equally guilty of the supposed “fallacy.”

In fact Moore’s argument was unoriginal. Not only earlier realists like Richard Price, but nearly everyone in the modern period who considered the question of definitions agreed that normative words could only be defined in terms of other normative words.²³ Indeed moral sense theorists such as Hutcheson considered this to be one of the strongest arguments in their favor. If normative terms cannot be defined then they must denote simple ideas, and according to the Lockean psychology these empiricist philosophers accepted, simple ideas can only come from sense, not from reason. Richard Price appreciated the force of this argument, and saw that to meet it, he had to undermine the empiricist assumption that underlies it. Price’s book therefore opens with an attempt to show that reason *can* be the source of simple ideas. It was Price, in this connection, who brought into prominence the fateful idea that we know ethical truths by “rational intuition.”²⁴ But Moore’s attack on the definability of goodness had more impact than Price’s because of the way it

intersected with broader trends in philosophy that were destined to change the shape of the subject.²⁵

The period starting in the late nineteen-thirties and -forties was the age of what we might call “high” analytic philosophy. One of the things it brought with it was a slight change in the emphasis of empiricism. The empiricists of the eighteenth century saw their debate with the rationalists as a debate about the intellectual sources of our concepts or “ideas”—whether we get our ideas from reason or from sense experience and sentiment. In the early twentieth century empiricism shifted from a view about the sources of our concepts to a view about their contents, about how they are to be analyzed. It is an interesting question how those two ideas are related—whether a concept can come from experience without being analyzable in terms of experience. The twentieth century emotivists who claimed Hutcheson and Hume for their philosophical ancestors assumed, without much argument, that it follows from the view that moral concepts come from sentiment that moral judgments may be analyzed as the mere expression of sentiments. As I have just pointed out, Hutcheson at least would have rejected this idea, and the non-cognitivism that goes with it, since Hutcheson did not think that moral concepts and judgments can be analyzed at all. But as the case of Richard Price shows, the realists already had some tendency to interpret any account of why and how human beings use moral concepts as a reductivist account of what those concepts mean. The shift in the emphasis of empiricism cemented that tendency.

And of course it had more alarming implications still. According to the “verificationist” theory of meaning popular among logical empiricists, a concept’s content is given by the way its application would be empirically verified, by the experiences we would use to tell whether the concept applies or not. This view raised important questions about moral language, for it looks as if the applicability of moral concepts, of good and bad and right and wrong, cannot be empirically verified. Under the influence of verificationism, many early twentieth century philosophers came to doubt whether moral concepts had what they called “cognitive content” at all. This is one philosophical expression of the famous Fact/Value distinction, and it sparked a debate about what the function of moral language is, if it is not to report facts about the world. Various “non-cognitivist” proposals about the nature of moral language were explored. Prescriptivists held that moral language is essentially prescriptive or imperative. Emotivists held that moral language is used to express our approval and disapproval of actions, and that moral judgments are no more true or false than cheering or booing are true or false. By the nineteen fifties this view was the prevalent moral theory in the United States.

All of this sharpened up the debate between the realists and everybody else. Now the alternatives appeared to be either having or lacking cognitive content, where having cognitive content meant being descriptive of some possible object of experience. If one filled in the content of moral concepts

with some item that is a possible object of experience—say, pleasure—then one ran afoul of the naturalistic fallacy. If one filled it in with some non-natural object of intuition, like Moore's intrinsic values, then one ran afoul of the scientific conception of the world. This seems to leave ethics in a real bind, and it resulted in extensive debates both about verificationism, and about whether the naturalistic fallacy is indeed a fallacy. I'm not interested in rehearsing those debates here. What is important for my purposes is this: Even when what we might call the verificationist element in verificationism was dropped—that is, even when philosophers reclaimed the intelligibility of propositions that cannot be verified through the empirical sciences—one element in the verificationist picture was retained. That element is the idea that it is function of all of our concepts, or anyway all of our authentically cognitive concepts, to describe reality.

We must go carefully here. In calling this into question, I do not mean to deny that there is a sense in which all of our concepts—that is, all of the concepts we have any business using—can be used in propositions which do in fact describe reality, in the sense that they are capable of being true or false. Rather, I mean to call into question the idea that this is what all of our concepts are *for*—that their cognitive job, so to speak, is to describe reality. So long as we retain that idea, it will continue to appear that moral realism is the only possible alternative to relativism, skepticism, subjectivism, and all of the various ways that ethics might seem hopeless. And so long as moral realism appears to be the only alternative to these skeptical options, the need to show that moral truth is as solid, as real, as objective, as scientific truth—and also that it is objective *in the same way* as scientific truth—will seem pressing. This was our situation in the early and middle years of the twentieth century.

I am aware that what I have said about the function of our concepts will seem vague until I articulate an alternative function. I intend to do that presently, but first I want to notice that from the start there was already a problem with the distinction between cognitivism and non-cognitivism in ethics. The distinction suggests that a moral judgment either articulates a description of some fact or is a disguised version of some alternative use of language—either expressive or prescriptive. But where does this leave theories like Aristotle's and Kant's, according to which moral judgments are the conclusions of practical reasoning? A conclusion of practical reasoning is not obviously a description of a fact about the world, but it hardly seems like some sort of emotional expletive, either. Where do these theories fit?

I believe that the answer is that they *don't* fit, but unfortunately this may not be obvious, because the issue raises a further question about the principles of practical reason themselves. We may say that it is true that an action is right just in case it accords with the categorical imperative, for instance, but what then are we to say about the categorical imperative itself? When we ask about the status of the principles of practical reason, the question of cognitivism and

non-cognitivism seems to come up once more. For instance one may suppose that the principles of practical reason must be self-evident truths known by intuition, and then Kant will come out looking like a traditional rationalistic realist. This is not just a fantasy—this is how Kant was actually read by many late nineteenth and twentieth century philosophers, especially in Britain. Sidgwick and Mill, who read Kant this way, will serve as sufficiently distinguished examples.

Alternatively, someone who takes seriously Kant's thesis that moral laws are the laws of autonomy, legislated by the agent's own will, may read him, as Hare sometimes seems to do, as a prescriptivist and so a non-cognitivist. So although practical reason theories might at first seem to fall between the cracks, there are ways of making them fit the mold.

I will come back to the question how practical reason theories, or at least Kant's theory, should be understood. What I want to do now is articulate a contrast between the theory of normative concepts that I believe stands behind the debate between the cognitivists and the non-cognitivists, and another theory of normative concepts which I take to be a genuine alternative. I call this alternative, in deference to prevailing usage, constructivism. In order to articulate the contrast, I am going to compare arguments from two of the giants of twentieth century moral philosophy, Bernard Williams and John Rawls. Both of these arguments concern a favorite twentieth century theme—the implications of the diversity of ethical opinion. It is going to be necessary for me to spell these arguments out in some detail. But what I am asking you to be interested in for the purposes of this essay is not their particular success, but rather the conception of moral concepts and along with it of moral philosophy that lies behind them.

III. CASE STUDY ONE: BERNARD WILLIAMS²⁶

Williams is certainly not a moral realist, so his position here as its spokesman may occasion some surprise. I choose him for two reasons: first because I think his attempt to articulate the idea behind realism is unrivalled in its clarity, and second because of a realist assumption that I think in the end still haunts his own account of moral objectivity. In *Ethics and the Limits of Philosophy*, Williams argues that there is a contrast between the kind of objectivity we can hope to find in science and that which we can hope to achieve in ethics. He frames this contrast in terms of convergence, that is, in terms of what might lead us to the best kind of agreement about the judgments in question. In science, the ideal form of convergence would be this: we come to agree with one another in our scientific judgments because we are all converging on a description of the way the world really is. In ethics, Williams thinks, this sort of convergence is unavailable, and so another must be found.²⁷

Williams's account of realism emerges when he explains what he means by "the way the world really is." One thing we might mean in talking about

“the way the world really is” is whether we have applied our concepts correctly. If we say the sky is blue on a day when it is blue then we have done that. But we can also query our conceptual scheme itself. We can ask whether it is the correct one or the best one or something along those lines. Since science leads us to modify our conceptual scheme, and we think of these modifications as improvements, it seems as if some such question must be in order. Williams supposes that the improved scheme is improved in the sense that it comes closer to describing the way the world really is.

Williams proposes that we can capture the notion of “the way the world really is” by formulating a kind of limiting conception which he calls “the ‘absolute conception’ of the world.”²⁸ The idea involves a contrast between concepts that are more and less dependent on the particular perspective from which we view the world. For instance, we use color concepts because we see in color, so color concepts are dependent on our own particular perspective. The concept of a certain wave length of light is supposed to be less dependent on our perspective.

Williams associates two other properties with a concept’s greater independence from particular perspectives. First, our use of concepts which are more dependent on our own perspectives will be both explained, and justified, in terms of a theory that employs concepts which are less dependent. So for instance our use of color concepts might be explained by a theory of vision which employs wave length concepts. Relatedly this theory will also *justify* our belief that color vision is a form of *perception*, that is, a way of learning about the world, by the way that it explains it.²⁹ Color vision is a way of learning about the world because it gives us information about wave lengths, or something yet more ultimate, which we take to be part of reality. Second, the more independent of our own perspective a concept is, the more likely it is that it could be shared by other rational investigators who were unlike us in their ways of learning about the world. Suppose that there are rational creatures on some other planet who cannot see colors but can hear them. They could not use color concepts but they might be able to use wave-length concepts. The more independent concepts are more shareable.

Williams thinks that the nearest thing we have to a conception of “the way the world really is” is the conception of the world that is maximally independent of our own perspective. And he thinks that if we and the alien investigators began to converge on such a conception (and of course to agree on what judgments are correct within it) then we would have reason to believe we were converging on the absolute conception. This would be the best case of convergence for science. Our theories would come to converge with the theories of other investigators because all of us were converging on a description of the way the world really is.

Now consider what the parallel would be in ethics. Here Williams first deals with a possible objection, namely that there is nothing analogous to

perceptual judgments in ethics. Seeing the facts is one thing, and evaluating them in a certain way is another. To counter this objection, Williams appeals to the existence of what he calls “thick” as opposed to “thin” ethical concepts. Thin ethical concepts—like right and good and ought—do not appear to be world-guided, in the sense that their application does not appear to be guided by the facts. (Notice the echo of verificationism here.) But thick ethical concepts—Williams’s examples are coward, lie, brutality, and gratitude—are world-guided and action-guiding at the same time.³⁰ Only an action that is in a particular way motivated by fear can be called cowardly, and yet to call an action cowardly is to suggest that it ought not to be done.³¹

Of course a prescriptivist or emotivist or a latter-day expressivist has his own account of these concepts. He thinks that their world-guidedness is one thing and that their action-guidingness is another. To say that an action is cowardly is simply to denigrate an action motivated by fear. The difficulty with this analysis, according to Williams, is that it suggests that it would be possible to use a thick ethical concept with perfect accuracy even if you were completely incapable of appreciating the value that it embodies. Williams thinks this is implausible. He does not mean that we can only use an evaluative concept when we ourselves actually endorse the value in question. But he thinks we can apply such concepts only by entering imaginatively into the world of those who have the relevant values, not merely by applying a set of factual criteria.³² We have to see the world through their eyes. This makes it natural to think of judgments employing thick ethical concepts as a kind of perceptual judgments, for they are a sort of lens through which we view the world. And that in turn makes it natural to think that, like other perceptual judgments, they may be a kind of knowledge.

I say that the sky is blue, and my visitor from Mars says that it makes a humming noise. Are we agreeing? Certainly we don’t *mean* the same thing, since I am talking about how the sky looks and he is talking about how it sounds. Yet when we reflect on these views we find that the things we both say have implications expressible in terms of a more absolute concept, that of wave-lengths. And when we look at those implications our judgments are found to converge. Here we find grounds for confidence that both of our perceptions are guiding us rightly: they are ways of knowing about the world.

To get the parallel in ethics, we would have to compare judgments made in alternative sets of thick ethical concepts. One person says that lying is sinful and another says it is dishonorable. They do not appear to *mean* exactly the same thing. But we might take both of their remarks to have implications describable in terms of what *we* think is a more absolute concept—that lying is *wrong*—and here we find that they converge. This would be evidence that their moral perceptions are guiding them rightly, and are ways of knowing about the moral part of the world. On the other hand, suppose what is in question is avenging an insult. The first person thinks it is sinful to avenge it,

while the second thinks it is dishonorable not to. Does this imply any disagreement, in particular a disagreement about whether avenging an insult is, in our sense, morally wrong? Williams does not think so, and this is partly because there is a world-guided side to the idea of something's being dishonorable. The second person may indeed be dishonored—he may be personally diminished, in his own eyes and those of his community—if he does not stand up for himself. Williams thinks that facts of this kind should make us doubt whether the two people are using concepts that converge on what is morally wrong, and so should also make us doubt whether “morally wrong” is a more absolute concept after all—that is, whether it describes some reality.

Instead Williams proposes a different way in which we might think about differences in ethical beliefs. He suggests that we should see the values of a culture not as their best approximations of the truth about right and wrong, but rather as a kind of *habitation*—although, as he emphasizes, and this is a point I will come back to—*not one that they have built*. Their values form a part of the structure of the social world in which they live. This does not mean that we cannot make any evaluative judgments about a culture's values. We can ask whether the social world that is made of those values is a good place for human beings to live. A theory of human nature, drawing on the resources of the social as well as the physical sciences, could guide our reflections about what conduces to human flourishing. Psychoanalytic theory, for instance, could guide our views about whether a social world structured by certain values was mentally healthy or not. Williams proposes that if we did find that a social world promoted the best life or at least a flourishing life for human beings, this would justify the values embodied in that social world.

I have two reasons for placing this argument before you. First, I want you to consider Williams's very clear articulation of the idea of realism, and of what moral realism would be if we had it. The values of different cultures would represent their attempts to discern the moral part of reality, perhaps through the lens of some more perspectival concept. Second, I am interested in what the kind of objectivity Williams does suppose ethics might have in common with scientific objectivity as he conceives it. This is the one piece of realism that I mentioned earlier as still haunting Williams's account, and it comes out clearly when Williams draws his conclusion. He argues that only one ethical belief might be objectively true in the ordinary, scientific, sense, namely, the belief that a certain sort of social world was the best one for human beings to live in. And then he says:

other ethical beliefs would be true only in the oblique sense that they were the beliefs that would help us to find our way around in a social world which . . . was shown to be the best social world for human beings.³³

Williams thinks that scientific beliefs are objective in the sense that they approximate as closely as possible to a representation of “the way the world

really is.” Such beliefs help us find our way around in the natural world. So when Williams goes looking for some remnant of objectivity in ethics, he looks for the world that ethical beliefs would help us to find our way around in. Theoretical beliefs constitute a kind of map of the natural world, and ethical beliefs constitute a kind of map of the social world.³⁴ To be knowledge, a belief must help you find your way around in some world, he seems to think, and to do that, it must represent that world. It must describe some kind of reality.

IV. AGAINST THE MODEL OF APPLICATION

I think that this view of the relations between science and ethics represents a rather deep confusion about the difference between knowledge and action—something almost amounting to a failure to tell them apart. One way to articulate this admittedly difficult thought is by the metaphor I have just used. If to have knowledge is to have a map of the world, then to be able to act well is to be able to decide where to go and to follow the map in going there. The ability to act is something like the ability to *use* the map, and that ability cannot be given by *another map*. (Nor can it be given by having little normative flags added to the map of nature which mark out certain spots or certain routes as good. You still have to know how to use the map before the little normative flags can be of any use to you.) To put the same point another way, goodness in action cannot just be a matter of applying our knowledge of the good—not even a matter of applying our knowledge of what makes action itself good. This is because the ability to apply knowledge *presupposes* the ability to act.

Let me try to make this last point in a less metaphorical way. Suppose we agree that ethics is about what makes action good. For now I intend that phrase—“what makes action good”—to be neutral between consequentialist, virtue-theoretic, and deontological accounts of what makes an action good. Whatever their differences, the proponents of all of these accounts have to agree that one thing that makes an action good is conformity to the principle of instrumental reason, or what Kant calls the hypothetical imperative. An action that does not succeed in achieving its end is a failure, after all, but a movement that does not even succeed in aiming at an end doesn’t succeed in being an action at all.

Now a realist account of the normativity of the instrumental principle is incoherent. For think how that account would have to work. The agent would have to recognize it, as some sort of eternal normative verity, that it is good to take the means to his ends. How is this verity supposed to motivate him? The obvious way to understand how facts motivate us is by means of a kind of extension of the instrumental principle itself. Philosophers have long acknowledged that the instrumental principle naturally extends to what is sometimes called “constitutive reasoning”—to use Williams’s example, your end is an entertaining evening, and you choose dinner and a movie as what will constitute that end. The same line of thought extends the instrumental principle

even further to cover the role of *judgment* in action quite generally. Your end is to have a cup of coffee and you choose *this* cup of coffee as what constitutes your end. But this cannot be how we employ the instrumental principle itself. We would have to say that an agent's end is to do a good action, and that he sees that an action conforming to the instrumental principle is good. Taking the means to his ends is therefore *itself* a means to his end of doing a good action, and he chooses to conform to the instrumental principle under the influence of—the instrumental principle itself. The point is that the instrumental principle cannot be a normative truth that we *apply* in practice, because it—or its natural extension to cover the case of judgment—is essentially the principle of application itself, that is, it is the principle in accordance with which we are operating when we apply truths in practice.

But the realist picture in fact works no better for moral principles than it does for the principle of instrumental reason. For even if we know what makes an action good, so long as that is just a piece of knowledge, that knowledge has to be applied in action by way of another sort of norm of action, something like an obligation to do those actions which we know to be good. And there is no way to derive such an obligation from a piece of knowledge that a certain action is good. A utilitarian thinks an action is good because it maximizes good consequences and a virtue theorist thinks it is good because it is kind or brave. But how is it suppose to follow that it is to be done?

Now a classical utilitarian (or for that matter, an intuitionist) might argue that I am ignoring the obvious. An agent is obligated to perform an action when there is a rule specifying that actions of that kind are to be performed. The classical utilitarian, for example, thought there was such a rule, the principle of utility. It is because of the existence of this rule that we characterize actions as obligatory or forbidden. But the trouble with that strategy is that it leaves us with two problems, which in the end come to the same thing. First, it does not tell us why there is such a rule. Nor, if this is a different question, does it tell us why we should conform to the rule. We seem to be caught on the horns of a dilemma when we confront this question: are we obligated to obey the rule? If one is obligated to obey the rule, then the notion of obligation must be prior to the existence of the rule. We cannot explain obligation in terms of the rule, as something that arises from it. On the other hand, if we are *not* obligated to obey the rule, then it seems we may permissibly ignore it, and so we have not after all explained why the actions it directs are obligatory. This is really a variant of the same problem that Clarke found in Hobbes, but I am now claiming that it also holds for any view that makes the goodness of action an item of knowledge. The knowledge itself, being something external to the will, is just a kind of sovereign.

The problem rests in thinking of the rules that define the obligatory and the forbidden as standards we *apply* when we are deliberating about what to do. Standards of goodness for things other than action do work this way. Having

decided, say, that you are going to buy a car, you then ask yourself what makes for a good one, or perhaps what makes for a good one in your own case. But these standards exert their normativity, if they do so at all, through action itself. Possession of a good car is the object, and if one is successful the product, of the action. But the performance of a good action cannot intelligibly be the object and product of action itself. And it's a related fact that in the case of things other than action, you don't absolutely have to apply the standards of goodness. Usually you have the same reason for choosing a good X (or a good X for you) that you have for choosing an X. But it is at least imaginable that you might just *pick* an X without reference to its goodness, like someone who is asked to pick a number from one to ten or someone randomly picking a cookie from a passing tray. Evaluative standards, taken by themselves, do not obligate. If we think of rules of action as something we may or may not *apply* when we deliberate about what to do, then either we are obligated to apply them or we are not. In neither case, as I have argued, can obligation be derived from the existence of a rule of action. So it is not just the principle of instrumental reason that gives rise to the problem.

Nor can we solve the problem by omitting the talk of rules and just saying that the goodness of action is directly normative for the agent. So long as the fact that the action is good is supposed to be a piece of knowledge that the agent applies when she acts, the same problem will arise. The argument I am making now is, in a way, the ultimate extension of the open question argument. If it is just a fact that a certain action would be good, a fact that you might or might not apply in deliberation, then it seems to be an open question whether you *should* apply it. The model of applied knowledge does not correctly capture the relation between the normative standards to which action is subject and the deliberative process. And moral realism conceives ethics on the model of applied knowledge.

V. CASE STUDY TWO: JOHN RAWLS

But if moral philosophy is not the search for some sort of knowledge that we might apply in action, what else can it be? What other model is available to us? We can find an alternative model in the work of John Rawls. Rawls, following Kant, treats the problems of political philosophy as problems that are practical all the way down.

Rawls, like Hobbes before him, thinks that justice is the solution to a problem.³⁵ Political philosophers have long been aware that there is a kind of paradox at the very heart of liberalism. The paradox emerges most starkly if you imagine someone trying to argue in favor of instituting liberal policies in a nation whose culture and beliefs are not liberal. Anyone who wanted to argue that liberal policies should be instituted in such a society would face an intractable problem, for it is an essential tenet of liberalism that political policies should be acceptable in the eyes of the people who are governed by them. If liberalism is the doctrine that

you can't push people around in the name of what you think is right, then liberals themselves are committed to the view that *they* can't push people around in the name of the doctrine that you can't push people around in the name of what you think is right. To put the point more simply, we cannot tyrannize over others in the name of liberalism and still be consistent liberals.

Now Rawls is not, in this sense, trying to justify the liberal state. That is, he is not trying to give arguments that would show that there are grounds for forcing a liberal regime on an illiberal population. But he is concerned about a parallel problem that arises when we try to justify policies *within* the liberal state, for even within the liberal state, we must use the coercive mechanisms of the state to enforce liberal policies. Since liberalism claims that political policies are justified only when they are acceptable in the eyes of the citizens, we must be able to offer reasons in support of these coercive policies that are acceptable to all the citizens. Suppose, for instance, a society contains a majority and a minority religion, and suppose the majority wishes to get their church accepted as the official state church. What reasons can they offer for this? They might say, "Our religion should be the state religion, because our religion is the one true faith." This is not a reason that the minority can reasonably be expected to accept. So in a liberal society, this will not be accepted as a good reason for coercive action. It does not meet the standard of "public reason."

Of course the majority are not just going to *insist* that theirs is the one true faith. They have certain arguments for their view—metaphysical, theological, and historical arguments—and they can marshal these and try to convince the minority that they are right. But the minority also has these kinds of arguments, and we may suppose that when all is said and done the disagreement remains. In a modern society, people may hold different philosophical, theological, and metaphysical doctrines, and with them different conceptions of a good life, which are all "reasonable" in the sense that they need not involve any obvious error or craziness. In the face of this fact it is inconsistent with liberalism to justify political policy on metaphysical or theological grounds, since such grounds may not be acceptable to all.

But of course it follows from this that we cannot justify liberal policies *themselves* on controversial metaphysical or philosophical grounds. For instance, in explaining to the majority why we must honor freedom of religion, we cannot consistently appeal to Kant's theory that autonomy is the supreme moral value or to Mill's arguments for the consequentialist value of open discussion and experimental living. These may be excellent reasons for believing in freedom of religion; but they do not meet the criterion of being acceptable in the eyes of everyone. But this seems to leave us at a loss. So there is the problem: how are we to give reasons that everyone can accept, in a society where people derive their reasons from radically different conceptions of the good?³⁶

Now before we go on, I want you to consider the parallel between this problem and the one that confronts Kant in the opening argument of the third section

of the *Groundwork*, when he is explaining the foundation of the categorical imperative.³⁷ Kant begins by defining a free will as a causality that is effective without being determined by any alien cause. Anything outside of the will counts as an alien cause, including the desires and inclinations of the person. The free will must be entirely self-determining. Yet, because the will is a cause, it must act according to some law or other: a lawless cause, Kant thinks, is a kind of contradiction.³⁸ Alternatively, we may say that since the will is practical reason, it cannot be conceived as acting and choosing for no reason. Since reasons are derived from principles, the free will must have a principle. But because the will is free, no law or principle can be imposed on it from outside. Kant concludes that the will must be autonomous: that is, it must have its *own* law or principle. But now we have a problem: for where is this principle to come from? If it is imposed on the will from outside then the will is not free. So the will must adopt a principle for itself. But until the will has a principle, there is nothing from which it can derive a reason. So how can it have any reason for adopting one principle rather than another? And indeed the problem is in a way even worse than that. For it looks as if the free will, by imposing some principle upon itself, must restrict its own freedom in some arbitrary way.

These two problems, Kant's and Rawls's, have the same structure. In both cases what we are looking for is principles themselves, for we need reasons, ways of choosing and justifying our actions or our policies, and reasons are derived from principles. Yet the very structure of the situation seems to forbid us to choose any *particular* principles. The liberal's need to avoid compromising the freedom of the citizens by forcing a particular conception of the good on them parallels the free will's need to avoid adopting a principle that will compromise its own freedom. In each case, it looks as if the choice of any particular principle will represent an act of arbitrary power. In Rawls's construction of his problem, it looks as if the choice of any particular principle of justice must be based on an arbitrary preference for one conception of the good over others. In Kant's construction of his problem, it looks as if the choice of any principle for the will must involve an arbitrary restriction of the will's freedom. And the solutions proposed by Kant and Rawls take a parallel form.

Kant's solution goes like this: The categorical imperative, as represented by the Formula of Universal Law, tells us to act only on a maxim that we could will to be a law. And *this*, according to Kant, *just is* the law of a free will. To see why, we need only compare the *problem* faced by the free will with the *content* of the categorical imperative. The problem faced by the free will is this: the free will must have a law, but because the will is free, it must be its own law. And nothing determines what that law must be. *All that it has to be is a law*. Now consider the content of the categorical imperative, as represented by the Formula of Universal Law. The categorical imperative merely tells us to choose a law. Its only constraint on our choice is that it have the form of a law. And nothing determines what that law must be. *All that it has to be is a law*. Kant

concludes that the categorical imperative *just is* the law of a free will. It does not impose any external constraint on the free will's activities, but simply arises from the nature of the will. It describes what a free will must do in order to be a free will. It must choose a maxim that it can regard as a law.³⁹

Rawls's solution to his problem can be put in parallel terms. Rawls's two principles of justice tell us that all citizens must have equal basic liberties, and that our society must otherwise be designed so that everyone has as large a share of primary goods as possible, with which to pursue his or her own conception of the good.⁴⁰ And *these*, Rawls might say, *just are* the principles of justice for a liberal society. To see why, we need only compare the *problem* faced by a liberal society with the *content* of Rawls's two principles of justice.⁴¹ Echoing Rousseau, we might say that the problem faced in the original position is this: to find a conception of justice which enables every member of society to pursue his or her conception of the good as effectively as possible, while leaving each member as free as he or she was before.⁴² The content of Rawls's two principles simply reflect this conception of the problem. So Rawls's two principles simply describe what a liberal society must do in order to *be* a liberal society, just as Kant's principle describes what a free will must do in order to *be* a free will. Rawls's principles are derived from the idea of liberalism itself, in the same way Kant's categorical imperative is derived from the idea of free volition.

In Kant's argument we arrive at the categorical imperative by thinking about the problem faced by a free will, just as in Rawls's we arrive at the principles of justice for a liberal society by thinking about the problem faced by a liberal society. In each case, in fact, a sufficiently detailed and accurate description of the problem actually yields the solution. And you should notice one implication of this—the categorical imperative is not a rule one *applies* in deliberation. We arrive at the categorical imperative by thinking about *how* a free will must deliberate: it must do so by choosing a law for itself. The categorical imperative is a principle of the *logic* of practical deliberation, a principle that is constitutive of deliberation, not a theoretical premise that is applied in practical thought. And one might also say that Rawls's principles are a development of the *logic* of liberalism.

VI. CONSTRUCTIVISM

Kant and Rawls are constructivists: I will now try to articulate what this means. Practical philosophy, as conceived by Kant and Rawls, is not a matter of finding knowledge to apply in practice. It is it is rather the use of reason to solve practical problems. The concepts of moral and political philosophy are the names of those problems, or more precisely of their solutions. This is made clear by the way Rawls employs the concept/conception distinction in *A Theory of Justice*.⁴³ There, the *concept* of justice refers to the solution to a problem. The problem is what we might call the distribution problem: people

join together in a cooperative scheme because it will be better for all of them, but they must decide how its benefits and burdens are to be distributed. A *conception* of justice is a principle that is proposed as a solution to the distribution problem, arrived at by reflecting on the nature of the problem itself. The concept *refers to whatever solves the problem*, the conception proposes a particular solution.⁴⁴ The normative force of the conception is established in this way. If you recognize the problem to be real, to be yours, to be one you have to solve, and the solution to be the only or the best one, then the solution is binding upon you.

The same structure is clear in Kant's argument. There the problem is the one set by the fact of free agency. It's nothing less than the problem of what is to be done. And Kant thinks that just by reflecting on the nature of that problem, we can arrive at the categorical imperative. The move from a negative to a positive conception of freedom in Kant's argument parallels the movement from concept to conception in Rawls's. Negative freedom is the name of a problem: what shall I do, when nothing determines my actions? Positive freedom proposes a solution: act on a maxim you can will as a universal law.

So according to constructivism, normative concepts are not (in the first instance—a caveat I will explain below) the names of objects or of facts or of the components of facts that we encounter in the world. They are the names of the solutions of problems, problems to which we give names to mark them out as objects for practical thought. The role of the concept of the right, say, is to guide action; the role of the concept of the good might be to guide our choice among options, or of ends. The “thinness” of these concepts, to use Williams's language, comes from the fact that they are, so far, only concepts, names for whatever it is that solves the problems in question. We need *conceptions* of the right and the good before we know how to apply the concepts. The task of practical philosophy is to move from concepts to conceptions, by constructing an account of the problem reflected in the concept that will point the way to a conception that solves the problem. To produce a constructivist account of the right or the good is to ask: is there some feature of the problem itself, or of the function named by the concept, that will show us the way to its solution? The feature of the problem of liberal justice that shows us the way to its solution, according to Rawls, is that a liberal society must respect the freedom of its members while enabling them to pursue their conceptions of the good. The feature of the problem of free action that shows us the way to its solution, according to Kant, is that free action must be determined by the agent herself.

It may seem at first glance as if on a constructivist account there is something very special about the moral or normative domain. All of our other concepts, one might think, name things in the world, or perhaps things which track things in the world, while ethical concepts gesture at practical problems and their solutions.⁴⁵ But even leaving aside the question whether this is a correct account of our scientific concepts, the idea that some of our everyday

concepts refer to the solutions to problems is perfectly familiar. To see this, consider artifact concepts—consider, for example, the concept of “chair.” Why do we have the concept of chair? Certainly not because it would form part of the absolute conception of the world, for those alien investigators with whom we are to share that conception might, for all we know, be oval creatures who swim through their environment like fish, and never sit down. We have the concept of “chair” because the physical construction of human beings makes it possible, and occasionally necessary, to sit down. So the concept of chair is functional; a chair is something that plays a certain necessary role in human life, and the conception of chair is filled out by asking what sorts of things can properly fill that role. The person who first came up with a conception of chair probably was also the first who constructed an object—a chair—to fit that concept, and so to solve the human problem that it represents. To this extent constructivism makes moral concepts like the concepts of artifacts. This doesn’t make them arbitrary or relative, for there are kinds of artifacts—“chair” is an example—that all human beings in all human cultures have some version of, and that have to have certain features given the problems that they solve. Williams, you may recall, said that the values of a culture are “part of their way of living, a cultural artifact they have come to inhabit (though they have not consciously built it).”⁴⁶ But why shouldn’t we consciously build our social world, or our political societies, or for that matter our practical identities? This, according to a constructivist, is what practical philosophy is all about.

And of course none of this means that moral language doesn’t admit of truth and falsehood, any more than it means that artifact language doesn’t admit of truth and falsehood. For the correct conception of a concept will be a guide to its correct application, and when a concept is applied correctly, what we get is truth. But what makes the conception correct will be that it solves the problem, not that it describes some piece of external reality. Rather, as the term “constructivism” suggests, our use of the concept when guided by the correct conception *constructs* an essentially human reality—the just society, the Kingdom of Ends—that solves the problem from which the concept springs. The truths that result describe that constructed reality.

It is important to note one feature of constructivist theories. Theories will vary in how extensively constructivist they are, because different normative “objects” are constructed in different constructivist theories. Rawls constructs the principles of justice. In doing so he takes certain other normative notions for granted. The parties in Rawls’s “original position” choose with a view to what will be *best* for them, but under constraints of information, the veil of ignorance. The notion of what is “best” for someone is also a moral or normative notion, and it, or rather the notion of “good” upon which it is based, is not constructed, according to Rawls.⁴⁷ In T. M. Scanlon’s theory, to take another example, moral principles are constructed; the problem they solve is one of justifiability. Scanlon, in constructing moral principles, asks which

principles people might reasonably reject. The notion of reason is also a normative notion, and Scanlon does not think that it is a constructed notion.⁴⁸ So there is a question about how “deep” constructivism can go. Can even our own most basic reasons themselves be constructed? Kant’s view, as I understand him, and as I have sketched his view in this paper, is that they can. To put it in my own terms, when an agent determines whether she can will a maxim as universal law, she is determining that she can endorse a certain consideration in favor of doing something and therefore can treat it as a reason. This constructed normative “object” is then available for use in other constructions. If this sort of Kantian argument doesn’t work, then constructivism cannot go “all the way down.” I of course think that it can.

VII. CONCLUSION

And now I will conclude with a bit of a twist. In this paper I’ve contrasted two theories of normative concepts and their function. I’ve identified, as the basic assumption behind the realist model, a general view of our concepts—the view that the function of a concept is to describe a piece of reality. I’ve suggested in opposition to this that some of our concepts—including justice, right, and good—are essentially names that refer to the solutions to problems. Now the objection inevitably arises: shouldn’t we be asking which of these two theories of our normative concepts *is true*? But we don’t have to ask that question, for considered in one way, constructivism and realism are perfectly compatible. If constructivism is true, then normative concepts may after all be taken to refer to certain complex facts about the solutions to practical problems faced by self-conscious rational beings. Of course it is only viewed from the perspective of those who actually *face* those problems in question that these truths will appear normative. Viewed from outside of that perspective, those who utter these truths will appear to be simply expressing their values.⁴⁹ Realism and expressivism are both true in their way. But establishing that realism is true in *that* sense is not the *end* of moral philosophy, in either sense of “end”: it is only the *beginning*.

The important difference between realism and constructivism rests then not in which one is true, but in how they orient us towards the tasks of practical philosophy. The moral realist thinks of practical philosophy as an essentially theoretical subject. Its business is to find, or anyway to argue that we can find, some sort of ethical knowledge that we can apply in action. According to constructivism, the only piece of knowledge that could be relevant here is knowledge that the problems represented by our normative terms are solvable, and the only way we can find out whether that is so is by trying to solve them. So for the constructivist practical philosophy is a practical subject. Its business is to work out solutions to practical problems.

For much of the twentieth century, just as for the three centuries or so that preceded it, philosophers remained in thrall to the view that the function of

all human concepts, and perhaps of all conceptual inquiry, is to describe the world. This, in my view, amounts almost to thinking that the function of human life is to describe the world, and if that's right, it is clear enough what is wrong with realism. Rawls's work in one way broke new ground, but in another way simply followed the lead of Hobbes and Kant, in thinking that philosophy can be practical, or to put the same point another way, that practice itself can be reflective. And after all, it is no accident that we still read Hobbes and Kant, while we have forgotten Clarke and Price and Whewell, just as the philosophers of the future will still be reading Rawls when they have forgotten G. E. Moore. For it is in this practical conception of moral and political philosophy that both our significant historical achievements and our hopes for making progress in the future can be found.⁵⁰

ENDNOTES

1. The Boyle Lectures, 1705, first published in London in 1706. Selections are available in D. D. Raphael, *British Moralists 1650–1800* (Indianapolis: Hackett Publishing Company, 1991). I will quote from this edition.
2. The ancestor of what I say here is my discussion of distinction between procedural and substantive realism in *The Sources of Normativity* (Cambridge: Cambridge University Press, 1996), §§1.4.4–1.4.9.
3. We may understand “normative entities or facts” capaciously here. Some contemporary realists who ground moral judgment in reasons like to say that they are not committed to the existence of anything like Platonic forms: they are only committed to the existence of irreducibly normative facts about what is a reason for what. I include “facts” to accommodate them. Indeed the view I am about to discuss in the text, Samuel Clarke's view, also seems to be about facts rather than entities: facts about what makes what fit. Once facts are included, however, naturalistic realists may then wish to add that the facts in question may supervene on natural facts, say about pleasure or desire, that can be used to explain our moral views and practices in a systematic way. I mean to include this too. What is essential to a realist position is the view that propositions employing moral concepts have truth value because they track certain independent facts which explain our use of moral concepts.
4. In *Leviathan*, 1651.
5. Thomas Hobbes, *Leviathan*, edited by Edwin Curley (Indianapolis: Hackett Publishing Company, 1994), 76. All references to Hobbes will be to this edition.
6. *Leviathan*, 78.
7. Raphael, *British Moralists*, vol. I, 194.
8. *Ibid.*, 193–194.
9. *Ibid.*, 194.
10. *Ibid.*
11. John Balguy, *The Foundation of Moral Goodness* (London: 1728 and 1729).

12. Richard Price, *Review of the Principal Questions and Difficulties in Morals* (London; first published 1758). Selections are available in Raphael, *British Moralists 1650–1800*. I will quote from this edition.
13. Thomas Reid, *Essays on the Active Powers of Man* (Edinburgh: 1788).
14. William Whewell, *The Elements of Morality* (Cambridge: 1845).
15. G. E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1903).
16. W. D. Ross, *The Right and the Good* (Oxford: Oxford University Press, 1930).
17. H. A. Prichard, *Moral Obligation* (1949) and *Duty and Interest* (1929) (Oxford: Oxford University Press, first published together in 1968).
18. Thomas Nagel, *The View from Nowhere* (New York: Oxford University Press, 1986), chap. 8.
19. Derek Parfit, *Rediscovering Reasons*. In manuscript, to be published by Oxford University Press.
20. Raphael, *British Moralists*, vol. II, 162.
21. Even Clarke's extremely vague gesture towards deriving the fitness of actions from the relations between things comes in for an implied criticism from Price: "the term *fit* signifies a simple perception of the understanding." Raphael, *ibid.*, 161.
22. The "naturalistic fallacy" is misnamed because Moore believed that his argument showed not only "naturalistic" definitions of "good" but also "metaphysical" ones to be impossible.
23. See Cudworth, in Raphael, *British Moralists*, vol. I, 106; Hutcheson, also in Raphael, vol. I, 295–296, 305; Price, in Raphael, vol. II, 141, 161; Reid, in Raphael, vol. II, 271.
24. Raphael, vol. II, 142.
25. Not accidentally, because Moore himself was one of the founders of analytic philosophy.
26. This discussion is largely extracted from *The Sources of Normativity* 2.3.1–2.3.9, 67–78.
27. Bernard Williams, *Ethics and the Limits of Philosophy* (Cambridge, Mass.: Harvard University Press, 1985) 136.
28. *Ibid.*, 139.
29. *Ibid.*, 149.
30. *Ibid.*, 140–141.
31. Williams says that thick concepts often provide reasons for action (or refraining) but of course in one way this is not true of "cowardly." To say that an action is cowardly is to suggest that there is a reason not to do it but not to mention what that reason is. Something in the situation is worth overcoming human fearfulness for, but the term does not tell us what. This is because courage is a so-called executive virtue. Williams's other examples suggest more directly why the action shouldn't be done, but they are still not strictly speaking action-guiding. One does not avoid an action because it is brutal, for instance. One avoids it because it will hurt someone.
32. Williams, *Ethics and the Limits of Philosophy*, 141–142.
33. *Ibid.*, 155.

34. The argument of this paper could give the reader the impression that I think a realist account is appropriate for science but not for ethics. This is the right moment to signal that that is not my view. But I do think that theoretical concepts have, as I suggest in the text, something like a map-making function: they enable us to find our way around. That should not be confused with a descriptive function; the qualities of a good map are by no means the same as the qualities of a good description. What is necessary for my argument is not “realism for scientific concepts, a different function for ethical ones,” but rather “concepts that enable us to find our way around in the world vs. concepts that enable us to solve practical problems.”

35. This discussion is extracted from “Rawls and Kant: On the Primacy of the Practical.” In the *Proceedings of the Eighth International Kant Congress, Memphis 1995*, ed. H. Robinson (Milwaukee: Marquette University Press, 1995), 1165–1173.

36. This problem is much more clearly in focus in *Political Liberalism* (New York: Columbia University Press, 1993) than it was in Rawls’s earlier book, *A Theory of Justice* (Cambridge, Mass.: Harvard University Press, 1971), and the modifications in the way Rawls presents his view in the later work are largely due to his increased appreciation of its depth and difficulty. Nevertheless, as I am about to argue and as Rawls himself believes, the strategy of the argument in *A Theory of Justice* does provide for its solution.

37. Kant, *Groundwork of the Metaphysics of Morals*. In the Prussian Academy edition pagination found in the margins of most translations, 446–447.

38. For an explanation of why Kant thinks this, see my “Self-Constitution in the Ethics of Plato and Kant,” *The Journal of Ethics* 3(1999): 1–29, 23–27.

39. I defend this interpretation of Kant’s solution at greater length in my essay “Morality as Freedom” in *Creating the Kingdom of Ends* (New York: Cambridge University Press, 1996).

40. This is a deliberately general statement of Rawls’s two principles, which he states with increasing specificity as he develops his view. For more exact formulations, see *A Theory of Justice*, 302–303, and *Political Liberalism*, 5, 271.

41. For the idea of constructing a solution from the “original position” see *A Theory of Justice*, 17–22, and *Political Liberalism*, 22–28.

42. Rousseau says that the problem solved by the social contract is to “Find a form of association which defends and protects with all common forces the person and the goods of each associate, and by means of which each one, while uniting with all, nevertheless obeys only himself and remains as free as before.” Quoted from *On The Social Contract, The Basic Political Writings of Jean-Jacques Rousseau*, trans. D. A. Cress (Indianapolis: Hackett Publishing Company, 1987), 148.

43. This discussion is extracted from *The Sources of Normativity*, §§3.4.1–3.4.3, 113–116.

44. When we think of the subject this way, we will not be inclined to think that there is a difference between doing “meta-ethics” and doing “normative” or practical ethics. The attempt to specify the meaning and reference of an ethical concept will point fairly directly to practical ramifications. This represents another way in which constructivists break with the platitudes of twentieth century ethics—and return to the more substantive ethical theorizing of the past.

45. But see note 34.

46. *Ethics and the Limits of Philosophy*, 147.

47. At least this is Rawls's view in *A Theory of Justice*. See §§68, 446.

48. See T. M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass.: Harvard University Press, 1998), 17.

49. I haven't said much in this paper about expressivism, our latter-day form of non-cognitivism. So I cannot adequately develop the ideas I am about to voice here. 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.

50. This paper was originally written for the meetings of the Australasian Association of Philosophy in the year 2000, and then revised for the American Philosophical Association's Special Sessions on Philosophy in America at the Turn of the Twentieth Century in 2001. I'd like to thank the audiences on both occasions, as well as the audiences at the Harvard/MIT Graduate Student Philosophy Conference and at Cornell University in spring 2002. I owe a special debt to Jay Wallace for his commentary at the American Philosophical Association, and to Nicholas Sturgeon for comments written when I delivered the paper at Cornell. And I'd like to thank Charlotte Brown and Arthur Kuflik for comments on the early draft version.