

2014 EDIT

Chapter 6

Because It's Right

Morality teaches us that, if we look on her only as good for something else, we never in that case have seen her at all. She says that she is an end to be desired for her own sake, and not as a means to something beyond. Degrade her, and she disappears.

F. H. Bradley (1927, 58)

1. DOES MORAL PHILOSOPHY REST ON A MISTAKE?

Morality can be painfully demanding, so much so that we sometimes question the wisdom of complying with it. Yet, arguments that we have good reason to be moral are as old as Plato's *Republic*. Indeed, according to H. A. Prichard, making this argument work is the central preoccupation of moral philosophy. But Prichard also believes that to the extent this is true, the whole subject of moral philosophy rests on a mistake (1968a, 1).¹

Prichard is neither the first nor the last person to dismiss an entire discipline as a mistake, but Prichard has an argument that poses a real challenge to moral philosophy, an argument that repays sympathetic analysis. Prichard's article emerges from a particular and peculiar philosophical tradition known as British intuitionism, yet the challenge it poses to moral philosophy is anything but parochial. On the contrary, the article has had and continues to have an influence independently of, even in spite of, the intuitionist tradition from which it emerges. For example, it anticipates and to some extent undoubtedly inspires the recent antitheory movement in ethics.² Nevertheless, although dozens of articles cite Prichard's famous essay, often with approval, it has seldom met with sustained criticism.³ This paper reconstructs and criticizes Prichard's argument, then uses that critique to lay foundations for the larger project of constructing a plausible moral philosophy.

Prichard says we begin to question whether we really ought to do our alleged duty—to keep a promise, for example—when we recognize that doing our duty will not give us what we

desire. We then question things we usually accept as duties. We ask if there is any proof that we truly have a duty to act in ways usually called moral.⁴ Prichard sees two ways of interpreting this request. We could be asking whether being moral is prudent. Alternatively, we could be asking whether being moral is good in some nonprudential sense—good for others, for example, or intrinsically good quite apart from its consequences (1968a, 2). Prichard thinks both versions of the question are mistakes, and I will look at each in turn.

How can we determine what is moral in the first place? We cannot simply check what is moral. At least, we cannot do so in the same way we can check who is Prime Minister. Nevertheless, like the word ‘Prime Minister,’ the word ‘moral’ is a word we inherit from an existing language. It comes to us laden with meaning. We can stipulate what we will be referring to when we say “brillig,” for that is not a term of ordinary language, but there are only so many things we could correctly call eggplant. Like the word ‘eggplant,’ the word ‘moral’ is more than a made-up sound. We cannot simply stipulate that it refers to, say, the property of maximizing utility, any more than we could stipulate that the word ‘eggplant’ refers to rutabagas.

A term’s *extension* consists of the set of things to which the term refers. The term ‘Prime Minister’ may under certain circumstances have Jean Chrétien as its extension. Even so, we would not want to say Jean Chrétien is the *meaning* of the term ‘Prime Minister.’ One implication is that we might not know who is Prime Minister, despite knowing exactly what the term means. Similarly, even if we settle what the word ‘moral’ means, we can still be uncertain about what is in fact moral.

As it actually happens, though, we tend to be surer of the word’s extension than of its meaning. We have a shared understanding that being moral involves being honest, kind, peaceful, etc. (I will refer to this consensus as common sense morality.) It may not be part of the meaning of ‘moral’ that honesty is moral, but honesty may be and commonly is understood to be part of the term’s extension.

Moreover, the consensus is not only that we should call these things moral, but also that we should *be* these things, which gives us a clue to the word’s meaning. When a person refers to an act by saying, “That’s immoral” listeners normally understand the speaker to be saying there is reason not to do the act.⁵ Further, listeners will interpret the speaker as saying something other than that the act will not satisfy an agent’s desires. When a person says lying

is immoral, listeners normally will understand the speaker to mean there is a *special* reason not to lie—special because it is grounded in something other than an appeal to the agent’s desires.

This way of understanding the term’s use may not fully capture the term’s meaning, any more than a set of injunctions to be kind, honest, and peaceful fully covers morality’s extension. The conclusion (so far) is only that moral reasons are understood to appeal to something other than the agent’s desires. Moral reasons are *categorical*, which means they have a claim on us independently of how they appeal to our interests and desires.⁶

When people argue about what is right, they may disagree about what constitutes this special kind of reason. Or they may agree that the property of maximizing pleasure constitutes a special reason for endorsement, but quarrel over which actions (or character traits or institutions, etc.) have this property. Even so, when people argue about whether something like affirmative action is right, they have a shared understanding that it *matters* whether affirmative action is right. People who argue about what is moral share an understanding that for an act to be morally required is for there to be a special reason to do it. That is why people care about what conclusion they reach regarding whether something like affirmative action is morally required (or forbidden). As they see it, whether they have special reasons to support (or resist) the practice goes hand in hand with whether the practice is morally required (or forbidden).

But do we need to prove we have such special reasons? As Prichard sees it, moral philosophy rests on the mistaken assumption that we do—a mistaken assumption that without proof that we have special reasons, we have no basis for saying we *ought* to conform to common sense morality. Why is this assumption a mistake? Prichard asks us to consider how we would prove that conforming to common sense morality (which I will refer to as being CS-moral) is moral. According to Prichard, there are two ways to try to prove that being CS-moral is moral, and each inevitably fails. The first way is to prove that being CS-moral will give us something we want (1968a, 3). The second way is to prove there is something good (not necessarily for us) either in right action’s result or in right action itself. Prichard’s objections to these two approaches are as follows.

The first way fails because proving that being CS-moral will give us what we want is beside the point. The demonstration may show that being CS-moral is prudent, but not that being CS-moral is moral. As Prichard puts it, the exercise might convince us that we want to be CS-moral but cannot convince us that we ought to be (1968a, 3).⁷ To show that being CS-

moral is moral, we have to show that we have characteristically moral reasons to be CS-moral, i.e., reasons that at a minimum do something more than appeal to our desires.

The second way, according to Prichard, boils down to saying happiness or working for happiness is good and therefore we should work for happiness in general (or if not for happiness, then for whatever the fundamental good happens to be). This answer has an advantage over the first approach, for at least it clearly does more than appeal to our desires. (Even if the act is for our own good only, this goes somewhat beyond mere appeal to desires.) But this second way also fails, Prichard says, for it presupposes the view that the rightness of acts has to do with what they accomplish. The “fatal objection” to any teleological theory “is that it resolves the moral ‘ought’ into the non-moral ‘ought’, representing our being morally bound to do some action as if it were the same thing as the action’s being one which we must do if our purpose is to become realized” (1968b, 117).⁸

So goes my reconstruction of Prichard’s argument.⁹ In summary, the rightness of keeping a promise, say, does not depend on whether keeping it will have good results at all, let alone on whether keeping it is in the promisor’s interest. Because attempts to prove we ought to do what we believe is right inevitably appeal in one way or another to the goodness of doing what we believe is right (1968a, 2), Prichard concludes that the only place to look for an answer to the question of why we should do what is right is manifestly the wrong place to look. The reductionist urge to ground rightness in something more fundamental is misguided, for rightness neither can be nor needs to be grounded in anything else. The sense of an action’s rightness is, in fact, absolutely immediate (1968a, 7). We see that being CS-moral is moral by direct apprehension, if we see it at all. Trying to *prove* that being CS-moral is moral is a mistake not unlike the epistemological mistake of trying to prove we are awake when we know we are awake by direct apprehension (1968a, 16). It is an instance of the mistake of seeking a grounding for that which is itself bedrock.

The next two sections respond to Prichard’s argument. I argue that there is no mistake in asking whether being moral is prudent. Then I argue that there is no mistake in asking whether it truly is moral to do things like keep promises.

2. MORALITY VERSUS PRUDENCE

Prichard concedes that it can be perfectly legitimate to ask why we should perform a certain act when the act is incompletely described in relevant ways. The question becomes illegitimate, in Prichard's view, when the act is described well enough that special reasons to perform the act are, in effect, built into the act's description. For example, it may not be obvious that Kate has reason to give her neighbor a hundred dollars, but it is perfectly obvious that she has reason to *repay a debt* by giving him a hundred dollars (1968a, 8). Described in this more complete way, the act carries its reason on its sleeve. When an act is described in such a way that asking why we should do it becomes tantamount to asking why we should do what is required, the answer becomes obvious: we should do it because it is required.¹⁰

Still, an act that is well described in moral terms may remain incompletely described in prudential terms. The question "what's in it for me?" may remain unanswered. We could dismiss the latter question as morally irrelevant, but this would be to ignore the question rather than answer it. Even if Prichard is correct that it is impossible to give an argument why we *morally* ought to do the right thing, this does not foreclose the possibility that philosophers might yet show that it is prudent to do the right thing. Nothing in Prichard's argument counts against undertaking the Platonic project of showing that being moral is profitable.

Prichard then engages the Platonic project more directly. Prichard says proving we have a prudential motive to do the right thing would be beside the point. If we are talking about being moral, we are not talking about doing what is right out of prudence. Rather, to be genuinely moral is to what is right *because* it is right. In Prichard's words, "a morally good action is morally good not simply because it is a right action but because it is a right action done because it is right, i.e., from a sense of obligation" (1968a, 10).¹¹

It may seem, as evidently it seemed to Prichard, that the project of reconciling prudence and morality cannot proceed unless this Kantian line of argument is rebutted. Not so. Even if we grant that being moral involves following a categorical imperative, we remain free to ask whether we are better off following a categorical imperative. And one way or another, the question has an answer. Whether or not moral imperatives are categorical, there remains a fact of the matter concerning whether following moral imperatives is to our advantage. To try to show that being moral turns out to be compatible with prudence is not to assume that moral

imperatives are prudential imperatives. On the contrary, we can try to prove a conditional of the form “If I want X, then I should be moral” without in any way presuming that moral imperatives have this same conditional form.

If we were asking whether prudence can be a proximate *motive* for being moral and if we took “being moral” to entail “being motivated by a sense of rightness rather than by prudence,” then Prichard’s objection would be decisive. The question would be a mistake. The actual question, however, is whether there is an extensional overlap between being moral (and thus being motivated by a sense of rightness) and being prudent, in which case Prichard’s objection misses its mark.¹² Asking whether doing the right thing is prudent does not presume only prudential answers could motivate our being moral. It does not presume prudence is even *among* the things that could motivate our being moral.

Demonstrating the existence of an extensional overlap need not motivate people to be moral. But really, that was never the point. The point is that even agents committed to doing what is right *because it’s right* might nevertheless wonder whether they would have done anything differently had they been more self-consciously prudent. Moral agents might care about this issue not because they seek a motivation for being moral but rather because they, like Glaucon, sometimes wonder whether they have prudential reasons to regret being moral, that is, whether their being moral is contrary to the self-interest. They simply wonder. It may not dampen their moral motivation in the slightest, but still, they would rather know, and the desire to know leads them to philosophize.

In summary, Prichard thinks it is a mistake to try to prove that being moral is for our own good, for the attempt presupposes that whether we ought to be moral depends on whether being moral is prudent.¹³ The presupposition may well be a mistake, but we need presuppose no such thing. Asking whether being moral is prudent does not imply that we view morality or moral motivation as reducible to mere prudence. This version of the question is no mistake.

3. WHAT DO WE DO WHEN WE DO THE RIGHT THING?

The previous section conceded that we should do what is right because it is right, but showed that this concession is hardly a conversation stopper. Whether it is prudent to be moral remains an issue. Further, even from the moral point of view, it is not enough to say we should do what

is right because it is right. As Prichard himself would stress, the question we face as moral agents is not about philosophical generalizations, but rather about what to do when we get face to face with particular situations. And saying we should do what is right would be to miss the point of our asking what we should do. The point is, we need to have concluded that X *is* right before the incantation “because it’s right” can single out X as something we have reason to do. Naturally, we should do what is right, and let’s concede that we should do so because it is right. But why should we keep promises? Why, in some rare cases, should we break them? Why should we tell the truth? Why, in some rare cases, should we lie instead? Why think keeping promises and telling the truth have anything to do with rightness?

“Why should I do what is morally required?” is the sort of question that wears its moral answer on its sleeve, even if it does not wear its prudential answer on its sleeve. But that is not the same kind of question as “why should I tell the truth?” Rightness may wear moral motivation on its sleeve, but what rightness patently does not wear on its sleeve is its extension. Indeed, the question of which particular actions are right remains wide open. So Prichard has not only left undone the legitimate task of identifying prudential reasons not to regret being moral; he has also left us the more fundamental task of identifying what morality requires. (I don’t mean to say it was Prichard’s job to do this, but Prichard seems to say moral philosophy rests on the mistake of thinking this task of identification can be a philosophical task.)

The latter was no accident, of course, for Prichard was, after all, an intuitionist. He says we intuit what is right. Be this as it may, the question in which we are actually interested is logically prior to this epistemological question. That is, even if we grant that there are times when we learn that act X is right by intuiting that X is right, we still want to know what it is about X that triggers our intuitions.¹⁴ Consider this: if we had no idea what triggers our intuitions, what grounds would we have for taking our intuitions seriously?¹⁵

One might insist that intuitionism is not only an epistemological thesis but also a thesis about what rightness is; a right action simply is an action that directly and immediately strikes us as something we have reason to do. I do not believe Prichard held this ontological thesis, but in any event, this ontological variant of intuitionism amounts to a rather sinister reductionism. It reduces rightness to the sphere of that which directly and immediately strikes us as required. Consider what it implies about things we do not directly and immediately apprehend as required—things whose rightness (or wrongness) we do not come to fully

appreciate merely by getting face to face with them. If we cannot directly apprehend that keeping a certain promise is required, may we rule out *on those grounds* the possibility that keeping the promise is required? Surely not.

If we take intuitionism to be addressing the question of what rightness is, we are taking it to be an alternative *kind* of reductionism rather than an alternative *to* reductionism. It is more charitable to accept that Prichard's intuitionist epistemology leaves open the ontological question about what properties occasion our intuitions.

Perhaps we learn general principles by generalizing from particular instances. We get face to face with particular instances, as Prichard says, and then learn general principles by induction.¹⁶ Even so, the order in which we learn particulars and general principles is not the issue here. Even if we learn particulars first, there must be something about particular requirements that makes them requirements. Regardless of whether we learn the particulars first, a question inevitably remains regarding what we are seeing in a particular act when we see it as required. What makes promise-keeping rather than promise-breaking required? And why do we think promise-keeping in some exceptional cases is not required after all, and may even be forbidden? What makes those cases different? That we see them differently is not what makes them different. We need to identify what is being seen when some cases of promise-keeping are seen as required and others as forbidden or at least not required.

The list of required acts has to be more than a mere list. If membership in the category were determined arbitrarily, then Prichard would be wrong, for in that case membership in the category of required acts would not imply any special reason to do the act. Prichard wants to say that an act being correctly labeled "required" is itself a good reason to do it—so good that we need no other reason. I am not quarreling with this. My point is only that if our intuitions are picking out some things as right and others as wrong, and doing so in a nonarbitrary way, this implies that acts we intuitively identify as right differ in some nonarbitrary way from acts we intuitively identify as wrong. What then is the difference?

One might think this misses the real point, which is that to call an act required is to *state* a special reason to do it. But suppose we *mistakenly* call an act required. In that case, we *think* we have stated a special reason to do it, whereas in fact there is no such reason to do it. We could say that to *correctly* call an act required is to state a special reason to do it, but then we still need to know what it is in an act that makes it true that the label "required" is correctly

attached. If Prichard is correct in saying special reasons for action are entailed by an act's being required, then we cannot label an action 'required' (or more precisely, we cannot know we have labeled the action correctly) until we know we have the requisite reasons for attaching the label, i.e., that there really are special reasons for doing the act in question. We do not create the special reason merely by (perhaps mistakenly) applying the label.

For an act to be right, there must be a reason why it is right. Prichard's concern—that deriving a sense of rightness from something else would run contrary to our actual moral convictions (1968a, 4)—is baseless. Indeed, if there were nothing in the keeping of a promise to ground our judgment that it is right, then the judgment itself would be baseless, which is contrary to our convictions if anything is.

Prichard is correct to say we already have a reason to perform an action when we see that it is required. We do not need to know what makes actions required in order to know we ought to do what is required. Still, one can ask what makes required actions required; in which case, we had better have something to say about when there is good reason to see an action as required. To answer questions of that sort, we need a rule of recognition for morals.

4. A RULE OF RECOGNITION FOR MORALS

I argued against Prichard on two fronts. First, we can have something to say about whether being moral is prudent. Second, while we would be mistaken if we thought we still needed to *prove* we have special reason to do X even after coming to see that X is required (since it would be hard to see X as morally required unless we already see the special reasons), we are not mistaken in hoping that philosophy can answer questions about what makes right actions right in particular cases. It is time to consider what these critical points tell us about the more positive task of constructing a moral theory, while trying to be mindful of Prichard's challenge to modern moral philosophy.

My approach to moral theory begins by borrowing from H.L.A. Hart. Hart's legal theory distinguishes between primary and secondary legal rules (1961, 89-93). Primary rules comprise what we normally think of as the law. They define our legal rights and obligations. We use secondary rules, especially rules of recognition, to determine what the law is.¹⁷ For example, among the primary rules in my neighborhood is a law saying the speed limit is thirty

miles per hour. The secondary rule by which we recognize the speed limit is: read the signs. Exceeding speed limits is illegal, but there is no further law obliging us to read signs that post the speed limit. So long as I stay within the speed limit, the police do not worry about whether I read the signs. In reading the signs, we follow a secondary rule, not a primary rule.

We can think of moral theories in a similar way. For example, utilitarianism's recognition rule is the principle of utility: X is moral if and only if X maximizes utility. As it stands, the principle defines a family of moral theories rather than any particular member thereof. The different flavors of utilitarianism are produced by replacing X with a specific subject matter. Act-utilitarianism applies the principle of utility to actions themselves. Act-utilitarianism's fully specified recognition rule—an act is right if and only if it maximizes utility—then translates directly into act-utilitarianism's single rule of conduct: maximize utility. Rule-utilitarianism applies the principle of utility to sets of action-guiding rules. The resulting recognition rule states that an action guide is moral if and only if following it has more utility than would following any alternative action guide. Of course, the utility-maximizing set of primary rules might boil down to a single rule of conduct saying "maximize utility." Then again, it might not.¹⁸

Deontological theories are harder to characterize. We could begin with a generic recognition rule saying X is moral if X is universalizable.¹⁹ Applying the rule to maxims yields a more specific recognition rule (something like: a maxim is moral if acting on it is universalizable), which in turn yields a set of imperatives, reverence toward which is grounded in considerations of universalizability. Perhaps the idea of universalizability does not have enough content to yield determinate imperatives on its own. Deontology may need a second recognition rule formulated in terms of respect for persons as ends in themselves, so that the two rules can converge on a set of concrete imperatives. But that is another story.

A moral *theory* consists of a recognition rule applied to a particular subject matter. Given a subject matter, a rule of recognition for morals specifies grounds for regarding items of that kind as moral. By "grounds" I do not mean necessary and sufficient conditions. In act-utilitarianism, the principle of utility presents itself as necessary and sufficient for an act's morality, but trying to contrive necessary and sufficient conditions is not the only way (and I think not the best way) to do moral theory. To have a recognition rule, all we need is what I call a *supporting* condition.

I introduced the concept in Chapter 1. As I said there, a *supporting condition* is a qualified sufficient condition, qualified in the sense of being a sufficient basis for endorsement in the absence of countervailing conditions. Formulating recognition rules in terms of supporting conditions rather than attempting to specify necessary and sufficient conditions is one way of acknowledging intuitionist claims that we could never fully articulate all of the considerations relevant to moral judgment. We can allow for that possibility (without letting it stop us from doing moral theory) by formulating recognition rules in terms of supporting conditions—conditions that suffice to shift the burden of proof while leaving open a possibility of the burden being shifted back again, perhaps by considerations we have yet to articulate.

As an example of a supporting condition, we might say, along the lines of act-utilitarianism, that an act is right if it maximizes utility, barring countervailing conditions. In two ways, act-utilitarianism properly so-called goes beyond merely offering a supporting condition. First, it denies there are countervailing conditions, thereby representing the principle of utility as a proper sufficient condition, not just a supporting condition. Second, act-utilitarianism says an act is right only if the act maximizes utility, thereby representing the principle of utility not only as sufficient but also as necessary for an act's morality.²⁰

5. THE NORMATIVE STATUS OF MORALITY'S RECOGNITION RULES

H.L.A. Hart, himself a legal positivist, argued that rules of recognition for law may or may not pick out what is moral when they pick out law. Herein lies a crucial disanalogy between rules of recognition for morals and for laws. Questions about legality are sometimes answered by simply "looking it up." Arguably, we do not need to know we have moral reason to obey a law in order to recognize it as law. Legal positivism is, roughly speaking, the thesis that a recognition rule can correctly pick out a rule of conduct as legal even though the rule is immoral. But there can be no such a thing as moral positivism, since it is not possible for a rule of recognition to *correctly* pick out rule of conduct X as right when X is *not* right. It may not be essential to *laws* that they have an inner morality, but we can entertain no such agnosticism about morality itself. It is in the nature of Prichard's conception of morality (unlike law) that a recognition rule can correctly identify actions as morally required only if there is decisive reason (absent countervailing conditions) to perform them. Only such a

recognition rule lets us stop the conversation—as Prichard would want to stop it—upon concluding that our recognition rule identifies an action as morally required.

We need to say more about what it means to regard X as right. I will approach this issue by starting with a different question; namely, what is being questioned when a person asks “why be moral?”

First, when asked in earnest, “why be moral” is a question about something that matters. “Why stand on one foot” is, on its face, an idle question, but “why be moral” is not. Second, the “why be moral” question matters despite the fact that it patently does not presume that being moral matters to people from their first person singular perspectives. Whether people have first person singular reasons to be moral is pointedly left open. Thus, the implicit urgency comes from another source.

It stems, I would say, from the fact that morality essentially is something that matters to us from a first person plural perspective. My endorsement begins to look like characteristically moral endorsement when grounded in the thought, not that *I* have reason for endorsement, but that *we* have reason for endorsement. While endorsement as rational need not go beyond the first person singular, endorsement as moral at a minimum goes beyond first person singular to first person plural.²¹

The second thing to say is that the transcendence of the singular perspective involved in moral endorsement cannot go much farther than this. If moral endorsement involves taking a plural perspective, then we can imagine how being moral could be disadvantageous for you or me and yet we could still have clear reason to endorse being moral. For example, many theorists now think of cooperating in a Prisoner’s Dilemma as a paradigm case of being moral.²² While disadvantageous from an *I*-perspective, cooperating remains rational in the sense of being to our advantage from a *we*-perspective. It is from a plural perspective that, in a Prisoner’s Dilemma, we find something irrational about individual rationality. When you and I each decide not to cooperate, I am doing the best I can given your noncooperation, and you are doing the best you can given mine, and yet *we* are not doing the best we can. However, if being moral were pointless not only from a singular perspective but from a plural perspective as well, then it would be pointless, period. Being moral would be something we would have reason to avoid in ourselves and condemn in others. Being moral, though, is not like that. Being moral

need not be prudent from a singular perspective, but part of the essence of being moral is that we have reason to endorse it from a plural perspective.

One thing that makes moral reasoning different from legal reasoning is that questions about how we recognize morality are hard to separate from questions about whether we have reason to endorse it. Morality's recognition rules pick out the extension of 'moral' just as the law's recognition rules pick out the extension of 'legal.' Morality's recognition rules, however, pick out X as moral by homing in on properties that, from a plural perspective, give us reason to endorse X. We sometimes can discern the rules of the road by reading the signs. We sometimes can discern the applicable law simply by looking it up. Analogously, we might sometimes be able to discern what is moral simply by consulting what we (correctly) take to be a moral authority. But in formulating a theory about what makes something moral, we are seeking to identify *truth-makers*. So, although recognition rules essentially serve an epistemological role, they serve that role by tracking moral ontology. Moreover, to constitute the sort of theory that could play a recognition rule's epistemological role in a moral agent's life, we have to be talking about *usable* truth-makers. A theory's recognition rules, then, have to direct us to look for a kind of truth; moreover, they have to direct us to look for a kind of truth that we are capable of finding. (What else could morality be?)²³

It is because morality is bound up with what we have reason to endorse from a plural perspective that "why be moral" is a pressing question. The "why be moral" question we inherited from Plato is a question about the relation between two kinds of *telos*—between what matters to us as individuals and what matters to us as a society. Because morality, as we conceive of it and as Plato conceived of it, matters to us from a first-person plural perspective, we have reason to hope it matters to us (or can be made to matter to us) from our first-person singular perspectives as well.

If, *per impossible*, morality did not matter from our plural perspectives, then neither would it matter whether morality could be reconciled with our singular perspectives. In different words, both Socrates and Glaucon care from a plural perspective about Glaucon's being moral. They are treating the question of whether morality can be reconciled with Glaucon's singular perspective as up in the air, yet there is some perspective, some other perspective, from which it is not up in the air. They *want* the answer to be that morality is reconciled with Glaucon's singular perspective. Analogously, we care whether people

cooperate in a Prisoner's Dilemma. It is in the interest of both players to decline to cooperate, so from their singular perspectives it makes to sense to be trying to convince them to cooperate. What makes sense of caring whether they cooperate is that there is a different perspective, a plural perspective, from which cooperating will make them better off.

Who Are We?

This takes us to one of the points at which satisfactory moral theorizing becomes really, really difficult. Unfortunately, while the scope of a person's *I*-perspective is more or less fixed (encompassing the person's own interests and preferences), the *we*-perspective does not have fixed borders, making it hard to characterize the *we*-perspective with any precision. It should go without saying, though, that the plural perspective is no mere fiction. (It is not for nothing that natural languages have words like 'we' and 'us' for plural self-reference.) When I speak of the *we*-perspective, what I have in mind is not the sort of group perspective you and I might take when we identify ourselves as fellow Diamondbacks fans, but rather the particular perspective we take when we worry about the "why be moral" question.

That perspective usually does not encompass the whole world.²⁴ If I see that my mowing the lawn will hamper your efforts to write your book, then my taking a *we*-perspective involves identifying with you as a member of the group of people who will be affected by my mowing the lawn. If I see that mowing the lawn will adversely affect people in a far away country (because they are waiting anxiously for your book), then my taking a *we*-perspective involves identifying with them as well. The scope of my *we*-perspective expands and contracts along with my awareness of whose interests are at stake.²⁵ This does not mean I should not mow the lawn. We could not live together if we did not allow ourselves the latitude to impinge on each other in various ways. Your latitude may not serve my ends, and mine may not serve yours, but what is relevant from the plural perspective is that *our* latitude serves our ends. We are better off in virtue of members of our group having that kind of latitude.

6. THE DESCRIPTIVE BOUNDARIES OF MORAL INQUIRY

This, then, is the normative status of morality's recognition rules. Being recognized as moral has normative force because, when morality's recognition rules pick out X as moral, they do so by recognizing that X has properties we have reason to endorse from a plural perspective.²⁶

Consider this objection. Kate has reason from a plural perspective to endorse Disneyland. "We'll have a lot of fun there. Nearly everyone does," she says to her friends. Yet, though Kate endorses Disneyland from a *we*-perspective, she is endorsing it not as moral, but as amusing, or something like that. To endorse something as moral is to endorse it from a plural perspective, but not everything endorsed from a plural perspective is thereby endorsed *as moral*.

I agree with the objection. Certainly, we should not equate endorsing Disneyland from a plural perspective with endorsing Disneyland as moral. How then should we think of the plural perspective's role in moral theory? From a plural perspective, we do not pick out maxims (for example) as moral. Still less do we pick out Disneyland as moral. Instead, we pick out a criterion for assessing maxims, given that maxims are subject to moral assessment.

Now, if something is a lot of fun for almost everyone, why is that not a property that we have reason to endorse from a plural perspective? Or if being a lot of fun is such a property, then what distinguishes endorsing something as fun from endorsing it as moral? Section 4 noted that it is not the task of recognition rules to circumscribe their own subject matter. On the contrary, any theory pretty much takes a subject matter as given. There has to be a subject that gives rise to moral questions before we can have occasion to devise theories to answer those questions. We begin with an intuitive understanding that subjects giving rise to moral questions include (roughly speaking) things that bear on human flourishing in communities, regarding which human action can make a difference. (Let me stress: I do not treat intuitions as recognition rules for morals. On my theory, intuition enters the picture as a source of questions, not as a tool for answering them.) The subject matters of moral inquiry are pretheoretically given, i.e., given in the sense of raising moral questions prior to our devising moral theories to answer them.

Accordingly, "X is moral if X is a lot of fun" is not a recognition rule for morals because, when applied to any of the specific subject matters over which moral theories range,

the property of being a lot of fun is not a reason for endorsement from a plural perspective. We do not in fact recognize it as reason to endorse capital punishment or promise-keeping or any of the subjects that normally raise moral questions.²⁷ We do recognize it as a reason for endorsement when the subject is amusement parks, but that would make it a basis for moral endorsement only if amusement parks as such were among the pretheoretically given subject matters of moral inquiry, and they are not. A moral perspective is more specific than a plural perspective not because it is a more narrowly defined perspective but rather because it consists of taking a plural perspective only with respect to issues already defined, intuitively and pretheoretically, as moral issues.

To summarize, a recognition rule like the principle of utility could embody a genuine reason for endorsement from a plural perspective and still fail to exclude Disneyland as a subject for moral assessment. However, it is not incumbent on recognition rules to have the internal resources to limit their subject matters. We test a purported recognition rule not in abstraction but rather as applied to a pretheoretically given subject matter. We test it by asking whether it homes in on a property that, given the subject matter, grounds endorsement from a plural perspective. For example, if we apply the principle of utility to Disneyland, and then afterwards decide that Disneyland as such is not a subject of moral inquiry, it would be a mistake to blame the principle of utility for the misapplication.

We considered how recognition rules distinguish what is right from what is not, given a subject matter with respect to which such questions arise. I have no theory to tell me what the subject matters of moral assessment are; on my theory, that is a *pre-theoretical* question. I have only a sense that morality and moral assessment concern what makes it possible for human beings to flourish together. Given this pretheoretical understanding of the general character of the subject matters of moral assessment, amusement parks are not among morality's subject matters, but *institutions* are. Thus, Disneyland is subject to moral assessment not as an amusement park but rather as an institution that has a bearing on whether people flourish within their communities. (Similarly, Michael Jackson is subject to moral assessment not as an entertainer but rather as a person whose choices had an impact both on himself and on many other people.) Likewise, acts, rules of conduct, and character traits are subjects of moral assessment because they affect whether people flourish within communities.

No doubt our intuitive conception of the proper subjects of moral assessment is more complicated than this, and I am not proposing to shed much light on our intuitive and pretheoretical understanding of the descriptive boundaries of moral assessment. It remains that, given an understanding of the subject matters of moral inquiry—of the kinds of things concerning which moral questions arise—we have something about which we can theorize. We can devise a theory about how those questions should be answered and why.

The descriptive boundaries of the subject matters of moral inquiry are given prior to our doing moral theory. They define what we want to have a theory about. Given a predefined subject matter, my proposal is that we capture the normative bite of morality's recognition rules when we say they home in on properties that, with respect to that subject, we have reason to endorse from a plural perspective. If amusement parks are not among the subject matters of morality, then morality's recognition rules do not range over amusement parks in the first place, which is why morality's recognition rules cannot pick out Disneyland as such as moral.

As with other intellectual endeavors, we need some sense of a subject matter and of questions to which it gives rise before we can have any reason to devise theories about it. Long before we begin to formulate moral theories, we already classify certain issues as moral issues. Roughly speaking, when an issue is crucial to human flourishing in communities, and when human beings can make a difference regarding that issue, we tend to see it as raising moral questions, and thus as a subject calling for moral theory. In this sense, the subject matters of moral inquiry are (at least provisionally) a pretheoretical given.²⁸

7. IS THE RIGHT PRIOR TO THE GOOD?

One might worry that if we analyze the rightness of acts in terms of the goodness of states of affairs, the concept of rightness loses its turf, so to speak. The concept becomes superfluous, and we may as well dispense with it entirely. But this worry is not well-founded. To explain our grounds for identifying an act as right is not to explain rightness away. The explanandum does not disappear merely in virtue of having been explained. In different words, giving an account of an action guide's normative force does not eliminate the need for an action guide. We cannot dispense with talk about what is right because we cannot dispense with talk about what we should do. We can, however, still speak of keeping promises because it is right (or

because breaking promises would violate rights). We thus go beyond Prichard's understanding of what moral philosophy can do without rejecting Prichard's understanding of morality and moral motivation per se.

Only when we ask how we *recognize* that keeping promises is right (or that breaking promises would violate rights) do we move from morality proper to moral epistemology: that is, from questions addressed at the action-guiding level to questions addressed by recognition rules. This is crucial. If we thought of recognition rules as part of morality's action guide, we would be missing the distinction between recognition rules and rules of conduct. To properly address Prichard's objections to teleology, a theory must isolate its teleology at the level of recognition rules, so that the concept of rightness can take on a life of its own at the action-guiding level. When a theory's teleology is embedded in recognition rules, it specifies terms by which we recognize what is required, in the process leaving moral agents with an action guide that tells them what is required and which they follow because doing so is required.

In short, recognition rules, which have a teleological spirit, support action guides, parts of which may well be categorical. In turn, action guides support particular actions or choices. For example, suppose the principle of utility is morality's recognition rule and that suppose it recognizes ten commandments against lying, stealing, etc., as morality's rules of conduct. If we saw the principle of utility as morality's ultimate rule of conduct, then we would interpret the ten commandments as rules of thumb—rules that give way to the principle of utility when it is obvious that following them will not maximize utility. A recognition rule, however, is *not* an ultimate rule of conduct. Rather, it identifies morality's rules of conduct, a set of ten commandments in the case we are imagining, and the ten commandments are thereby certified as the ultimate rules of conduct. Conduct is judged according to whether it follows the ten commandments, not whether it maximizes utility.

a) *Reading the Signs*

Consider a legal analogy. "Read the signs" may be the rule by which we recognize rules of the road, but if we found ourselves in a situation where obeying a speed limit would somehow prevent us from reading a traffic sign (imaging needing to speed up to pass a semi-trailer that otherwise will block your view of the upcoming sign), that would not be enough to make the

speed limit give way. It would not even begin to make the speed limit give way. The highway patrol judges our conduct by the rules of the road, and would be properly unimpressed if we said we violated the rules of the road out of commitment to a “higher law” bidding us to read the signs. “Read the sign” is not a higher law. It’s a mere recognition rule, nothing more.

Given that recognition rules are not rules of conduct, ultimate or otherwise, it is entirely possible that some of morality’s rules of conduct are deontological (that is, they make no appeal to consequences) even if morality’s rule of recognition is teleological.²⁹ A rule by which we *recognize* deontological imperatives can be teleological without in any way affecting the deontological force of the imperatives thus identified. An imperative may dictate an action without appealing to the action’s role in serving the agent’s purposes; indeed, it may dictate action without appealing to *anyone*’s purposes. This leaves open whether the imperative has teleological support. It may serve a purpose to be committed to keeping promises come what may, even though it sometimes happens that keeping a promise serves no purpose. It serves a purpose to keep regular office hours even though some of those hours predictably will be spent waiting in vain for students to drop by.

A teleological recognition rule, applied to imperatives, is analogous to a rule of recognition in law; it is beside the point except when we are pondering whether a particular imperative is moral. Upon recognizing an imperative as moral, we thereby know what we need to know to see that we have a moral reason to follow it. Having settled that the imperative is morally imperative, the rule of recognition has no further role to play. It drops out, leaving us with action-guiding imperatives that may well present themselves to us in deontological form. In any event, the action-guiding imperative, not the rule of recognition, is what guides action.

A “soft” deontological prohibition is insensitive to consequences in normal cases but makes exceptions in extraordinary cases. We saw how there could be a teleological grounding for imperatives that are normally insensitive to consequences. In contrast, *absolute* imperatives are insensitive to consequences even when the universe is at stake. It is conceivable, just barely, that we could have teleological grounds for recognizing an absolute imperative as moral. It could be good to internalize the rule “I will not lie—not even to save the universe,” so long as it never actually happens that we need to lie to save the universe. I doubt there are any teleologically well-grounded absolute rules of conduct, but the idea is perfectly coherent.

The idea that morality is teleological at the level of recognition rules does not preclude the possibility of there being absolutely exceptionless rules of conduct.

b) *When Good Reasons Are Redundant*

There is, of course, a controversy in moral philosophy over whether the right is prior to the good. Some theorists, if they considered the question at all, would dismiss the idea that morality's recognition rules are teleological; they would assume it contradicts their belief that the right is prior to the good. It would be a mistake to dismiss my theory on that basis, though. My theory is entirely compatible with the view that the right is prior to the good at the *action-guiding* level. We should keep promises because it is right, and at the action-guiding level nothing more needs to be said. But that does not tell us what makes promise-keeping right, or even (in cases of doubt) whether promise-keeping is right. When it comes to recognizing what is right, the good is prior to the right, and must be so. We judge acts in terms of right, but when we need to explain what makes an act right, or whether it is right in a doubtful case, we can do so only in terms of good. So, regarding the controversy over the relative priority of the right and the good, the truth is, (1) the right is prior at the action-guiding level and (2) the good is prior at the level of recognition rules.³⁰

Teleological considerations need not enter a moral agent's deliberations about what to do. If we cannot act without breaking a promise, then under the circumstances that may be all we need to know in order to know we categorically should not act. Sometimes, though, we do not know what morality requires of us. Some promises should not be kept, and we do not always know which is which. When we do not know, we need to fall back on recognition rules, which identify the point of being categorically required (that is, required on grounds that do not appeal to the agent's interests and desires) to act in one way rather than another.

Prichard thinks that if one understands that keeping a particular promise is required, one thereby recognizes reason to keep it. In that case, pointing out that keeping promises has good consequences would be irrelevant. Prichard is right, and now we can see why. When we already see that we are required to keep a promise, pointing out good reasons to keep it is *redundant*. The redundancy of pointing out good reasons to keep a promise, when we already see that keeping it is required, makes the good reasons irrelevant.³¹ But what if we have not yet

recognized that keeping a particular promise is required? In that case, coming to see that breaking the promise would have bad consequences is not redundant at all. In that case it is Prichard's point that is irrelevant, for in that case we are not asking why we should do what is required. Rather, we are asking whether keeping *this* promise is required in the first place.

c) *It's Not Just a Good Idea. It's the Law*

One might be troubled by the idea of keeping a promise simply "because it's right." "Because it's right" may seem oddly abrupt as a reason for action. However, it certainly is not peculiar to morality. For example, when a motorist's impatient passenger asks her why she is driving at twenty miles per hour, it would not be peculiar for her to reply by saying "because it's the law." Her passenger now knows why she is driving at twenty miles per hour and might go on to ask how she knows that it is the law. She might answer that she read the speed limit sign. In a more philosophical if still somewhat impatient frame of mind, the passenger might then ask what is the *telos* of the twenty miles per hour speed limit. What justifies it? The driver may not know. But she still knows the law. Further, if she knows that there is a school in the neighborhood, then she can add that the school's presence justifies the law (and she can say this even though she has no idea whether the school's presence is what actually motivated authorities to impose the speed limit). A conversation about morality might unfold in the same way. Asked why she keeps promises, a person might say "because it's right." She might be asked about the *telos* of promise-keeping, or about how she knows keeping promises is right, but those will be different questions.

In summary, Prichard denies that the good plays a role in determining the right. He infers this from the premise that we keep promises because doing so is right, not because doing so is good. I accept the premise, but the inference is invalid. Of course we should keep promises because it is right, and at the action guiding level this is all that needs to be said. But this is different from asking why promise-keeping is right, or (in cases of doubt) whether promise-keeping is right. To answer the latter questions, noting that we have intuitions about promise-keeping is not enough. (Perhaps our intuitions are tracking the good.) We need to articulate reasons to keep promises. And pointing out that promise-keeping is right is to *imply*

there are reasons rather than to *specify* them. Explaining why we ought to do what is right and identifying what is right in the first place are different tasks.

8. CONCLUSION

We examined H. A. Prichard's argument that the question "why be moral" is fundamentally confused. It turns out that there is no confusion involved in asking the question from a prudential point of view. Asking the question from a prudential point of view does not presuppose any reduction of morality to a system of prudential imperatives. On the contrary, we can intelligibly ask whether following categorical imperatives is to our advantage. One way or another, the question has an answer.

A recognition rule cannot be constituted in such a way that the action guide it picks out is as likely to lead us to do bad as to do good. Morality's recognition rules cannot be arbitrary with respect to goodness. Otherwise, arbitrarily identifying an act as right is not a reason to do it. And the idea that we could identify an act as morally imperative without in the process coming to have a reason to do it is contrary to Prichard's supposition that we should do what is right because it is right. A recognition rule for right action essentially picks out, as right, actions for which there are good reasons, which is precisely what lets us conclude, as Prichard wants us to conclude, that to recognize their rightness is to recognize good reason to do them.

We have not explored any particular theory about the content of morality's recognition rules beyond saying they recognize a thing as moral by recognizing reason to endorse it from a plural perspective. This chapter's burden has been to show why we can safely reject H. A. Prichard's conclusion that undertaking to identify such rules is a mistake. Further, we can reject his conclusion while allowing that his premises are not without merit.

Epilog 2014: So, Why Be Moral?

This Chapter serves two masters. One objective was to come to terms with Prichard and in the process aggressively claim that there is more to say in response to the "why be moral" question than Prichard anticipated. A second objective was to show that there can be an argument for the right even if, in some crucial way, the right is morally prior to the good. I wanted to explain

how the right could have a grounding, and yet it could still be coherently built into the concept of rightness—rightly understood—that one does not necessarily appeal to what grounds rightness when one appeals to rightness. *If one has reason* to believe that x is right, then one can have reason to see "because it's right" as a conversation-stopping reason to do x .

¹ At one time, Philippa Foot (1978, 126) agreed with Plato that “if justice is not a good to the just man, moralists who recommend it as virtue are perpetrating a fraud.” Likewise, David Gauthier (1986, 2) says “the acceptance of duty is truly advantageous.” Kurt Baier (1978) and Kai Nielsen (1989) agree that “why be moral” is a legitimate question, although Nielsen’s answer is pessimistic. On the other side, Prichard’s view that the question itself is illegitimate is endorsed by J. C. Thornton (1970), Dan Brock (1977), and John McDowell (1978), among others. Prichard was not the first to take such a view, either. For example, see Essay II of F. H. Bradley’s *Ethical Studies* or Henry Sidgwick’s introduction to *Methods of Ethics*.

² Antitheorists characterize (and consequently reject) moral theorizing as an attempt to mechanically deduce all particular moral conclusions from a single universal principle. Robert Louden (1992, Chapters 5 and 6) agrees that any theory fitting that description ought to be rejected, but argues that the best and historically most prominent moral theories (i.e., those of Aristotle and Kant) do not fit the description.

³ W. D. Falk (1986) accuses Prichard of equivocating between internalist and externalist moral “oughts.” The only other substantial critique of Prichard, to my knowledge, is Baier (1978, 231-38). For an acute critique of intuitionism more generally, see Stephen Darwall (2002).

⁴ Actually, I suppose today we frame the question in terms of right action rather than moral action. My main reason for moving between these terms is stylistic, speaking of doing what is moral when I am trying to follow Prichard and of doing what is right otherwise. When speaking of persons rather than of actions (*being* moral as opposed to being prudent), or of separating the subject of morality from other subjects such as prudence, I find ‘being moral’ and ‘morality’ more natural than ‘being right’ and ‘rightness,’ but again, my reasons for choosing one word rather than another are in most cases stylistic rather than deeply philosophical.

⁵ It seems easier here to speak of wrongness rather than rightness as being associated with special reasons for action. That one course of action involves telling the truth does not imply that one should take that course, but that another course of action involves telling a lie has clear implications. I thank Rod Wiltshire for the intriguing conjecture that wrongness is a natural kind and rightness is not. Rightness is simply the logical complement of wrongness, in the way “non-dog” is the logical complement of “dog.”

⁶ I use the terms ‘categorical’ and ‘deontological’ almost interchangeably. An imperative is categorical if it makes no appeal to the agent’s interests and desires, and deontological if it makes no appeal to consequences of any kind. Thus, as I use the terms, a categorical imperative is a kind of deontological imperative.

⁷ And, as Stephen Toulmin (1970, 417) adds, making us want to do what we ought is not the philosopher’s task.

⁸ Prichard’s point applies to theories grounding rightness in collective prudence as well. So Prichard’s objection not only challenges the Platonic project but most contractarian theories as well. For example, the objection cuts against the view expressed by Baier (1970) that we should be moral because being moral makes us all better off.

⁹ Prichard also rejects the idea that an action’s rightness lies not in its actual result but rather in its intended result. I tend to disagree, depending on whether we really are talking about evaluating the action’s rightness as contrasted with, say, its praiseworthiness, but present purposes do not require us to engage this further argument.

¹⁰ To call an act right is ambiguous. One might be saying the act is required or that it is permitted. The former sense is more relevant here. I use ‘right’ and ‘required’ interchangeably in what follows.

¹¹ This may be why Prichard saw the connection between a sense of rightness and a reason to be moral as “absolutely immediate.” If anything intrudes between the two, one is no longer doing the right thing for the right reason.

¹² Prudence involves acting in one’s best interest simpliciter rather than acting in one’s best interest *because* it is in one’s best interest. Otherwise, if we interpret prudence in the latter sense, prudence and morality exhibit a particularly uninteresting kind of incompatibility; the real issue about the overlap between moral and prudent behavior will inevitably resurface, cast in other terms.

¹³ One could see Prichard as rejecting rationalism in ethics in the same way Michael Oakeshott rejects rationalism in politics. That is, we understand and appreciate ethical traditions only from the inside, by living within them and by knowing their history. It is hubris to criticize traditions on the grounds that they fail to serve purposes we

think ought to be served, or that they do not serve their purposes as well as imaginable alternatives. Such criticism is from the outside in, which is not a legitimate critical perspective. Instead, one must get inside the institution and experience the duties it imposes face to face and case by case. See the title essay in Oakeshott (1991). This theme also runs through the work of Alasdair MacIntyre. (The thesis that modern moral concepts are holdovers from earlier traditions, in which they had a significance that has since been lost, finds one of its earliest and most concise expressions in G.E.M. Anscombe [1958, 1, 5-8]).

Now, there is merit in the Anscombe-MacIntyre-Oakeshott line of argument. Nevertheless, moral philosophy is itself a body of traditions and practices. Distancing oneself from the practice of criticizing ethical traditions and viewing that practice with a critical eye amounts to taking an outside-in approach to philosophy's core tradition. Thus, to indulge in such criticism is tacitly to endorse outside-in criticism: criticizing philosophy from outside in by pointing out that philosophy partakes of outside-in criticism. A more telling critique would say something interesting about how to distinguish between the use and misuse of outside-in criticism.

¹⁴ Although Prichard's article does not say what triggers our intuitions, those who worked within the intuitionist tradition had a great deal to say about it. The point, though, is not that nothing can be said, but rather that something needs to be said. And when we begin to say what warrants us in intuiting that X is wrong, we begin to leave Prichard's brand of intuitionism behind.

¹⁵ With more ordinary intuitions, the answer might be experience. That is, we may have learned from experience to trust that sort of intuition. ("No, I do not want to get into that person's car. I see no reason not to, but *something* is telling me not to.") Still, the lesson of experience will not be simply that we should trust intuition, but that we have *reason* to trust intuition—doing so is for our own good, and we have a history of regretting the consequences of failing to do so. (So, I am not intending this as a concession to Prichard.) I owe the thought to Paul Bloomfield.

¹⁶ This is one of intuitionism's core insights. Another is that, in forming moral judgments, we draw upon tacit knowledge, some of which we cannot fully articulate. Similarly, a wine taster may have an astonishing ability to discern when and where the grapes came from, yet the information she finds in the wine's taste may be too subtle to put in words. These two ideas—that our knowledge is fundamentally of particulars rather than universals and that much of what we know is incorrigibly inarticulate—are also central tenets of the moral antitheory movement. (See also footnote #2.) A reading of Prichard thus is a useful introduction to the antitheory literature.

¹⁷ We speak here primarily of determining law in an epistemic sense, but in Hart's theory, secondary rules also determine the law in an ontological sense. For a discussion of the different senses in which secondary rules determine the law, see "Negative and Positive Positivism" in Jules Coleman (1988).

¹⁸ I think the most defensible reading of the principle of utility would say not. Recognition rules are not ultimate rules of conduct; primary rules are not mere rules of thumb. Primary rules do not defer to the "ultimate" rules in cases of conflict. Again, consider the legal analogy. In a situation where obeying the speed limit somehow interferes with reading the signs, the primary rule is still what binds. The speed limit does not give way to a "higher" law bidding us to read the signs. Likewise, in ethics, if we recognize that, in the world we actually live in, following the rule "keep promises no matter what" has better consequences than following alternative rules like "keep promises if and only if doing so maximizes utility," then the principle of utility (qua recognition rule) picks out "keep promises no matter what" as being among morality's rules of conduct.

¹⁹ I acknowledge that there are broader conceptions of deontology than this, revolving around a more general idea that being moral is a matter of having reverence for the moral law.

²⁰ Samuel Scheffler (1982) defends a "hybrid" theory, which departs from act-utilitarianism by holding that maximizing utility is sufficient but not necessary for an act's morality.

²¹ The idea that to endorse something as moral is to endorse it from a plural perspective does not beg the question against egoist or otherwise individualist moral theories. It is within the realm of possibility that I might endorse a refined enough egoism not only for me but for you too, and not only because it is best for me but because it is best for you too. In short, I can from a plural perspective endorse that we each tend our own gardens. Again, to belabor a distinction that I know from experience is not obvious, recognition rules are not rules of conduct. Rules of conduct are what we look *at*. Recognition rules are what we look *with*.

²² See especially David Gauthier (1986). A Prisoner's Dilemma is a game in which individuals make separate decisions about whether to contribute to a cooperative venture. In essence, the problem is, if an individual contributes, the benefits will be dispersed in such a way that the marginal benefit per unit of contribution is less than one unit to the contributor but more than one unit to the group. See Schmitz (1991, 105). In an obvious way, people are better off as a group if they contribute but better off as individuals if they do not.

²³ I thank Philip Pettit for pressing this point in conversation.

²⁴ Geoffrey Sayre-McCord (1994) directly addresses this issue. This essay originally went to press around the same time as Geoff's, but the ideas in this paragraph are similar enough to Geoff's to make me wonder whether I got them from him. In any case, I thank Geoff for his influence on me over three decades of conversation.

²⁵ The scope of my plural perspective need not coincide with the scope of yours, which is one reason why we can disagree about what is moral. Discussing our differences often helps us extend our perspectives in ways that bring them into alignment, though, so disagreement that can be traced to differences in perspectival scope need not be intractable. If you convince Kate that her *we*-perspective until now has failed to encompass the interests of members of other races, for example, then she broadens her perspective accordingly. Or if she willfully refuses to do so, then her kind of *we*-perspective reveals itself to be distinctly unlike the perspective that I am attributing to people who earnestly ask the "why be moral" question.

²⁶ Note that this is a characterization of the perspective from which we formulate recognition rules. Whether *being* moral necessarily involves taking a plural perspective is a separate question. (Do morality's rules of conduct include an injunction to take a plural perspective? No, they do not, any more than the rules of the road include an injunction to read the signs.)

²⁷ I suppose that if all we knew was that X is fun and does no harm, we might treat that as grounds for endorsing *x* as morally permissible. Given additional information that *x* = Disneyland, we might not retract our endorsement. We would not have deemed *x* morally required, though, even before learning that *x* = Disneyland.

²⁸ Partly for this reason, a method of seeking "reflective equilibrium" seems practically unavoidable in moral theorizing. I do not regard seeking reflective equilibrium as a meta-principle or a moral theory or even a formal method, really. It is simply a matter of remaining responsive to what is pretheoretical. With respect to a given subject matter, we assess candidate action-guides (for example) by the light of our recognition rules. In turn, we assess our recognition rules by asking whether the action-guides they yield are plausibly responsive to what pretheoretically seems important about that particular subject matter. I thank Thomas Pogge for discussion.

²⁹ At best, Prichard says, the element of truth in the view that rightness is tied to goodness is that unless we recognize that an act will give rise to some good, we would not recognize that we ought to do it. But, he adds, this does not mean pain's badness is the reason not to inflict it (1968a, 5). This seems like a massive concession, but Prichard mentions it in passing as if it were unimportant. In a footnote, Prichard claims that if pain's badness grounded the wrongness of inflicting it, then inflicting pain on oneself would be as wrong as inflicting it on others. Not so. If two rules of conduct (Do not inflict pain on others; do not inflict pain on yourself) are grounded in pain's badness, that implies nothing about whether the two rules of conduct are equally stringent.

³⁰ Although Rawls's official view is that the right is prior to the good (1971, 31), his theory's recognition rule is paradigmatically teleological. We recognize a principle as just by asking whether people behind a veil of ignorance would see the advantages of a basic structure informed by the principle. "The evaluation of principles must proceed in terms of the general consequences of their public recognition and universal application" (1971, 138). This is not the sort of statement one expects to find at the core of a theory in which the right is prior to the good, but I think what Rawls wants to say is that the right is prior to the good at the action-guiding level.

³¹ Even so, we should not concede to Prichard that pointing out reasons to keep promises is *always* irrelevant to someone who believes promise-keeping is required. Even someone who believes promise keeping is required might be unable to articulate good reasons to keep promises and might learn something from discussion.