4
Attitude, Affect, and Authority

Introduction

Moral acceptance is noncognitive. Specifically, moral acceptance centrally involves a certain kind of affect, what Scanlon (1998) describes as a desire in the directed attention sense. In accepting a moral sentence that he understands, a competent speaker reconfigures his affective sensibility so as to render salient, in a phenomenologically vivid manner, the moral reasons apparently available in the circumstance, as he understands it. In accepting a moral sentence that he understands, a competent speaker quite literally decides how he feels about things. It is the structure of a person's moral consciousness and not some further fact that constitutes the relevant kind of affect. The relevant affect is nothing over and above the tendency for certain features of the circumstance to become salient in perception, thought, and imagination, and for these to present a certain complex normative appearance. Specifically, certain features of the circumstance become salient and appear to be reasons for acting, while other features potentially cease to be salient and can appear to be

outweighed or even ruled out as reasons for doing otherwise, even if, in normal circumstances, they would count as such reasons. The salient features appear to be reasons that are not contingent upon our acceptance of them. Moreover, potentially distinct features of the circumstance become salient and appear to be reasons for accepting the moral sentence, and these reasons directly or indirectly involve grounding reasons, i.e. reasons that ground the deontic status of the relevant practical alternatives. These grounding reasons appear to be reasons not only for the speaker, but for everyone else as well. They appear to be sufficient for accepting that sentence on behalf of others. From this perspective, the competent speaker can seem justified in demanding that others accept the moral sentence and so come to respond affectively in the relevant manner. The effects centrally involved in moral acceptance are in this way essentially other regarding.

In uttering a moral sentence that he understands, a competent speaker conveys the relevant affect and implicitly demands that others come to respond affectively in the relevant manner. A moral utterance frames the perspective of its audience so as to induce the relevant affect. Notice that such framing effects are a hallmark of figurative language. Suppose that Bernice, in remarking about Edgar's *gravitas*, represents Edgar as a pear. Emma, in appreciating the aptness of Bernice's remark, cannot now but see Edgar as pear-shaped. Bernice's figure of speech has framed Emma's perspective so as to render salient, in a phenomenologically vivid manner, certain relevant dimensions of similarity between Edgar and pears. So, whereas the fictional content of Bernice's remark represents Edgar as a pear, its real content merely represents Edgar as pear-shaped. Moreover, the fictional content plays a role in this framing effect by focusing on the relevant features of the circumstance and by representing them in a certain qualitative light. So in accepting the aptness of this remark, Emma has a tendency to focus on certain features of the

circumstance in perception, thought, and imagination, and a tendency for these features to have a certain qualitative appearance. Similarly, in accepting a moral sentence that he understands, the moral proposition expressed by the accepted sentence frames the perspective of the competent speaker so as to render salient, in a phenomenologically vivid manner, the moral reasons apparently available in the given circumstance as he understand it. So, whereas the fictional content of a moral utterance is the moral proposition expressed, its real content is plausibly limited to the morally salient features of the circumstance. Moreover, the fictional content plays a role in this framing effect by focusing on the relevant features of the circumstance and representing them in a certain normative light. So in accepting a moral sentence, a competent speaker has a tendency to focus on certain features of the circumstance in perception, thought, and imagination, and there is also a tendency for these features to have a certain complex normative appearance. The moral proposition expressed by the accepted moral sentence functions as an apt moral trope that frames the perspective of a virtuous moral sensibility.

A moral utterance conveys how a speaker feels about things. It does so not by virtue of an expressivist semantics that determines a nonfactualist interpretation for it. In reasoning from the nonrepresentational function of moral utterance (in this instance, its framing effects) to the accepted moral sentence having a nonrepresentational content, the expressivist plausibly conflates distinct senses of 'represent.' Moral discourse may be fully representational, moral sentences may express propositions that attribute moral properties to things, but the acceptance of a moral sentence might not be belief in the moral proposition expressed but, rather, the relevant kind of affect. Thus, for example, in accepting that the rights people have over their own persons ground the permissibility of abortion, Edgar has a tendency to focus on a certain feature of the circumstance, the

embryo being essentially a part of the mother's body, and a tendency for this feature to have a certain normative appearance. In uttering this claim, Edgar conveys this affect and implicitly demands that others come to respond affectively in the relevant manner. Edgar conveys the relevant affect in part by conveying the relevant normative appearance. Moreover, Edgar conveys the relevant normative appearance by representing what it is an appearance of. In representing the embryo's being essentially part of the mother as the ground of her right to make an uncoerced decision to abort, Edgar makes public and so conveys the relevant normative appearance. The represented moral property is an objective correlative of the relevant affect. Moral propositions may constitute the fictional content of moral acceptance and utterance, but they do not constitute their real content (which is plausibly limited to representing the morally salient features of the given circumstance). Moral propositions, propositions that attribute moral properties to things, play a role in moral acceptance and utterance, not by being the objects of belief and assertion, but by being apt moral tropes that frame the perspective of a virtuous moral sensibility.

What makes it fictionally true that things instantiate moral properties? What makes an action good or just within the moral fiction? A schematic answer is available:

> Suppose that moral predicate 'F' denotes moral property p. It is fictionally true that x is F iff x instantiates nonmoral property p^* that would elicit the relevant affective response in a person with a virtuous moral sensibility.

Since different attributions of moral properties differently structure the apparent reasons available in a given circumstance, we can be sure that different moral properties are paired with different affects. However, as we discussed in the previous chapter, there is no noncircular way to accurately specify the affect paired with the moral property. The affect is a desire in the

directed attention sense, so accurately specifying that affect would involve accurately specifying the constituent normative appearance, and the only way of accurately specifying the constituent normative appearance is in terms of what it is an appearance of. However, the semantic indispensability of moral properties is no more problematic for a fictionalist interpretation of moral discourse than the semantic indispensability of metaphor is for a fictionalist interpretation of metaphorical discourse (see Hills, 1997, and Walton, 1993). Just as there is no noncircular way to accurately specify the relevant affective response, there is no noncircular way to accurately specify the kind of sensibility from which this response is elicited. It is the reasons apparently available from the perspective of a *virtuous* moral sensibility that constitute the relevant affect. (Different conceptions of moral virtue are possible and, if actually implemented in moral practice, would generate different and potentially competing moral fictions.) This is unobjectionable. Suppose that theology is a fiction. It would not be surprising that a range of affective responses were available only to the participants of a theological fiction—that it is only within the theological fiction that one could regard things as sacred or holy. Similarly, it should not be surprising that a range of affective responses are available only to the participants of a moral fiction—that it is only within the moral fiction that one could feel beneficent or just.

Within the moral fiction, there are facts about the existence and distribution of moral properties. Moreover, competent speakers accept sentences that attribute moral properties to things. It is natural, then, that within the moral fiction the acceptance of a moral sentence is belief about the attributed moral property—that it is fictionally true of a competent speaker that he believes the moral proposition expressed by the moral sentence he accepts. In accepting a moral sentence, a competent speaker does not so much believe the moral proposition expressed as he makes as if to believe that proposition, where

making as if to believe, in this context, is to be disposed to respond affectively in the relevant manner. Acceptance in moral inquiry functions less to represent moral reality than to transform moral character by enabling competent speakers to respond affectively the way a virtuous person would. The aim of moral inquiry is not moral truth, but moral transformation.

Could We Discover that Morality was a Fiction?

What would be the rational response to the discovery that morality is a fiction?

There are both general and more specific questions here. The general question is: What would the rational response be to the discovery that morality is a fiction where at issue is the fictional status of morality and not the specific character of the moral fiction? A specific question is: What would the rational response be to the discovery that morality is a fiction if it had the character described in the previous section? Whatever the answer to that question is, the verdict might not be general. It might be the specific character of the moral fiction and not its fictional status that prompts the relevant response. In this chapter we will consider both questions.

This presupposes that we could, in fact, discover that morality is a fiction, but a doubt might be registered about this. Competent speakers unhesitantly describe their acceptance of a moral sentence S as believing that S. Could it really be the case that our attributions of moral belief are simply an unwitting pretense? So before we proceed, we will consider the prior question: Could we discover that morality is a fiction?

In Chapter 3 I argued that there were limits to the hermeneutic ambitions of moral fictionalism. While moral fictionalism is not necessarily committed to revising any of our distinctively moral commitments, it is committed to revising certain sophis-

ticated epistemological and linguistic beliefs. In accepting a moral sentence, competent speakers take themselves to believe the moral proposition expressed; and in uttering a moral sentence, competent speakers take themselves to be asserting the moral proposition expressed. But if fictionalism is true, then competent speakers are systematically misconceiving what they are in fact doing when they accept and utter a moral sentence. According to hermeneutic moral fictionalism, moral acceptance is best understood as some attitude other than belief in a moral proposition, and moral utterance is best understood as some linguistic action other than the assertion of a moral proposition. We might be justified in accepting most of the central moral claims that we in fact accept, at least by the norms of acceptance internal to moral practice, but we might nevertheless be wrong in describing our acceptance of these claims as beliefs and our utterance of them as assertion. Notice, however, that the claim that the acceptance of a moral sentence is not belief in the moral proposition expressed is not itself a moral claim: rather, it is an epistemological claim about the nature of the attitude involved in accepting a moral claim. Similarly, the claim that the utterance of a moral sentence is not the assertion of the moral proposition expressed is not itself a moral claim: rather, it is a linguistic claim about the nature of the action performed by moral utterance. Competent speakers may be right about the moral claims that they in fact accept but may systematically misunderstand what they are in fact doing when they accept and utter such claims.

While hermeneutic moral fictionalism is not necessarily committed to any distinctively moral error, nevertheless, the fact that competent speakers so badly misconceive what they are in fact doing in accepting and uttering moral sentences might strike some as incredible. Don't we know what we are doing when we accept and utter moral sentences? (This objection is not confined to the special case of moral fictionalism. Indeed, Stanley, 2001, presses this objection against any form of hermeneutic

fictionalism.) The challenge is to make it intelligible that competent speakers mistake what they are doing in accepting and uttering moral sentences without making the mistake impossible to discover. Indeed, there are several reasons why we might intelligibly be the unwitting participants of a moral fiction.

First, our attitudes and actions are not always fully transparent to us. Thus, it is plausible, for example, that a friend could be in a better position than you are to know that you are envious of a colleague and that a pattern of behavior that you have engaged in is an expression of that envy. Indeed, if your friend were to confront you and baldly assert that you are envious, you might initially deny it. However, if he were then to discuss the history of your relationship with your colleague and your behavior as seen from without, you might come, over time, to recognize what you initially denied. If this is indeed possible, then it is possible as well that an alien ethnographer could be in a better position than a native speaker to discover that in accepting moral sentences competent speakers do not believe the moral propositions expressed, and that in uttering moral sentences they are not asserting the propositions expressed. If the alien ethnographer were to confront a native speaker and baldly assert that his moral practice is fictionalist, the speaker might initially deny it. However, if the alien ethnographer were carefully to review the linguistic and ethnographic evidence with an open minded member of the linguistic community, such a speaker might come, over time, to believe that he was wrong all along about moral acceptance and utterance—that moral acceptance is not, in fact, belief in a moral proposition and moral utterance is not, in fact, the assertion of a moral proposition.

Second, one could not discover that morality was a fiction merely by reflecting on the content of moral vocabulary. Consider two possible worlds, w and w^*. Both are near duplicates, the population of each speaks a moral language with identical moral vocabulary with the same representational content. However, in

w, when a competent speaker accepts a moral sentence he believes the moral proposition expressed, and when he utters a moral sentence he asserts that proposition. In contrast, in w^*, when a competent speaker accepts a moral sentences he does not believe the moral proposition expressed but adopts some other attitude, and when he utters a moral sentence he does not assert that proposition but performs the distinct linguistic action of quasi-assertion. If a speaker were to determine whether he was an inhabitant of w or w^*, he could not do so merely by reflecting on the content of moral vocabulary. By hypothesis, moral content is invariant across w and w^*. Semantic knowledge would not determine his location in modal space.

Third, it is not just that reflection on moral content would not reveal that morality was a fiction; in addition, reflection on moral content might tend to conceal that morality was a fiction if in fact it was. If, in reflecting on the contents of the moral sentences that he accepts, a competent speaker recognizes the representational nature of that content, he might naturally take himself to believe the propositions expressed and to be asserting them when making moral utterances. If moral content were representational, then in accepting and uttering moral sentences a competent speaker might naturally take himself to be representing the moral facts to himself and others. The speaker's mistake might be encouraged by the systematic ambiguity in our representational idiom. Sometimes by 'representing o as F' we mean that the proposition that o is F is being put forward as true. Sometimes by 'representing o as F' we mean that the proposition that o is F is expressed whether or not that proposition is being put forward as true. In the former sense a representation is being put forward as true; in the latter sense the content of a representation is being specified whether or not that representation is being put forward as true. However, if a speaker were unclear about the distinction, then he might mistake the purported representation of putative moral facts for the successful

representation of such facts, and so mistake a noncognitive moral fiction for the cognition of the moral facts.

Moreover, there is a sense in which the cognitive appearance of moral discourse is both intelligible and predictable, given a fictionalist interpretation. After all, according to moral fictionalism, when a competent speaker accepts a moral sentence, he does not so much believe as makes as if to believe the moral proposition expressed. Making as if to believe, in this context, essentially involves adopting the relevant kind of affective response. Mistaking making as if to believe for belief is facilitated by the fact that attributions of moral belief are true within the moral fiction. The conditions that make a moral claim pretense-worthy for a competent speaker, if they obtain, also make the ascription of moral belief to that speaker pretense-worthy. That people believe the moral claims that they accept would be part of the extended moral fiction. Moreover, if attributions of moral belief are fictionally true, and moral pretense is unwitting, then it would be easy to mistake the fictional truth of such attributions for genuine truth, and so mistake a noncognitive moral fiction for the cognition of the moral facts.

What If Morality Were a Fiction?

Suppose we discovered that morality was a fiction. What would the rational response to the discovery be? Two observations are relevant here.

First, while a moral claim might be acceptable, given the norms internal to moral practice, nevertheless the acceptance of that claim might be illegitimate or unjustified by some relevant norm external to that practice. Compare: That Mercury rising has an unsettling effect on a person's psychology might be acceptable by the norms internal to astrology, but that does not mean that the claim is acceptable full stop. Suppose we reject

astrology; then the claim is not acceptable given the norms external to astrology that we accept and regard as authoritative. So it is one thing for a claim to be acceptable by the norms internal to a practice and another for it to be genuinely acceptable from the perspective of our overall practice.

Second, in Chapter 3 I remarked that it was unobvious what the commitments embodied in our use of moral language have to do with the metaphysics of morals. After all, the metaphysical commitments of a person as embodied in his use of language is one thing, and reality quite another. I argued that moral realism and its alternatives are better understood as epistemological postures or stances that are articulated in terms of the commitments involved in moral discourse. This point is presently important because moral fictionalism is consistent with the existence of moral facts. A competent speaker's accepting a moral sentence might consist in a noncognitive attitude, but this is nonetheless perfectly consistent with the existence of moral facts.

These two observations combine in an obvious way. Suppose the norms governing moral acceptance were noncognitive. A moral claim might be acceptable by the norms internal to moral practice, but might not itself be acceptable because, let us suppose, it did not correctly represent the moral facts. Even though a moral claim is acceptable by the norms internal to moral practice, it might not be genuinely acceptable by norms appropriate to the cognition of the moral facts if indeed there were any. Moreover, if there were, then, while the acceptance of a moral claim might be justified by the norms internal to moral practice, the acceptance of that claim might be unjustified nonetheless.

This possibility, if actual, would constitute not only an epistemic difficulty, but a normative difficulty as well.

In accepting a moral sentence, a competent speaker accepts as well what reason is thereby provided. The reason involved in

accepting a moral sentence (and conveyed in uttering that sentence) differs importantly from other reasons that the speaker might have. Moral reasons seem to have an authority that nonmoral reasons lack. While I did not give anything like a full account of the authority of morality in Chapter 1, I did describe the role it plays in moral discourse and in the cognitive psychology of competent speakers. One important feature of the authority of morality is the precedence that moral reasons take over nonmoral reasons. Specifically, in accepting a moral sentence, a competent speaker accepts as well a reason to act or to refrain from acting in a given circumstance that potentially overrides or cancels any conflicting nonmoral reasons available in that circumstance. So moral reasons differ from nonmoral reasons in that the former possesses an authority that the latter lacks—a fact that is manifest in the precedence that moral reasons take over nonmoral reasons. The authority of morality is manifest in other ways as well. Not only does a competent speaker, in accepting a moral sentence that he understands, accept a reason that takes precedence over nonmoral reasons, but he also accepts a reason that is not contingent upon his acceptance of it, for which there are grounds not only for him but for everyone to accept, where such grounds potentially justify demanding that others accept that sentence.

Suppose that, in accepting a moral sentence by the noncognitive norms of acceptance internal to moral practice, a competent speaker accepts as well a genuine reason. There nevertheless remains the question whether the accepted reason is a moral reason with the requisite authority. Suppose that a moral reason is a moral fact, or, at the very least, that the features of a given circumstance that count as a moral reason constitute a moral fact. Then in accepting a moral sentence by noncognitive standards of acceptance, a competent speaker would be mistaking a nonmoral reason for a moral reason. Indeed, he would be endowing the nonmoral reason involved in the acceptance of a

moral sentence with the authority appropriate only to moral reasons. He would, for example, mistakenly treat the reason he accepts as having precedence over other reasons available in the given circumstance. Moreover, he would fail to recognize when the nonmoral reason that he accepts is overridden or cancelled by a competing moral reason available in the given circumstance. So not only would a competent speaker be misconceiving the nature of the reason that he accepts in accepting a moral claim, but he might be accepting no legitimate reason at all.

Morality in Plato's Cave

This is the situation that MacIntyre (1981: chapter 2) describes the Bloomsbury group as being in. According to MacIntyre, there was a radical discrepancy between meaning and use in the moral discourse of the Bloomsbury group. Given the meaning of moral vocabulary, the acceptance of a moral sentence seemed to involve the acceptance of a reason with the requisite authority. However, given the use of moral vocabulary, the acceptance of a moral sentence actually involved the acceptance of a nonmoral reason that lacked this authority. Specifically, their acceptance of a moral sentence was not governed by norms appropriate to the cognition of the moral proposition expressed (where the nature of the represented fact was the alleged ground of the authority of the accepted reason); rather, their acceptance of a moral sentence was governed by noncognitive norms. (For present purposes, what really happened in WC1 is irrelevant. What is presently important is that MacIntyre's account might be true, if not of the Bloomsbury group as they actually were, then of how they nearly might have been.)

According to G.E. Moore (1903), moral properties are nonnatural properties that can be intuited by persons with the appropriate moral sensibilities. Not only was Moore a nonnaturalist

and an intuitionist, but he was a consequentialist as well: an action is right in a given circumstances just in case it produces more good consequences than any alternative action that is open in that circumstance. Moreover, Moore held a specific conception of the good: those things that instantiate nonnatural goodness to the greatest degree are personal intercourse and the beautiful.

According to MacIntyre, the Bloomsbury group embraced Moore's moral philosophy not on the strength of Moore's arguments, but rather because Moore's moral philosophy reflected the values they antecedently accepted. Their preferred form of life privileged the values of personal intercourse and the beautiful, just as Moore prescribed. Not only did the Bloomsbury group share Moore's conception of the good, but they were also disposed towards consequentialist forms of moral reasoning. Keynes reports that discussions of value involved the explicit ranking of states of affairs. He cites the following questions put forward for discussion:

> If A was in love with B and believed that B reciprocated his feelings, whereas in fact B did not, but was in love with C, the state of affairs was certainly not as good as it would have been if A had been right, but was it worse or better than it would become if A discovered his mistake?
>
> If A was in love with B under a misapprehension as to B's qualities, was this better or worse than A's not being in love at all? (MacIntyre, 1981: 16–17)

Moreover, such questions were resolved by appeal to intuition. The Bloomsbury group would focus on the target state of affairs and attempt to discern as best they could the presence and degree of nonnatural goodness instantiated in the target state of affairs. If there was disagreement, then either the disputants were focusing on different subject matters, or the moral sensibility of one was better placed to discern the presence and degree of nonnatural goodness than that of the other.

So it would seem that the moral practice of the Bloomsbury group was thus explicitly intuitionist—at least on the surface:

> But, of course, as Keynes tells us, what was really happening was something quite other: 'In practice, victory was with those who could speak with the greatest appearance of clear, undoubting conviction and could best use the accents of infallibility' and Keynes goes on to describe the effectiveness of Moore's gasps of incredulity and head-shaking, or Strachey's grim silences and Lowes Dickenson's shrugs. (MacIntyre, 1981: 17)

If Keynes is to be believed, this is plainly the kind of manipulative noncognitivism for which Moore's student, Stevenson, has been criticized. On this view, the Bloomsbury group, in accepting an attribution of goodness, was not in fact tracking the presence and degree of nonnatural goodness. Rather, they were engaged in an unwitting pretense in which things have, in addition to their natural properties, certain nonnatural properties that supervene on them and that can be intuited by persons with the appropriate moral sensibilities.

Suppose that Edgar was a minor member of the Bloomsbury group. Being a faithful student of the *Principia*, Edgar understands the sentence

> A's being in love with B under a misapprehension of B's qualities would be a better state of affairs than A's never being in love

as Moore does—as representing a difference in the degree of nonnatural goodness instantiated by two potential states of affairs. Moreover, Edgar accepts this sentence. In what does Edgar's acceptance of this sentence consist? From within, Edgar's coming to accept this sentence occurred just as Moore describes: Edgar contemplates A's being in love under a misapprehension and A's never being in love and intuits that the former state of affairs instantiates nonnatural goodness to a degree greater than

the latter. However, Edgar's intuition can be explained independently of the actual intuition of any nonnatural properties. Edgar accepts that A's being in love under a misapprehension is better that A's never being in love because, given his sensibility, Edgar approves of the former state of affairs more than he does the latter.

So there were two complementary principles governing this pretense. First, Moore's *Principia*, regardless of the truth of its doctrines, functioned as the master fiction of the moral pretense. In accepting and uttering moral sentences, the Bloomsbury group were acting as if the *Principia* doctrines correctly described the moral facts. If we confine our attention to attributions of nonnatural goodness, a rough statement of one principle governing the moral fiction would be:

> It is fictionally true that x instantiates nonnatural goodness iff according to the *Principia*, x instantiates nonnatural goodness.

Not only did the *Principia* prescribe, at least in general outline, which attributions of nonnatural goodness were fictionally true, it also prescribed an independent procedure for determining which individual attributions were fictionally true. According to the *Principia*, attributions of nonnatural goodness are accepted on the basis of intuition. What makes it fictionally true that a person is intuiting instances of nonnatural goodness is that, given their sensibility (a sensibility shaped by Moorean doctrine), they approve of that thing. Within the moral fiction, while nonnatural properties are distinct from natural properties, they nevertheless supervene on them. If a thing instantiates the (fictionally subvenient) natural properties, thereby endowing it with the tendency to elicit approval from persons with the appropriate sensibility, then it is fictionally true that it instantiates nonnatural goodness:

> It is fictionally true that x instantiates nonnatural goodness iff x instantiates natural properties that would elicit the relevant emotional attitude in a person with the appropriate sensibility.

Putting these principles together, we get a principle connecting the emotional attitudes of the Bloomsbury group with the content of the *Principia*:

> x instantiates natural properties that would elicit the relevant emotional attitude in a person with the appropriate sensibility iff according to the *Principia*, x instantiates nonnatural goodness

In this way, the *Principia* both controlled and gave expression to the emotional attitudes of the Bloomsbury group.

But why did the Bloomsbury group express their amorous and aesthetic ends in the language of morality? Why this masquerade? MacIntyre suggests that, in rejecting the moral culture of the late nineteenth century in favor of a form of life that privileged the values of personal intercourse and the beautiful, what the Bloomsbury group lacked was a means justifying their preferences to others. Given the practical conflict between Victorian morality and their preferred form of life, the Bloomsbury group needed a means of rejecting at least those claims of Victorian morality that were incompatible with their ends; for, if the claims of Victorian morality were genuine, and so had the requisite authority, then, given the precedence of moral reasons, the reasons provided by these amorous and aesthetic ends were potentially overridden or even cancelled. In order to justify their rejection of Victorian morality and so pursue their preferred form of life, the Bloomsbury group needed to endow their ends with the authority of morality.

Moore's moral philosophy seemingly allowed them to do just that. The acceptance of Moore's moral philosophy was a means of reconceiving the nonmoral reasons provided by their

preferences as moral reasons with the requisite authority. So doing seemingly allowed the Bloomsbury group to justify their form of life to their Victorian critics. Far from being at odds with morality, the privileging of the aesthetic and the amorous was precisely what morality required—at least from the perspective of the Moorean fiction that they accepted. However, insofar as these amorous and aesthetic ends provided reasons for acting in a given circumstance, what reason they actually provided lacked the authority of morality. Apparent instances of nonnatural goodness were merely shadows cast by amorous and aesthetic ends held independently of morality.

If moral practice were in this way fictionalist, then a cognitive reconstruction of moral practice would be required. We would need to turn from the shadows cast on the Platonic cave and look into the light: the noncognitive norms governing moral acceptance would need to be replaced by norms appropriate to the cognition of moral facts. Such a replacement would be not only epistemically required, but normatively required as well. The noncognitive norms governing moral acceptance would need to be replaced not only to justifiably believe the moral proposition expressed but also to justifiably act on the accepted moral sentence. The normative difficulty, after all, was that, in accepting moral sentences by noncognitive norms, competent speakers were accepting nonmoral reasons that were potentially overridden or cancelled in the given circumstance and so were potentially accepting no legitimate reasons at all. In the situation that MacIntyre describes, practical rationality requires justified moral belief, and a cognitive reconstruction of moral practice is needed before competent speakers can justifiably believe any moral proposition. Given the normative difficulty envisioned by MacIntyre, moral fictionalism, if accepted as a correct description of the way moral acceptance actually functions, would naturally lead to a revised and reconceived moral practice.

Benign Moral Fictionalism?

Moral fictionalism, however, is not necessarily committed to the normative difficulty envisioned by MacIntyre.

The difficulty with there being moral facts that moral acceptance fails to track is an instance of a more general difficulty. If there were moral facts and the acceptance of a moral sentence was at variance with the norms appropriate to their cognition, then in accepting a moral sentence a competent speaker would potentially be mistaking a nonmoral reason for acting in the given circumstance for a moral reason with the requisite authority. The more general difficulty is mistaking nonmoral reasons for moral reasons—reasons with the requisite authority. The authority of morality is manifest in the functional role that moral acceptance plays in moral discourse and in the cognitive psychology of competent speakers. Specifically, in accepting a moral sentence S that he understands, not only does a competent speaker accept a reason that takes precedence over nonmoral reasons, that is not contingent upon his acceptance of it, for which there is a grounds not only for him but for everyone to accept, but the competent speaker in uttering S demands that everyone accept S. Moral realists maintain that cognition of the moral facts best explains and renders intelligible the authoritative role that moral acceptance plays. However, moral authority need not be understood as the moral realist understands it. It is at least conceivable that, even if there were no moral facts, there would nevertheless be a legitimate distinction between reasons that possess the authority characteristic of morality and those that lack it. If this distinction can be made without commitment to the moral facts, and if in accepting a moral sentence competent speakers accept a moral reason, then a fictionalist moral practice need not be the kind of cultural disaster that MacIntyre describes.

Let us begin with the notion of a reason. Even if one denied that there were moral facts, one might nevertheless claim that

there were facts about reasons. However, if one denied that there were moral facts one might deny as well that there were facts about reasons—perhaps the normative character of each presents similar obstacles to regarding both moral and normative discourse as being fully factual. Let us consider these options in turn.

Suppose there are no moral facts but there are facts about reasons. Nevertheless, a distinction might be drawn between moral and nonmoral reasons. Among the reasons that there are, some have the authority constitutive of being a moral reason. We typically convey moral reasons by means of moral utterance, but in so doing we are engaged in an act of quasi-assertion. We invoke a moral fiction in order to describe a particular kind of reason that is allegedly available in the given circumstance. Competent speakers convey that a feature of their circumstance has the normative significance that it does by invoking the metaphysical fiction of moral properties that supervene on those features and that ground their normative significance. So, while the fictional content of a moral utterance is a moral proposition, its real content represents a particular kind of reason—a reason with the requisite authority. There is a potential difficulty with this position, though perhaps not an insurmountable one. On the present view, there are facts about reasons and some of them are moral reasons. How is it that moral facts are not simply identified with the moral reasons? How are we to distinguish moral reasons from moral facts? If we cannot, then the present view collapses into a rationalist moral realism.

This difficulty might favor the second option. Suppose there are no moral facts and no facts about reasons either. Suppose, more specifically, that talk of reasons had a noncognitivist use, though not necessarily a nonfactualist one. How might this work? Recall Gibbard's (1990) strategy. Suppose that something's being a reason is understood as treating it as a reason. Suppose, moreover, that treating something as a reason is to accept a norm

that prescribes that it count in favor of something. Then there will be norms corresponding to reasons, and whatever can be expressed in terms of reasons can be expressed in terms of norms. Suppose a competent speaker accepts that a feature of his circumstance is a reason to perform an action. While the fictional content of that claim involves the representation of a reason, it conveys the speaker's acceptance of a system of norms that requires treating that feature as counting in favor of the relevant action. Among the reasons, so understood, a distinction might be drawn between moral and nonmoral reasons.

So there are two ways to distinguish moral and nonmoral reasons while denying the existence of a distinctively moral range of fact. One might accept that there are facts, about reasons and that some are distinctively moral, or one might deny not only that there are moral facts, but that there are facts about reasons as well (assigning, instead, a noncognitive use to our talk of reasons) and accept that some of the reasons that we recognize are distinctively moral. While the former option is a kind of weak noncognitivism, the latter option is a kind of strong noncognitivism. On the former option, a competent speaker in accepting a moral sentence accepts a moral reason where moral reasons are conceived as a kind of fact. While moral acceptance is not belief in a moral proposition, it is belief in a proposition that represents a kind of reason. On the latter option, a competent speaker in accepting a moral sentence accepts a moral reason where moral reasons are not conceived as a kind of fact. While moral acceptance is not belief in a moral proposition, neither is it belief in any other proposition. Mixed accounts are possible as well. Thus, on the present account, moral acceptance not only involves thoughts and perceptions that represent the morally salient facts about the relevant circumstance, but crucially involves a phenomenologically vivid sense of the moral reasons apparently available in the circumstance as it is understood to be. However, if minimalism is correct, these attitudes are not distinct: the

thoughts and perceptions involved in moral acceptance are events in a person's consciousness whose structure constitutes the relevant affect. For present purposes, however, it does not matter which of these options are embraced so long as the denial of moral fact can be combined with the claim that some reasons are distinctively moral reasons and that these are the reasons involved in moral acceptance. As long as this distinction can be marked without postulating moral facts, and noncognitive moral acceptance involves the acceptance of moral reasons, then we can avoid the normative difficulty that MacIntyre envisions.

Attitude and Authority

Suppose that authoritative reasons can be operationally characterized in terms of the functional role they play in moral discourse and in the cognitive psychology of competent speakers. Perhaps a suitably complex yet coherently integrated system of noncognitive attitudes could implement that functional role in a way that would render intelligible why moral reasons would exhibit that role. Indeed, Blackburn (1998) and Gibbard (1990) each give accounts of roughly this form. Blackburn vividly expresses this approach as follows:

> We should think in terms of a staircase of practical and emotional ascent. At the bottom are simple preferences, likes, and dislikes. More insistent is basic hostility to some kind of action or character or situation: a primitive aversion to it, or a disposition to be disgusted by it, or to hold it in contempt, or to be angered by it, or to avoid it. We can then ascend to reactions to such reactions. Suppose you become angry at someone's behaviour. I may become angry at you for being angry, and I may express this by saying that it is none of your business. Perhaps it was a private matter. At any rate, it is not a moral issue. Suppose on the other hand, I share your anger or feel "at one" with you for so reacting. It may stop there, but I may also feel strongly disposed to encourage

others to share the same anger. But then I am clearly treating the matter as one of public concern, something like a moral issue. I have come to regard the sentiment as legitimate. Going up another step, the sentiment may even be compulsory in my eyes, meaning that I become prepared to express hostility to those who do not themselves share it. Going up another level, I may also think that this hostility is compulsory, and be prepared to come into conflict with those who, while themselves concerned at what was done, tolerate those who do not care about it. I shall be regarding dissent as beyond the pale, unthinkable. This should all be seen as an ascending staircase, a spiral of emotional identifications and demands. The staircase gives us a scale between pure preference, on the one hand, and attitudes with all the flavor of ethical commitment, on the other. (Blackburn, 1998: 9)

Suppose that something's being a reason is a matter of treating it as a reason. Suppose, moreover, that treating something as reason is to accept a norm that prescribes that it count in favor of something. Then, there will be norms corresponding to reasons, and it might seem that whatever can be expressed in terms of reasons can be expressed in terms of norms understood as noncognitive attitudes. So instead of grounding reasons we might speak of higher-order attitudes. Suppose that, in accepting the wrongness of abortion, Bernice accepts a norm forbidding abortion if pregnant. Suppose, moreover, that she accepts higher-order norms that prescribe that she accept that norm, that everyone accept that norm whether or not they in fact accept it, and that she demand that others accept that norm. The authority of the demand might then be grounded in the higher-order attitudes that prescribe it. The general idea is that higher-order attitudes regulate which lower-order attitudes to accept and hence which features of the circumstance are to count as reasons. It is this regulative role in determining what counts as a reason that explains why the authority of the demand is grounded in the higher-order attitudes that prescribe it. Thus, the demand conveyed by Bernice's moral utterance would not be manipulative,

and accommodating that demand could be a correction of attitude.

It is doubtful, however, whether higher-order attitudes could be the grounds of authority. There are three closely related grounds for doubt.

Suppose that Edgar accepts a higher-order norm prescribing that he accepts a norm prescribing guilt if he frustrates the expectations of others. Suppose, however, that Bernice expects Edgar to take the blame for her wrongdoing. Edgar may be socially obliging, but he is nobody's patsy: he is not disposed to feel guilty for not taking the blame. Indeed, he is not disposed to feel guilty because he accepts a norm that forbids guilt in those circumstances. So Edgar accepts a higher-order norm that conflicts with a lower-order norm that he also accepts. How might this conflict between higher- and lower-order attitudes be resolved? Edgar might revise the lower-order norm forbidding guilt since it is inconsistent with a higher-order norm that he accepts. Indeed, this is part of the reason for thinking that higher-order attitudes could be the grounds of impersonal authority: the higher-order norms regulate which lower-order norms to accept and hence which features of the circumstance are to count as reasons. This is an illusion, however. Edgar might equally revise the higher-order norm. He may be obliging, and continue to be, but his confrontation with Bernice might reveal that obligingness has its limits, and he might revise his higher-order attitudes to reflect this. Everything else being equal, it is good when higher- and lower-order attitudes cohere, but when they conflict coherence can be achieved by revising either. But this undermines the idea that higher-order attitudes are authoritative: if higher-order attitudes regulate which lower-order attitudes to adopt, then coherence should be achieved only by revising lower-order attitudes but coherence may be achieved by revising higher-order attitudes as well. (Compare Watson's, 1975: 108–9, criticism of Frankfurt, 1971. Scanlon, 1998: 54–5, makes essentially the same

criticism of desire models of reasons that appeal to higher-order desires.)

There is a second related ground for doubt. Higher-order attitudes differ from lower-order attitudes. Specifically, they differ in their objects: While higher-order attitudes have lower-order attitudes as objects, lower-order attitudes do not. But how could attitudes of fundamentally the same kind differ in authority when the only relevant difference is an intrinsic difference in object? If there was a puzzle about how certain lower-order attitudes could be authoritative all by themselves, it is hard to see how this puzzle could be resolved by appealing to attitudes of fundamentally the same kind that differ only in object. This difference in object could not ground the authority that the latter allegedly has over the former. The point is easier to appreciate if, instead of the higher-order attitudes, that a person bears to his own attitudes we consider the higher-order attitudes that he bears to the attitudes of others. Suppose that Bernice is angry at Edgar. Suppose that Edgar feels that Bernice's anger is unwarranted. He might be angry at her for being angry. In this emotional conflict, it is wrong to think that Edgar's anger is authoritative just because it has Bernice's attitude as an object—Edgar, after all, may be being unreasonable. Higher-order attitudes are higher-order not in the sense that lower-order attitudes answer to them, but only in the sense that they have lower-order attitudes as objects.

There is a third related ground for doubt. There are two ways to describe the case where Edgar revises the lower-order attitude incompatible with the higher-order attitude that he accepts: the case might be described as a mere change in attitude, or it might be described as a correction of attitude. Suppose that the higher-order norm is authoritative. Then revising the lower-order norm is not merely a change in attitude, but a correction of attitude. However, if, as the noncognitivist conceives of it, the conflict is between attitudes that cannot be jointly satisfied where the only

relevant difference is an intrinsic difference in object, then there is no reason to think that the revision is anything other than a change in attitude. If the conflict were between accepting something as a reason for an attitude and accepting a reason that discounts that thing as a reason for that attitude, then the revision would be a correction of attitude. In the latter case, the relevant difference is not an intrinsic difference in object but a difference in the reasons for attitudes that intrinsically differ in object. The point is easier to appreciate if we consider the case where Edgar has all the relevant higher-order attitudes without regarding them as authoritative. Suppose that Edgar was raised to be socially obliging by domineering and psychologically adroit parents. Though Edgar cannot shake these attitudes, he can see no reason for acting on them. In these circumstances, if Edgar revises the lower-order attitude forbidding guilt, the revision would be a change of attitude and not a correction of attitude, since the revision is not a reflection of the reasons he accepts. Michael Smith describes a similar case in objecting to Blackburn (1998):

> [W]e can readily imagine someone who (say) has a desire that people keep their promises, and who shares many other people's anger at those who fail to keep their promises, and who feels disposed to encourage others to share the same anger too, and who feels disposed to be angry at those who don't share that anger, and yet who doesn't regard any of these sentiments as being in the least legitimate. We need simply to imagine someone who, in addition, regards all his various attitudes towards promising in much the same way as the unwilling addict regards his addiction. He might think, for example, that these attitudes were all simply caused in him by social forces, in much the same way as the ingestion of drugs caused the unwilling addict's desire to take drugs in him, and that no reasons can therefore be given in support of acting on the basis of these attitudes, much as the addict thinks that no reason can be given for acting on his desire to take drugs. (Smith, 2001: 111–12)

Thus, in representing a correction of attitude in terms that could only represent a change in attitude, the noncognitivist fails to account for the authority of moral utterance. The demands they convey may be prescribed by a coherently integrated system of higher-order attitudes, but, nonetheless, such demands may be manipulative, and accommodating them might be a mere change of attitude.

It is doubtful whether noncognitivism, even in its fictionalist guise, could account for moral authority in terms of a coherently integrated system of higher-order attitudes. Even if no such account could succeed, perhaps the authority of morality could be vindicated in some other way by the noncognitivist. Nothing I have said so far has suggested otherwise.

A noncognitivist that maintains that moral acceptance is desire in the directed attention sense, and maintains as well a minimalist account of that attitude must—on independent grounds—provide another account of moral authority. Recall that desire in the directed attention sense can be characterized in terms of its functional role: in terms of the tendency for certain features of the circumstance to become salient in perception, thought, and imagination and the tendency for these features to present a certain normative appearance. The tendency for certain features of the circumstance to become salient and the tendency for these features to present a certain normative appearance would be both intelligible and well explained by the acceptance of norms that prescribe that these features have that normative significance. Thus, accounts of moral authority in terms of a coherently integrated system of higher-order attitudes would be the basis of an explanation for the functional role of the relevant affect. The minimalist, however, denies that desire in the directed attention sense is an attitude whose nature can be specified independently of its functional role, and that can explain and render intelligible why this attitude has that role. Given this, a noncognitivist that maintains that moral acceptance is desire in

the directed attention sense and maintains as well a minimalist account of that attitude must account for moral authority in some other way.

Intransigence and Authority

Without speculating about how this might be done, let us consider a specific challenge to moral authority, given the present grounds for noncognitivism.

Recall that the argument from intransigence has three premises. First, public cognition is noncomplacent: if acceptance is cognitive and on behalf of others, then in the context of a disagreement in reasons, if a person is interested in accepting S on behalf of others, then he would thereby have a reason to re-examine the grounds of acceptance. Second, moral acceptance is authoritative: given its authority, moral acceptance is always acceptance on behalf of others. Third, moral acceptance is intransigent: in the context of a disagreement about reasons, a person interested in accepting S on behalf of others does not thereby have a reason to re-examine the grounds of moral acceptance.

To see how these claims constitute an argument for noncognitivism, first consider noncomplacency. If acceptance is cognitive and on behalf of others, then, in the context of a disagreement in reasons, if a person is interested in accepting S on behalf of others, he would have a reason to re-examine the grounds of acceptance. It follows that, if *moral* acceptance is cognitive and on behalf of others, then, in the context of a disagreement about reasons, if a person is interested in accepting a moral sentence on behalf of others, he would have a reason to re-examine the grounds of moral acceptance. This in turn entails that if, in the context of a disagreement about reasons, a person lacks a reason to re-examine the grounds of moral acceptance and moral acceptance

is on behalf of others, then the moral acceptance is noncognitive. Notice that the antecedent of this conditional is just the conjunction of intransigence and authority. Moral acceptance is intransigent: in normal circumstances we are under no obligation to reexamine the foundations of moral claims that we accept, even if they are disputed by otherwise rational and reasonable, informed, and interested people who accept reasons that, if genuine, would undermine them. Moreover, given its authority, moral acceptance is acceptance on behalf of others. This could be only so if moral acceptance were noncognitive.

Even if a cognitivist were to resist this argument by denying intransigence, there would, nevertheless, be an important normative lesson to be learned. Being unmoved to further inquiry is subject to epistemic criticism since it violates norms appropriate to moral belief. The envisioned cognitivist would claim that being unmoved to further inquiry is subject to epistemic criticism because moral acceptance is cognitive, and if moral acceptance is cognitive, then, in the context of a disagreement about reasons, an interested person has a reason to inquire further into the grounds of moral acceptance. Not only is being unmoved to further inquiry subject to epistemic criticism, but it is subject to moral criticism as well. The claim that if moral acceptance is cognitive then, in the context of a disagreement about reasons, an interested person has a reason to inquire further into the grounds of moral acceptance is a consequence of noncomplacency and authority. So the noncomplacent character of moral acceptance is, in part, a manifestation of the authority of a cognitive moral practice. So from the cognitivist's perspective, being unmoved to further inquiry could only be a symptom of an underlying moral debility since the authority of morality would thereby be undermined.

Can an independent question be raised—one not involving any cognitivist assumptions—about the compatibility of intransigence and authority? Perhaps.

An interpretation of the formula of humanity makes vivid the problem. Consider again Putnam's (1981: 165) reaction to fundamental disagreement:

> To be perfectly honest, there is in each of us something akin to *contempt*, not for the other's *mind*—for we each have the highest regard for each other's minds—nor for the other as a *person*—, for I have more respect for my colleague's honesty, integrity, kindness, etc., than I do for that of many people who agree with my ''liberal'' political views—but for a certain complex of emotions and judgments in the other.

What Putnam holds in something akin to contempt is Nozick's moral sensibility—a moral sensibility that privileges property rights over what Putnam regards as the compassionate treatment of the less well off. The question is whether something akin to contempt is the right attitude to adopt towards someone who in your view is lacking 'in a certain kind of sensitivity and perception.' Even if someone were lacking in this way, to treat him as an end is to treat him as capable, at least in principle, of acquiring the requisite sensitivity and perception. Moreover, to treat someone as an end is to allow for the possibility, however remote, that you yourself are lacking in this way. The difficulty of course is that contemptuousness is inconsistent with both these attitudes.

Treating someone as an end involves offering them reasons and treating them as capable of assessing those reasons. Conversely, it is to treat the reasons they offer as potentially genuine reasons that they are in a position to assess. It is this latter aspect of the formula of humanity that is presently relevant. What would it be, in the context of a disagreement about reasons, for Edgar to treat the reasons that Bernice offers as potentially genuine reasons that she is in a position to assess? It would involve, at a minimum, an openness to reflective doubt about his own grounds for the permissibility of abortion. This, in conjunction with an interest in accepting on behalf of others a

claim about the moral status of abortion, is sufficient to motivate further inquiry into the grounds of moral acceptance. Edgar would have a motive to inquire further into the grounds of moral acceptance to determine whether, in light of his discussion with Bernice, his reasons for acceptance are good reasons. He would also have a motive to inquire further to determine, in light of his discussion with Bernice, what, if anything, there is to Bernice's reasons for rejection. Bernice, after all, might be onto something that so far eludes Edgar. Adopting the end of further inquiry is not only to strive to be responsive to what reasons there are, but to treat Bernice as an end and not merely as a means.

Of course, there is latitude in the fulfilment of this end. Further inquiry is one end among many and a person's ends must be rationally ordered—perhaps Edgar has more compelling immediate concerns. Particular actions taken to fulfil this end are epistemically meritorious, while particular failures to fulfil this end merely lack epistemic merit and are not in any way epistemically blameworthy (though perhaps adopting the policy of never acting to fulfil this end would be). There is an additional reason why adopting the end of further inquiry should display this normative structure. In this context, striving to be responsive to what reasons there are is to strive for moral perfection, to better respond to authoritative reasons. So, not only are actions taken to fulfil this end epistemically meritorious, but such actions are morally meritorious as well. Similarly, not only do particular failures to fulfil this end lack epistemic merit, such failures lack moral merit as well. Moreover, just as particular failures are not epistemically blameworthy, such failures are not morally blameworthy (though perhaps adopting the policy of never inquiring further would be). It is not surprising, then, that striving to be responsive authoritative reasons should display this normative structure, a normative structure plausibly assigned to perfectionist duties.

Edgar, of course, need not revise his moral opinion. Treating Bernice as an end need not involve Edgar's abandoning the claim that abortion is morally permissible, only that he be prepared to bracket his full acceptance of that claim when inquiring further. Nor need it involve a partial normative accommodation of Bernice's position—perhaps on due reflection her position on abortion has nothing to recommend. What it does require is that Edgar adopt the end of further inquiry. In this context, striving to be responsive to what reasons there are is a manifestation of moral virtue.

Intransigence is plausibly incompatible with moral authority. In the context of a disagreement about reasons, for Edgar to treat the reasons that Bernice offers as potentially genuine reasons that she is in a position to assess would involve, at a minimum, an openness to reflective doubt about his own grounds for the permissibility of abortion. This, in conjunction with Edgar's interest in accepting on behalf of others a claim about the moral status of abortion, would be sufficient reason to inquire further into the grounds of moral acceptance. Notice that the requirement that Edgar be open to reflective doubt is a normative and not an epistemic requirement—it is part of what it is, in this context, to treat Bernice as an end. Moreover, the plausibility of this normative requirement is independent of the cognitive status of moral acceptance. It merely presupposes that there are reasons for acceptance, but this would be plausible even if moral acceptance were noncognitive. So, given an interpretation of the formula of humanity, intransigence can be shown to be incompatible with moral authority without making any cognitivist assumptions.

If, according to the norms that govern moral acceptance, moral intransigence is intelligible, then such norms are subject to normative criticism. It is arguable that, under certain conditions, the apparent intelligibility of moral intransigence would fail to appropriately value the humanity of others. If that is right,

then moral inquiry should be revised so as to become noncomplacent. The norms governing the acceptance of a moral sentence should be revised at least to the extent that, in the context of a disagreement about reasons, a competent speaker interested in the acceptability of S would be under a lax obligation to inquire further into the grounds of moral acceptance, at least if his disputant is otherwise rational and reasonable, informed, internally coherent, and similarly interested in the acceptability of S.

Renewed Moral Inquiry

In accepting a moral sentence that he understands, a competent speaker accepts as well what reason is thereby provided. Suppose a question is raised about the authority of these reasons given the norms that govern moral acceptance. Suppose, that is, that a questioned is raised about whether there are authoritative reasons that moral acceptance fails to track. Given the claims such reasons make on us, there would be reason to renew moral inquiry. Notice that the motivation is normative and not merely epistemic. If it were an open question whether there were authoritative reasons that moral acceptance fails to track, then moral inquiry, as it is actually conducted, would be potentially subject to normative criticism. There would thus be a normative and not merely epistemic reason to renew moral inquiry.

Moral fictionalism is consistent with the existence of the moral facts, and so it is logically possible that there are moral facts that moral acceptance fails to track. Moreover, if there were, there would be authoritative reasons that moral acceptance fails to track. However, this logical possibility is not sufficient to engender reflective doubt any more than the logical possibility that I am in the Matrix is. However, if moral intransigence were intelligible, then, as I have argued, a question could indeed be

raised about whether there are authoritative reasons that moral acceptance fails to track. If, on whatever grounds, it were an open question whether moral acceptance tracks authoritative reasons, there would be reason to renew moral inquiry.

What form would such renewed moral inquiry take?

In order for renewed moral inquiry to be noncomplacent, it would need to be self-consciously conducted as a public inquiry. After all, it is partly for the sake of others that that one should strive to be responsive to what reasons there are—that, in the context of a disagreement about reasons, one should adopt the end of further inquiry. It is nonaccidental that the results of such deliberation and, indeed, the deliberation itself can be presented in the medium of public language. Moral conversation, broadly conceived, is the proper medium of any such inquiry.

Renewed moral inquiry, so conceived, would not necessarily have as its aim the construction of a general and comprehensive moral theory. While a general and comprehensive theory for which there was noncollusive agreement among reasonable and rational people engaged in the joint endeavor of moral inquiry would be theoretically satisfying, it is unlikely to be achieved; nor should the aim of morally inquiry be anything so ambitious. The aim of renewed moral inquiry, rather, is to clarify our moral vocabulary and the grounds upon which we accept moral claims, to increase the coherence of the moral claims that we accept, and so on. This might result in a general and comprehensive theory, but then again it might not.

In clarifying moral vocabulary and the grounds on which we accept moral claims, and in increasing the overall coherence of our moral views, moral inquiry would rely on ordinary forms of public moral reasoning supplemented, where appropriate, by philosophical reflection. A philosophical theory of morality is by no means the grounds of moral inquiry. Rather, an adequate philosophical theory of morality would itself be grounded in the deliveries of moral inquiry.

As an illustration of this, consider the following: Suppose that renewed moral inquiry were undertaken not to redress the apparent intelligibility of moral intransigence, but to discover the moral facts if there are any. Having made a significant advance in the clarity and coherence of our moral views, even if large areas of disagreement remain, philosophical reflection on what has in fact been achieved might determine the cognitive status of that inquiry. In reflecting on the deliveries of the philosophically refined, public, moral reasoning we might be in a position to determine whether such reasoning was a form of moral cognition or whether the acceptance of moral claims on the basis of such reasoning remained noncognitive. In this way, a renewed moral inquiry might discover the moral facts if there are any.

If renewed moral inquiry essentially relies on ordinary forms of public moral reasoning, then what hope is there in its making any advance over actual moral debate? Two features of renewed moral inquiry are relevant here. First, moral inquiry is self-consciously conducted as a public inquiry—a cooperative venture whose end is acceptance on behalf of others. Moreover, it is a public inquiry that is motivated in a certain way. The point of engaging in renewed moral inquiry is for the participants in a moral fiction to assure themselves that there are not authoritative reasons that moral acceptance fails to track. This is an instance of what Rawls (1999: essay 22) has described as the burdens of reasons. The burdens of reasons are obstacles to a reasonable assessment of the moral reasons available in a given circumstance. The fact that moral acceptance is so burdened has normative implications for the conduct of moral discussion. After all, a fair-minded appreciation of the difficulties involved in assessing the authoritative reasons available in a given circumstance will affect how one interacts with others who disagree—even those who would disagree on fundamental matters. So renewed moral inquiry would be governed by reasonable

precepts not only because it is a cooperative venture, but because of the motivation for embarking on that venture in the first place. That renewed moral inquiry is subject to the precepts of reasonableness is what distinguishes it from all too familiar forms of moral combat.

What are the precepts of reasonableness that are a reasonable response to the burdens of reasons?

Since it is a cooperative venture that aims at reasonable consensus insofar as that is possible within the moral domain, it should be conducted in a manner conducive to that aim. While disagreements may reasonably persist, there is no place for intransigence here. In the face of reasonable disagreement, allowing for the possibility, however remote, that one's grounds for accepting a claim are not decisive is not only a reasonable precept of cooperative inquiry but also a rational response to the burdens of reasons.

Not only should moral inquiry be conducted in a manner conducive to that aim, but basic disagreement should be reasonably accommodated as well. Suppose that reasonable people engaged in the joint endeavor of moral inquiry assess an action in a given circumstance by fundamentally different principles that practically conflict. One way in which such basic moral disagreement might be reasonably accommodated is to build a partial consensus on the basis of what can be agreed to. Working from this partial normative accommodation, the parties should try, insofar as possible, to understand what, if anything, the other is responding to. Persistent disagreements, even if basic, should be approached from a perspective that emphasizes what is undoubtedly a large measure of agreement. Doing so not only lessens the temptation to see the other as a moral monster but also provides a starting point for reasonable discussion.

As an example of this, consider how the abortion debate has changed. In the early stages of that debate, no defender of a woman's right to an abortion would concede that abortion was

a bad thing, a fit object of regret even if justified. But that much is now conceded. Similarly, in the early stages of the debate, no prolifer would concede that abortion ought to be legally permitted even if morally forbidden. But that much is conceded, at least by many, no doubt in recognition that a decision to abort is a hard moral choice, combined with the conviction that people should make up their own minds about hard moral choices. This might suggest that the abortion debate has changed because the parties have partially accommodated the moral insights of one another. And this partial normative accommodation is plausibly a response to the reasons brought to bear by each side.

(I am unsure, however, whether this is the right account of the way in which the abortion debate has changed. The failure, early on, of the defender of abortion rights to concede publicly that an abortion is a fit object of regret might also plausibly be a rhetorical omission. Perhaps it was not conceded, not because it was rejected, but rather because to so concede would weaken the moral case for legalizing abortion. After all, it is hard to imagine a reasonably sensitive woman who actually had an abortion who did not at least concede the potential for legitimate regret, and so hard to believe that abortion being a fit object of regret was actually rejected. If that is right, then there was in fact no accommodation in this respect, and hence the change was not a response to the reasons brought to bear by pro-life advocates. Similarly, the concession by many that abortion ought to be legally permitted even if morally forbidden might merely be the counsel of despair prompted by the realization that the state-sanctioned mass slaughter of the innocent is a permanent feature of modern society. If that is right, then there was in fact no accommodation in this respect, and hence the change was not a response to the reasons brought to bear by pro-choice advocates. If this is the right account of the way in which the abortion debate has changed, then there was no tendency towards partial normative accommodation.)

These are not the only precepts of reasonableness. Rawls (1999) mentions, in addition, the reasonable expectation of disagreement and the crediting of one's interlocutor with good faith. There are undoubtedly others. I will not attempt to give anything like a comprehensive list, not least because such precepts should be determined, at least in part, from within the renewed moral inquiry. Just as other disciplines, whether psychophysics or econometrics, determine the methodology appropriate to their aim, so too should moral inquiry determine the methodology appropriate to its aim. So there is no saying in advance of such an inquiry what all of the precepts of reasonableness would be.

Moral inquiry here described is in some ways ideal. It is not ideal in the way that the Kingdom of Ends or Plato's Republic are; such an inquiry might be actually implemented in a way that the Kingdom of Ends or the Republic might not. Rather, it is ideal in that it is no substitute for practical deliberation. Due to inevitable practical exigencies, a decision to act or refrain from acting in a given circumstance might not wait on a consensus that may or may not emerge from moral inquiry. The participants of a renewed moral inquiry may have to act on moral reasons that they accept even if it is controversial whether such reasons have the authority that they take them to have. However, while moral inquiry is no substitute for practical deliberation, it is not entirely independent from it. Moral inquiry would depend, at least in part, on practical deliberation in that such deliberation is a potentially fruitful object of reflection for such an inquiry. However, just as importantly, practical deliberation would depend, at least in part, on moral inquiry in that the practical deliberation of the participants of such an inquiry would inevitably be informed by that inquiry. Moral inquiry would inform practical deliberation in at least two ways. First, moral inquiry would have the tendency to modify what one takes to be a morally relevant consideration in a given circumstance and so

would affect how one would react to finding oneself in that circumstance. Second, the virtues involved in a renewed moral inquiry would naturally generalize beyond this initial setting. So, while moral inquiry is no substitute for practical deliberation, neither is it independent from it.

Conclusion

Renewed moral inquiry might have a number of outcomes. At one end of the spectrum, a revision of moral practice is both theoretically and practically required. At the other end of the spectrum, no such revision is required. But there are interim possibilities. Perhaps moral practice would remain fictionalist even after a renewed moral inquiry but the character of the fiction would change. Perhaps, while benign moral fictionalism is a legitimate possibility, the moral fiction that competent speakers actually accept is not itself benign. Or perhaps the actual moral fiction is in many ways benign but renewed moral inquiry suggests ways in which that fiction could be improved. There is no telling, in advance, what such an inquiry would reveal.