

I

Moral Pyrrhonism and Noncognitivism

Introduction

WHAT is the nature of moral commitment?

Cognitivists claim that a person's moral commitments consist, at least in part, in their moral beliefs. So if Edgar is a moral vegetarian, then Edgar believes that it is wrong to kill animals for food. After all, we regularly describe our moral commitments as moral beliefs, and thus it is plausible that we hold such beliefs when we are so committed. Notice that Edgar's belief that it is wrong to kill animals for food is a moral belief, a belief with a moral proposition as its object. A morally committed person may hold moral beliefs, but that is not to say that a person's moral commitments involve no attitudes other than moral belief. So, for example, a cognitivist could concede that, in being a moral vegetarian, Edgar is averse to killing animals for food, where Edgar's aversion is a noncognitive attitude. Indeed, a cognitivist might plausibly claim that Edgar is averse to eating animals for food because he believes that it is wrong, and that Edgar's aversion is only an aspect of his moral commitment

insofar as there are moral grounds for it. Cognitivists need not deny that moral commitments involve noncognitive attitudes; cognitivists claim only that moral beliefs are indispensable to explaining a person's moral commitments.

In contrast, noncognitivists claim that a person's moral commitments are best explained by attitudes other than moral belief. While we may describe our moral commitments in terms of the moral beliefs we hold, the sense in which we hold such beliefs is best explained in terms of our noncognitive attitudes. Moral beliefs may be dispensable to explaining moral commitment, but that is not to say that a person's moral commitments involve no belief. So, for example, a noncognitivist could concede that, in being committed to the claim that it was wrong of her to lie to Edgar, Bernice believes that she lied to Edgar. It is just that this belief is not a moral belief. Bernice may believe, but Bernice believes no distinctively moral proposition. Noncognitivists need not deny that moral commitments involve cognitive attitudes; noncognitivists claim only that a person's moral commitments are best explained by attitudes other than belief in a moral proposition.

So cognitivists claim that a person's moral commitments are best explained by the moral beliefs he holds while noncognitivists claim that a person's moral commitments are best explained by attitudes other than moral belief. This difference is partly a difference in the nature of the attitudes involved in moral commitment (whether they are cognitive or noncognitive) and partly a difference in the content of these attitudes (whether the cognitive attitudes involved, if any, have moral propositions as their objects). This explanatory difference, however, is a manifestation of a more fundamental normative difference. Cognitivists and noncognitivists fundamentally differ about the norms governing moral commitment: Whereas cognitivists maintain that they are appropriate to belief in a moral proposition, noncognitivists deny this, emphasizing instead the nonrepresentational function of

moral commitment. (That the difference is fundamentally znormative is substantiated by the argument of Section 1.9.)

Acceptance

In this dispute between cognitivists and noncognitivists, how are we to determine the nature and content of the attitudes involved in moral commitment?

Moral commitments are expressed by moral sentences and incurred by competent speakers accepting moral sentences. As this linguistic observation is common ground between the cognitivist and the noncognitivist, one might begin with it and see if moral commitment can be neutrally characterized in terms of what a competent speaker is committed to in accepting a moral sentence. The idea is that one might then determine, in a non-question-begging manner, the nature and content of the attitudes involved in moral commitment, by determining the nature and content of the attitudes involved in accepting a moral sentence. ‘Acceptance’ here is a technical term and is explicitly stipulated to be neutral between cognitivist and noncognitivist understandings of the commitment incurred. It is useful to have a neutral term in order to pose the question whether the acceptance of a sentence in a given area is best understood as belief in the proposition expressed or as some other attitude. (Here I am following van Fraassen’s, 1980, usage.) The rationale for the procedure is twofold: First, the content of the accepted moral sentence is evidence about the contents of the relevant attitudes, since it plausibly constrains the contents of the attitudes involved in accepting that sentence. Second, the functional role of moral acceptance, the role that moral acceptance plays in moral discourse and in the cognitive psychology of competent speakers, is evidence about the nature of the relevant attitudes. Given this terminology, the dispute between the cognitivist and the noncognitivist becomes the following: given the norms governing

moral acceptance, is the acceptance of a moral sentence belief in a moral proposition expressed?

Acceptance is best understood in terms of its role in inquiry. Inquiry is not a solitary activity, at least not primarily. All successful forms of inquiry, such as physics, economics, literary theory, and so on, are *public* endeavors. Since successful forms of inquiry are public endeavors, their results are nonaccidentally presented in the medium of public language. Thus, for example, biology is a domain of inquiry, and associated with it is a region of discourse that involves a class of public language sentences couched in the distinctive vocabulary of that discipline. Morality itself constitutes a public domain of inquiry, albeit a distinctively practical one. When a person deliberates about his obligations to others, he inquires about his obligations to others and so engages in moral inquiry. Given the point of a person's deliberating about his obligations to others, it is nonaccidental that the results of such deliberation—and, indeed, the deliberation itself—can be presented in the medium of public language. Thus, associated with moral inquiry is a region of discourse that involves a class of public language sentences couched in a distinctive vocabulary. For the sake of simplicity, I will assume throughout that moral vocabulary consists solely in a class of moral predicates (such as 'good' and 'just'). This is a deliberate idealization: moral vocabulary also includes substantives (such as 'goodness' and 'justice') and, arguably at least, modal auxiliaries (such as 'must'). However, these important grammatical distinctions will be irrelevant to what follows.

The sentences of at least some regions of discourse express propositions relative to a context of utterance. Normally, the proposition expressed by a sentence is what is conveyed, among other things, in uttering that sentence in that context. Propositions are the primary bearers of truth and falsity. Sentences may be evaluated as true or false, but they inherit their truth-value from the proposition they express in the context of

utterance. Moreover, propositions are the objects of attitudes such as belief and assertion. This is not, of course, to claim that all attitudes are propositional. Edgar may be averse to eating meat, but the object of that attitude is not a proposition. I was careful to say that the sentences of *some* regions of discourse express propositions. Perhaps in other regions of discourse sentences fail to express propositions, but rather have a nonrepresentational content. So, while the sentences of at least some regions of discourse express propositions relative to a context of utterance, I will not assume that the sentences of every region of discourse are representational in the sense of expressing propositions that represent the putative subject matter of the given domain of inquiry.

We are assuming that moral vocabulary consists solely in a class of moral predicates. If moral predicates denote moral properties, then sentences that contain moral predicates express moral propositions—propositions that attribute moral properties to things. So conceived, the subject matter of moral inquiry would be the existence and distribution of moral properties. It is in this sense that we will understand the cognitivist's claim that the norms governing moral acceptance are appropriate to belief in the *moral* proposition expressed by the accepted moral sentence, and it is in this sense that we will understand the noncognitivist's denial. Instead, standard noncognitivists maintain that moral predicates do not denote moral properties, and hence that the sentences that contain them fail to express moral propositions; rather, moral predicates play a nonrepresentational role in moral discourse, and hence have a nonrepresentational content.

Public inquiry involves the production and the use of sentences from the associated region of discourse. There are two aspects of acceptance corresponding to these two features of public inquiry. Accepting a sentence from the region of discourse is both the object of inquiry and its grounds. Acceptance is the object of inquiry in the sense that it is a state that represents the end of

inquiry: in accepting a sentence, a person no longer takes himself to have a reason to investigate further, to continue to inquire whether or not to accept that sentence. Moreover, a person is justified in accepting a sentence if he possesses sufficient reason to inquire no further. Not only is acceptance the object of inquiry, but it is also the grounds of inquiry: in accepting a sentence *S*, not only does a person no longer take himself to have a reason to continue to inquire about *S*, but he also relies on his acceptance of *S* as grounds for further theoretical and practical inquiry. Acceptance is the grounds of inquiry in the sense that a person relies on the acceptance of the sentence in theoretical and practical reasoning and takes himself to be justified in so doing. These two aspects of acceptance are related: a person could be said to accept a sentence only if he was prepared to rely on it in theoretical and practical reasoning over a wide range of contexts. Moreover, a person would be justified in relying on *S* in theoretical and practical reasoning if he were justified in accepting *S*, if he had sufficient reason to inquire no further.

Acceptance is governed by norms and so is subject to criticism. Some of these norms are internal to the domain of inquiry; others are external to it. In a certain cultural and historical context, the acceptance of heliocentric astronomy may be criticized as impious, at least by one relevant norm or standard of impiety. However, the acceptance of heliocentric astronomy is not bad astronomy. The charge of impiety is not an astronomical criticism and relies on norms external to astronomical inquiry. Given the norms internal to astronomical inquiry, the acceptance of a heliocentric astronomy is not subject to criticism. Of course, being acceptable by internal norms need not guarantee genuine acceptability. The claim that Mercury rising has an unsettling effect on a person's psychology might be acceptable, if it is, by the norms of acceptance internal to astrology, but it is not acceptable by external norms that many of us accept and regard as authoritative.

Acceptance can be tentative or full (see Harman, 1986: 46–7). A person tentatively accepts a claim when, for example, he accepts a hypothesis in order to work out its implications. Thus, for example, while Edgar denies the axiom of choice, he might nevertheless tentatively accept that axiom in order to work out the implications of conjoining it with a standard set theory. Though Edgar tentatively accepts the axiom of choice, he does not fully accept it given his explicit denial. Notice that Edgar has a reason to accept the axiom of choice only while he has a reason to inquire after its implications for a standard set theory. Once he discovers some relevant set of implications, inquiry ends and there is no further need to accept the axiom. Tentative acceptance is not limited to supposition. To see this consider the following. Bernice only tentatively accepts General Relativity. Her acceptance of General Relativity is less than full acceptance in the sense that she is self-consciously prepared to give it up: she regards General Relativity as a very good approximation of the truth but an imperfect approximation nonetheless. Bernice has reason to accept General Relativity only while there is no significantly more accurate alternative. Tentative acceptance, while distinct from full acceptance, is a matter of degree. The degree of tentative acceptance depends on the extent to which a person relies on the acceptance of a sentence in theoretical and practical reasoning and the range of contexts in which a person does so rely. Thus, Bernice’s tentative acceptance of General Relativity is significantly more extensive than Edgar’s tentative acceptance of the axiom of choice. If, over time, and over a wide range of contexts, a person comes to rely sufficiently on the acceptance of a sentence in theoretical and practical reasoning, he may come to fully accept that sentence. Thus, the distinction between tentative and full acceptance is best understood as an approach to a limit.

In contrast to tentative acceptance, full acceptance ends inquiry. In fully accepting a sentence, the issue is closed, in the

sense that there is no reason to inquire further. A person is justified in fully accepting a sentence if, by the norms internal to inquiry or by authoritative norms external to it, he possesses sufficient reason to end inquiry. At issue in debate between cognitivists and noncognitivists is the nature and content of the attitudes involved in the full acceptance of moral sentences, and so only the norms governing full moral acceptance are relevant. To see this, consider the following. Bernice only tentatively accepts General Relativity. Suppose that her tentative acceptance falls short of belief. This would not establish that scientific acceptance consists wholly in attitudes other than belief in the accepted theory. Similarly, suppose that Edgar only tentatively accepts that it is wrong to kill animals for food. Suppose, moreover, that his tentative acceptance falls short of belief. This would not establish that moral acceptance consists wholly in attitudes other than belief in a moral proposition. It is the nature and content of the attitudes involved in the full acceptance of a moral sentence that is at issue in the debate between cognitivists and noncognitivists. While cognitivists maintain that full moral acceptance involves belief in the moral proposition expressed by the accepted moral sentence, noncognitivists deny this. Thus, only the norms governing full moral acceptance are relevant to determining the cognitive status of moral commitment. Henceforth, by 'acceptance' I will mean full acceptance.

The Argument from Intransigence

Why think that moral acceptance is noncognitive?

According to familiar internalist arguments for noncognitivism, there is a *functional* difference between moral acceptance and belief: accepting a moral sentence motivates a person to act in a way that belief does not. I will not argue in this way for noncognitivism; rather, I will argue that there is an *epistemic*

difference between moral acceptance and belief. I will not argue that moral acceptance has a functional property that belief lacks; rather, I will argue that belief has an epistemic property that moral acceptance lacks. This epistemic difference is brought out by the commitments incurred by reasonable and interested people engaged in a certain kind of disagreement, what Scanlon (1995) describes as a ‘disagreement about reasons.’ In cognitive inquiry, under certain conditions, people engaged in a disagreement about reasons have a motive that, in moral inquiry, under similar conditions, they would lack. This could only be so if moral acceptance were noncognitive. Or so I will argue.

Disagreements about Reasons

Suppose that Edgar and Bernice disagree about some sentence *S*. While Edgar accepts *S*, Bernice rejects *S*. Though she rejects *S*, Bernice strikes Edgar as an otherwise rational and reasonable human being—she can at least think and talk as well as Edgar. The mere fact of disagreement need not bother Edgar, for he might plausibly think that their disagreement derives from Bernice’s ignorance of the relevant evidence. Suppose, however, that Edgar engages Bernice in discussion and rules out this possibility: Edgar and Bernice share a common body of evidence. Not only is Bernice fully informed about the evidence that Edgar accepts, but she is also internally coherent in taking that evidence as a reason for rejecting *S*—just as Edgar is internally coherent in taking that evidence as a reason for accepting *S*. While they share a common body of evidence, they nonetheless disagree about its epistemic significance and are internally coherent in doing so. Given that each is internally coherent, each can offer what the other would regard as a question-begging argument for their acceptance or rejection of *S*. So both are otherwise rational and reasonable, fully informed, and can offer what the other would

regard as a question-begging argument for their acceptance or rejection of S.

Edgar and Bernice's positions conflict: they disagree about whether to accept or reject S. However, if we focus solely on the fact of conflict, we will miss something important about their disagreement. For Edgar and Bernice disagree not only about which sentence to accept in the given circumstance, they apparently disagree about what would count as a reason for acceptance in the given circumstance. Edgar and Bernice disagree about the norms governing acceptance: they implicitly accept distinct principles that determine what would count as a reason for acceptance in the given circumstance. This is manifest in a phenomenological difference between them. From Edgar's perspective, certain features of their circumstances are salient and have a certain normative appearance—they seem to be reasons for accepting S. From Bernice's perspective, potentially distinct features of their circumstances are salient and have a different normative appearance—they seem to be reasons for rejecting S. The world, as they commonly understand it, differs in the normative appearance it presents to each. Since Edgar and Bernice disagree not only about S but also what would count as a reason for accepting or rejecting S, their disagreement is a disagreement about reasons.

To get a better sense of this, consider how the traditional problem of induction can be recast as a disagreement about reasons. Suppose that Edgar is an inductivist: he believes that the regularities manifest in his experience are representative of the regularities that obtain in nature generally, even in the unobserved portions of nature. Suppose Edgar discovers that other people with distinct scientific traditions and cultures disagree with him. Bernice, for example, is a counterinductivist. Far from accepting the uniformity of nature, Bernice's beliefs about the unobservable are guided by a different principle, the non-uniformity of nature: Bernice believes that her experience is

positively misleading in the sense that an observed regularity is evidence that it does not obtain in nature generally—indeed, that it will fail in the very next instance. Though she denies the uniformity of nature, Bernice strikes Edgar as an otherwise rational and reasonable human being—she can at least think and talk as well as Edgar. The mere fact of disagreement need not bother Edgar, for he might plausibly think their disagreement derives from Bernice’s ignorance of the relevant evidence. Bernice might have been raised in an idiosyncratic, environmental niche where observed regularities are an unreliable guide to the regularities that obtain more generally in that environment, or she might have somehow failed to reflect adequately on what must have been a track record of predictive failure. Suppose, however, that Edgar engages Bernice in conversation and rules out these possibilities. It seems possible that Edgar may come, over time, to think that, just as his acceptance of the uniformity of nature is coherent given all the evidence, so is Bernice’s acceptance of the nonuniformity of nature. So both are otherwise rational and reasonable, fully informed, and can offer what the other would regard as a question-begging argument for their epistemic positions.

When Edgar and Bernice disagree about some theoretical sentence *S*, their epistemic positions conflict: they disagree about whether to accept or reject *S*. However, if we focus solely on the fact of conflict, we will miss something important about their disagreement. For Edgar and Bernice disagree not only about which sentence to accept in the given circumstance, they apparently disagree about what would count as a reason for acceptance in the given circumstance. Edgar accepts a principle according to which observable regularities count as a reason to accept that such regularities obtain in nature more generally, even in the unobserved parts of nature. Bernice, in contrast, accepts a principle according to which observable regularities count as a reason to *reject* that such regularities obtain in nature

more generally. So Edgar and Bernice not only disagree about *S*, they implicitly accept distinct principles that determine what would count as a reason for acceptance. This is manifest in a phenomenological difference between them. From Edgar's perspective, observable regularities are salient and appear to be a reason for believing that they obtain in nature more generally. From Bernice's perspective, observable regularities are salient and appear to be a reason for believing that they do not obtain in nature more generally. Since Edgar and Bernice disagree not only about *S* but also about what would count as a reason for accepting or rejecting *S*, their disagreement is a disagreement about reasons.

The disagreements described above, where the participants are otherwise rational and reasonable, fully informed, and can offer what the other would regard as a question-begging argument for their acceptance or rejection of *S*, are highly idealized. Indeed, so described, no such disagreements ever occur. No two people ever share precisely the same information, and it is impossible to say in advance of inquiry what information will be relevant to the acceptance or rejection of *S*. So no actual disagreement involves full information in the way described. These idealized cases of disagreement are nonetheless useful in dramatizing what is at issue in disagreement about reasons. What is at issue is not only whether to accept or reject *S*, but what would count as a reason to accept or reject *S*. What is at issue is the correct relevant principle that determines what counts as a reason for accepting or rejecting *S* in the given circumstance:

Disagreement about Reasons

In a disagreement about reasons, the disputants not only disagree about whether to accept or reject some sentence *S*, they disagree about what would count as a reason to accept or reject *S* in the given circumstance. Specifically, in a disagreement about reasons, the disputants, at least impli-

citly, accept distinct principles that count potentially distinct features of the circumstance as reasons for the acceptance or rejection of S.

Disagreements about reasons, however, need not occur in so idealized a form. So, for example, otherwise rational and reasonable palaeontologists can agree about the fossil record and yet disagree about what that record establishes. If they do, they are engaged in a disagreement about reasons: each implicitly accepts distinct principles that count potentially distinct aspects of the fossil record as reasons for the acceptance or rejection of the target claim. Similarly, constructivists and classical mathematicians disagree about what counts as a reason for accepting a mathematical sentence. Not only are methodological disputes in the special sciences disagreements about reasons, but so are disagreements that result from different styles of inductive reasoning. Disagreements about reasons may be theoretical, but they can be practical as well. Thus, Scanlon writes:

[Disagreement about reasons] is surely possible and perhaps even common. I think that it is plausible to suggest that we have an example of it in the contemporary disagreement between secular liberals like me, who see nothing morally objectionable about homosexuality, and conservative Christians who believe that it is a serious wrong. (Scanlon 1995: 352)

While in their idealized form disagreements about reasons plausibly never occur, in less idealized form such disagreements are plausibly ubiquitous.

Reacting to Disagreement

What is the rational response to a disagreement about reasons?

If we confine ourselves to what can be deontically described, then not only is it rationally permissible for Edgar to persist in his acceptance of S, but it is also rationally permissible for Edgar to revise—to reject or suspend judgment concerning S.

Edgar's persistence in his acceptance of S might be rationally permissible on a number of grounds. So, for example, in a cognitive domain Edgar might persist in his acceptance of S, despite the disagreement about reasons, on the grounds of doxastic conservatism.

Doxastic conservatism is the epistemic policy of persisting in one's beliefs unless presented with a positive reason to change one's mind (see Harman, 1986: chapter 4). Since the evidence's having a different normative appearance for Bernice is not a positive reason for Edgar to change his mind, if doxastic conservatism is a genuine epistemic norm, then it is rationally permissible for Edgar to persist in his belief that S. Doxastic conservatism is not the only grounds for the rational permissibility of persistence. Suppose, owing to some psychological necessity, Edgar simply cannot give up his acceptance of S. Since he must accept S, and is not self-contradictory or otherwise internally incoherent in so doing, it might be rationally permissible for him to persist in his acceptance of S.

Just as it is rationally permissible for Edgar to persist in his acceptance of S, it is rationally permissible for him to revise—to reject or suspend judgment concerning S. Revision might be rationally permissible on a number of grounds. So, for example, it might be rationally permissible for Edgar to revise if, upon reflection, he came to accept a debunking explanation for the disagreement between himself and Bernice, i.e. if he came to explain their disagreement in terms accidentally connected to reasons for acceptance. (See Cohen's, 2000, discussion of the paradox of conviction.)

Coming to accept a debunking explanation is not the only grounds for the rational permissibility of revision. Suppose that Edgar came to believe that there is a perfect symmetry between his epistemic position and Bernice's. Edgar could not coherently be a cognitivist and persist in accepting S, and in accepting that Bernice is wrong in rejecting S, while maintaining that there is a

perfect epistemic symmetry between them. (See Rosen's, 2001, discussion of the dispute between realists and fictionalists about abstracta.) If reflection on the disagreement about reasons prompts Edgar to accept a debunking explanation of their disagreement, or to accept that there is a perfect epistemic symmetry between himself and Bernice, then it would be rationally permissible for Edgar to revise—to reject or suspend judgment concerning S.

If we confine ourselves to what can be deontically described, then it would seem that persistence and revision are both rationally permissible. However, there is an important aspect of the rational response to a disagreement about reasons that has so far been left out of account. While, in the context of a disagreement about reasons, persistence and revision are both rationally permissible, sometimes at least, if acceptance is cognitive, there is something epistemically admirable about at least considering revising. After all, retaining belief on the grounds of conservatism, psychological necessity, and the like can seem like a reluctant capitulation to epistemic necessity. At any rate, acquiescing on such grounds is hardly a cognitive achievement. In contrast, a decision to reconsider manifests a responsiveness to reasons that is itself manifestly reasonable. Upon determining that his disagreement with Bernice is, at bottom, a disagreement about reasons, Edgar might be motivated to re-examine his reasons for accepting S. Edgar might inquire further into the grounds of his acceptance to determine whether, in light of his discussion with Bernice, his reasons for acceptance are good reasons. He might also inquire further to determine, in light of his discussion with Bernice, what, if anything, there is to Bernice's reasons for rejection. After all, Bernice might be onto something that so far eludes Edgar. While Edgar is not rationally required to inquire further into the grounds of acceptance, in the sense that his failure to do so would not be epistemically blameworthy, there would be something epistemically admirable about his

inquiring further. There is something cognitively virtuous about being motivated to inquire further into the grounds of acceptance in the face of a disagreement about reasons. If that is right, then there is a normative aspect of belief that is not describable in deontic vocabulary. Belief involves a cognitive virtue not describable in terms of rational permissibility. (See Rosen, 2001, for a similar suggestion.)

Two Kinds of Rational Norm

How are we to understand this? Here is one suggestion.

There are two kinds of rational norm.

The first kind of rational norm are those norms governing combinations of sentences that are candidates for acceptance. They take the form of principles determining whether combinations of sentences are rationally permissible, forbidden, or obligatory to accept. In a cognitive domain, they represent substantive judgments about the requirements of explanatory coherence on the epistemic state of a person at a time. A failure to conform to such a norm (by accepting a rationally forbidden combination of sentences, say) is irrational or, at the very least, epistemically blameworthy.

However, if acceptance is conceived as part of the broader activity of inquiry, where inquiry is one activity among many, then it is plausible that inquiry involves other kinds of norms as well. Just as there are rational norms governing combinations of sentences that are candidates for acceptance, it is plausible as well that there are rational norms governing the ends involved in inquiry. They represent the requirements on the ends to be adopted in changing one's epistemic state over time. So, for example, it is plausible to suppose that people who are motivated to inquire whether to accept S are rationally obliged to adopt means for determining whether to accept S consistent with other ends that they have adopted. Not only does inquiry, conceived as

a complex activity, involve the adoption of appropriate instrumental ends, but it is plausible as well that there are noninstrumental ends internal to inquiry that people engaged in that activity are rationally obliged to adopt.

To see this, consider the following. Suppose that Edgar accepts certain sentences that are indirectly inconsistent. The sentences that he accepts are not directly inconsistent—he does not simultaneously accept both a sentence and its negation. Rather, he accepts certain sentences such that there is a possible argument, each step of which involves immediate implications of the sentences he accepts to a conclusion that leaves this pattern of acceptance directly inconsistent. Is Edgar thereby irrational? No. He might have good reasons for accepting each of these sentences and might not recognize that they are indirectly inconsistent. Suppose that Edgar comes to recognize that the sentences he accepts are indirectly inconsistent. Perhaps Bernice has explicitly given him the argument leading to direct inconsistency. Edgar would be rationally obliged as an inquirer to adopt the end of resolving this inconsistency. It is in this sense that resolving such inconsistency is an end internal to inquiry.

While Edgar may be obliged as an inquirer to adopt the end of resolving this inconsistency, it would not be irrational for Edgar to persist in his acceptance, saying to Bernice: ‘That’s really interesting, I’ll have to think about that later, but right now I have to pick up the kids.’ In the meantime, Edgar may persist in accepting the indirectly inconsistent sentences, taking care not to infer everything that is implied by them. (Recall, if you can, your initial reaction to the Liar paradox, or to Zeno’s paradox.) Edgar’s discovery that the sentences he accepts are indirectly inconsistent may rationally oblige him as an inquirer to adopt a certain end—namely, the end of resolving this inconsistency—but there is latitude in the fulfilment of this end. Resolving this inconsistency might be hampered by inevitable practical exigencies. After all, a person has a plurality of ends and the fulfilment

of these must be rationally ordered. Rationally ordering adopted ends involves, among other things, prioritizing ends, scheduling means towards their fulfilment, and choosing means compatible with other ends. So the end of resolving the indirect inconsistency must in this way harmonize with other ends that Edgar has adopted (such as executing the daily routines involved in child-care). Whereas the content of the rational norm is purely epistemic—it requires only the adoption of an end internal to inquiry, the conditions for the fulfilment of this norm are not purely epistemic—acting on the obligatory end is constrained by considerations of practical coherence. Thus, there is latitude in the fulfilment of this end, and a failure to act towards its fulfilment merely lacks epistemic merit and is neither an instance of irrationality nor in any way epistemically blameworthy (though perhaps adopting the policy of never acting towards its fulfilment might be). I may not have resolved to my satisfaction the Liar paradox; I may never do so. But that does not make me irrational. My resolution of the paradox might be epistemically admirable, it might have epistemic merit; but my failure to do so is neither an instance of irrationality nor in any way epistemically blameworthy.

Inquiry thus involves two kinds of rational norm. On the one hand, there are rational norms governing combinations of sentences that are candidates for acceptance. They represent the requirements of explanatory coherence on the epistemic state of a person at a time. These are *strict* obligations. On the other hand, there are rational norms governing the ends involved in inquiry. They represent requirements on the ends to be adopted in changing one's epistemic state over time. These are *lax* obligations. Their laxity consists in the latitude involved in the fulfilment of these obligations, since an obligatory end is one end among many and a person's ends must be rationally ordered. The failure to fulfil a strict obligation is irrational or, at the very least, epistemically blameworthy. In contrast, any particular

action taken to fulfil a lax obligation is epistemically meritorious, while any particular failure to act merely lacks epistemic merit and is neither an instance of irrationality nor epistemically blameworthy (though perhaps adopting the policy of never acting to fulfil the lax obligation would be). *Cognoscenti* will recognize this as an application to the epistemic case of the distinction between perfect and imperfect duties (see Kant, 1785/1999: 4:421–3; 1797/1999: 6: 390–1.)

In a disagreement about reasons, persistence and revision are both rationally permissible. As such, the epistemic states involved in persistence and revision are in violation of no strict obligation. However, there is something admirable about at least considering revising, and this suggests the presence of a lax obligation, i.e. the presence of a rational obligation to adopt a certain end. What end could this be? Upon determining that his disagreement with Bernice is, at bottom, a disagreement about reasons, Edgar is under a lax obligation to inquire further into the grounds of acceptance. More precisely, given that he is interested in the truth of S, Edgar, in the context of a disagreement about reasons, has a reason to re-examine his reasons for accepting S, at least if his disputant is otherwise rational and reasonable, informed, and similarly interested in inquiring about S.

Let me explain. Even in the context of a disagreement about reasons, whether a person has a reason to inquire further depends on his interest in the truth of S. After all, ‘The truth, the whole truth, and nothing but the truth’ has never been a reasonable norm of inquiry. Absent some special interest, there is no reason to know whether Genghis Khan ever suffered from a hangnail, say. (See Harman, 1986: 55–6, for an explanation of the indispensability of interest in terms of the finite nature of inquirers.) However, given his interest in the truth of S, in the context of a disagreement about reasons, Edgar would have a motive to inquire further into the grounds of acceptance—or, at least, he would be so motivated if Bernice were otherwise

rational and reasonable, informed, and similarly interested in inquiring about S. Obviously, Edgar would lack this motive if Bernice were irrational, or unreasonable, or ignorant, or were moved by ulterior motives unconnected with reasons for acceptance. But if she is none of these, Edgar would have a motive to inquire further into his grounds for acceptance to determine whether, in light of his discussion with Bernice, his reasons for acceptance are good reasons. Edgar 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. After all, Bernice might be onto something that so far eludes Edgar. To inquire further is to strive to be responsive to what reasons there are. This would involve seriously considering the alternatives and so questioning the evidential status of initial appearances. While persistence is rationally permissible, Edgar must be prepared to bracket his full acceptance of S when re-examining his reasons for acceptance. 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. If, however, Edgar were to fulfil this end, he might satisfy himself with his acceptance of S, or he might suspend judgment concerning S, or might even reject S. Whatever the outcome, Edgar's noncomplacency in inquiring further would be epistemically admirable. Moreover, a failure to act towards the fulfilment of this end, to become responsive to what reasons there are, would merely lack epistemic merit and would be neither an instance of irrationality nor in any way epistemically blameworthy. Striving to be responsive to what reasons there are is, in this sense, a manifestation of cognitive virtue.

Two Kinds of Acceptance

The discussion so far provides preliminary support for the following claim:

Noncomplacency

If acceptance is cognitive, then, in the context of a disagreement about reasons, a person is under a lax obligation to inquire further into the grounds of acceptance. Specifically, if a person is interested in the truth of S, then, in the context of a disagreement about reasons, he would have a reason to re-examine his grounds for accepting S, at least if his disputant is otherwise rational and reasonable, informed, and similarly interested in inquiring about S.

Why should a disagreement about reasons motivate a person to inquire further into the grounds of acceptance? Perhaps the value of being reasonable constitutes a reason to inquire further. After all, in the face of a disagreement about reasons, to inquire further into the grounds of acceptance is to strive to be responsive to what reasons there are. Perhaps, then, in the context of a disagreement about reasons, if a person is interested in the truth of S, the value of being reasonable constitutes a reason to inquire further into the grounds of acceptance in the sense that part of what it is to be reasonable is to be so motivated in such circumstances (just as part of what it is to be benevolent is to be motivated by the good of others).

While initially plausible, noncomplacency is nevertheless controversial. A reasonable person interested in the truth of S would be motivated to inquire further if, in that context, it were open to reflective doubt about whether his reasons for acceptance were genuine reasons. But why must a disagreement about reasons invariably generate reflective doubt about the disputants' reasons? A disagreement about reasons would generate reflective doubt about the disputants' reasons if there were an acknowledged epistemic symmetry between them. But, in the context of any actual disagreement, it is plausibly always open to a person to simply deny that the symmetry obtains. Suppose that Bernice accepts that the Earth is flat and is ideally coherent in

so accepting: Bernice is otherwise rational and reasonable, informed, and can offer what Edgar would regard as a question-begging argument for her acceptance of the flat earth hypothesis. If Edgar is like us, he would reject any suggestion that there is a perfect epistemic symmetry between himself and Bernice but would maintain, instead, that he, and not Bernice, is appropriately related to the shape of the earth. Thus, it is implausible to suppose that Edgar's disagreement about reasons with an ideally coherent flat-earther would generate reflective doubt about his reasons for rejecting the flat earth hypothesis, and hence it is implausible to suppose that he would be motivated to inquire further into the grounds of his rejection, if interested. In the absence of such reflective doubt, a reasonable person interested in the truth of S may not be motivated to inquire further into the grounds of acceptance. In the absence of a special reason to doubt his reasons, a reasonable person interested in the truth of S may be satisfied that he is simply better placed to appreciate the facts.

The credibility of noncomplacency is thus subject to two apparently conflicting reactions: While initially plausible, it is controversial upon reflection. These apparently conflicting reactions would be reconciled, however, if it turned out that they were reactions to different things. Indeed, I believe that they are reactions to different things, for noncomplacency is ambiguous. They are two kinds of acceptance that are governed by distinct norms. Thus, noncomplacency can be understood as a claim about the norms governing one kind of acceptance or the other. Understood one way, this is plausible; understood the other way, it is controversial and plausibly false.

When a competent speaker accepts a sentence, he may accept that sentence for himself, but, importantly, he might do more than that. Not only may he accept the sentence for himself, but he might also accept that sentence on behalf of others. Acceptance for oneself is the object of individual inquiry. If

a competent speaker accepts a sentence *S* for himself, then he takes himself to have sufficient reason to end his individual inquiry about *S*. So if Edgar accepts ‘The UCL Philosophy Department is located at 19 Gordon Square’ he has no further reason to inquire about the address. (Of course, he might still have a reason to ask Bernice what that address is—say, in order to determine whether Bernice knows that address. However, in asking Bernice, Edgar is not inquiring after the address; rather, he is inquiring after Bernice’s knowledge of that address.) Acceptance for oneself is also the grounds of individual inquiry: if a competent speaker accepts *S* for himself, then he takes himself to have sufficient reason to rely on his acceptance of *S* in further theoretical and practical reasoning. So, for example, Edgar may rely on his acceptance of the address to estimate the time it would take to get there from his present location. Whereas acceptance for oneself is both the object and the grounds of individual inquiry, acceptance on behalf of others is the object and grounds of public inquiry: if a competent speaker accepts *S* on behalf of others, he takes himself to have sufficient reason to end public inquiry about *S*. Suppose Bernice asks Edgar for the address of the UCL Philosophy Department. If Edgar accepts the address on behalf of others, then, by his lights, there is no need for Bernice to inquire further—she may simply take his word for it. By his lights, his acceptance of the address can stand proxy for her own reasoning in inquiring about that address. Acceptance on behalf of others is also the grounds of public inquiry: if a competent speaker accepts *S* on behalf of others, he takes himself to have sufficient reason for others to rely on his acceptance of *S* in their own theoretical and practical reasoning. So, if Edgar accepts the address on behalf of others, then, by his lights, Bernice can rely on that address in her own theoretical and practical reasoning—she may, for example, rely on that address as a partial means of getting there.

The two kinds of acceptance involved in individual and public inquiry are governed by distinct norms. Thus, if a competent speaker accepts *S* on behalf of others, then he must coherently suppose, at least implicitly, that others do not accept reasons that would undermine his acceptance of *S*. (Harman, 1986: 51, observes that, in this respect, acceptance on behalf of others is like speaking on behalf of a group.) Suppose that Bernice asks Edgar where the UCL Philosophy Department is and he says that it is at 19 Gordon Square. Suppose, however, that Bernice has seen a flyer announcing that the Philosophy Department has moved from that address but she cannot now remember the ‘new’ address. Bernice would then accept a reason that undermines Edgar acceptance of that address. Thus, Edgar would not be justified in accepting that address on behalf of others because others, who are otherwise rational and reasonable, informed, and interested in inquiring about the address, accept undermining reasons and so reasonably reject that address. Nevertheless, Edgar could be justified in accepting the address for himself. Suppose that Edgar does not give full credence to the flyer. Perhaps he coherently supposes that it is a prank. Though Edgar cannot give Bernice a reason sufficient to rule out the evidence provided by the flyer, he does not give it credence, is coherent in not giving it credence, and continues to accept the ‘old’ address. While Edgar would not be justified in accepting the address on behalf of others, he would be justified in accepting that address for himself. This could be so only if acceptance for oneself and acceptance on behalf of others were governed by distinct norms.

This is further confirmed by the following. Suppose that Edgar is motivated to accept the address not only for himself but on behalf of others as well. (Suppose that Edgar and Bernice have a joint appointment there.) He might then look for evidence that would explain away the flyer. What explains his further inquiry is his motivation to accept the sentence on behalf of others coupled

with the fact that he takes himself to be justified in accepting it for himself. After all, if Edgar did not take himself to be justified in accepting the sentence for himself, then why bother looking for evidence that would explain away the undermining reason provided by the flyer?

In a cognitive domain, being justified in accepting a sentence for oneself and being appropriately related to the facts is no guarantee that one has sufficient reason to accept the sentence on behalf of others. Suppose that Edgar were correct in supposing the flyer to be a prank. Edgar would be justified in accepting the 'old' address and is appropriately related to the facts. However, Edgar would not be justified in accepting this on behalf of others since others who are otherwise rational and reasonable, informed, and similarly motivated coherently accept undermining reasons and, hence, reasonably reject the 'old' address. So just because a person is appropriately related to the facts is no guarantee that he possesses sufficient reason to accept *S* on behalf of others.

Noncomplacency can be understood as a claim about the norms governing acceptance for oneself or as a claim about the norms governing acceptance on behalf of others. Understood as a claim about acceptance for oneself, noncomplacency is false. Absent reflective doubt about his reasons for accepting *S*, Edgar is under no obligation to inquire further into the grounds of acceptance even if he is interested in the truth of *S*, and the mere fact of a disagreement about reasons is in general insufficient to generate such reflective doubt, even if his disputant is otherwise rational and reasonable, informed, and interested in the acceptability of *S*. Suppose, however, that Edgar is interested in accepting *S* not only for himself, but on behalf of others as well. (Perhaps the truth or falsity of *S* is relevant to a joint endeavor that Edgar is undertaking with Bernice, or perhaps they are participants of a public inquiry engaged in competing research programs.) In accepting *S* on behalf of others, Edgar

must coherently suppose, at least implicitly, that others who are otherwise rational and reasonable, informed, and interested in inquiring about S do not accept undermining reasons and hence, reasonably, reject S. Unfortunately, Bernice, who is otherwise rational and reasonable, informed, and interested in inquiring about S, accepts a reason that would undermine Edgar's acceptance of S. So, in the context of a disagreement about reasons, Edgar would lack sufficient reason to accept S on behalf of others even if he had sufficient reason to accept S for himself and were appropriately related to the facts. Since Edgar is interested in accepting S on behalf of others but lacks sufficient reason to do so, he would be motivated to inquire further into the grounds of acceptance to discover, if he can, grounds for accepting S that otherwise rational and reasonable, informed persons who are interested in inquiring about S could not reasonably reject. Inquiring further into the grounds of acceptance is an obligatory end of public inquiry for those engaged in a disagreement about reasons.

Moral Authority

Noncomplacency should be understood as a claim about the norms governing acceptance on behalf of others. The idea is that, if a person is interested in the truth of S, then, in the context of a disagreement about reasons, he would have a reason to re-examine his grounds for accepting S, at least if his disputant is otherwise rational and reasonable, informed, and similarly interested in inquiring about S. One distinctive feature of morality (or at least that part of morality that Gibbard, 1990, describes as 'morality in the narrow sense' and that Scanlon, 1998, describes as the domain of 'what we owe to each other') is its authority. Given the nature of its authority, moral acceptance is always acceptance on behalf of others. This is epistemically significant—for, taken together with noncomplacency, it has the following

important consequence: if moral acceptance is cognitive, then, in the context of a disagreement about reasons, a person is under a lax obligation to inquire further into the grounds of moral acceptance.

Morality is authoritative. After all, while morality in some sense answers to our concerns, it is also in some sense independent of them. The authority of morality is manifest in the role it plays in moral discourse and in the cognitive psychology of competent speakers. A full account of that authority would involve specifying its source in a way that made it intelligible that it should exhibit that role. However, without giving a full account of moral authority, a partial description of the role it plays in moral discourse and in the cognitive psychology of competent speakers suffices to establish that moral acceptance is always acceptance on behalf of others.

In sincerely uttering a moral sentence that he understands, a competent speaker accepts the uttered moral sentence, and, in accepting it he accepts as well what reason is thereby provided. Bernice accepts that abortion is wrong and thereby accepts as well a reason not to have an abortion if pregnant. Moreover, the reason that Bernice accepts, if genuine, potentially overrides whatever reason she might have to have an abortion if pregnant. Not only do moral reasons potentially override whatever conflicting nonmoral reasons we have for acting in the given circumstance, but they can also potentially cancel such reasons in that circumstance (see Frankfurt, 1988, chapter 13; McDowell, 1998: 55–6, 91–3; and Scanlon, 1998: 156–7). Sometimes a moral reason doesn't so much outweigh nonmoral inclination as discounts it as a reason for acting in the given circumstance. It is implausible to claim that moral reasons necessarily override or cancel all conflicting nonmoral reasons. Suppose Bernice promises to meet Edgar, but an important and rare opportunity arises such that Bernice cannot avail herself of that opportunity if she fulfils her promise to Edgar. If the opportunity were important

enough, and the promise was lightly given and of no great consequence to Edgar, then what reason there is to avail herself of that opportunity might outweigh the reason she has to meet Edgar (which is not, of course, to say that amends should not be made). While moral reasons do not always override or cancel conflicting nonmoral reasons, they very often do, and it is part of their nature and importance that they do. Thus, the moral reasons conveyed by our moral utterances often take precedence over conflicting nonmoral reasons:

Precedence

In uttering a moral sentence that he understands, a competent speaker conveys a reason to act in a given circumstance that potentially overrides or cancels any conflicting nonmoral reasons available in that circumstance. (Precedence is a variant of what Rawls, 1971, and Scanlon, 1998, describe as ‘the priority of right.’)

In accepting the wrongness of abortion, not only does Bernice accept a reason that takes precedence over nonmoral reasons, but she also takes the reason not to be contingent upon her acceptance of it. The acceptance of a moral reason is not a matter of taste. Of course, that something is to your taste is often a reason to prefer it. Bernice has a taste for westerns and distastes musicals. In deciding to watch one of two movies, a western and a musical, satisfying her taste for westerns constitutes a reason to watch the western rather than the musical. Moral reasons differ from matters of taste not in the sense that the former are reasons whereas the latter are not, but in the kind of reason they are.

According to Gibbard (1990: 164–6), if something is a matter of taste, satisfying that taste would not constitute a good reason if one lacked that taste. Bernice would not have a reason to watch the western if instead musicals were to her taste. In accepting something as a matter of taste, a person does not take that reason

to apply independently of his accepting it: if he lacked that taste, he would lack that reason. (N.B.: The sense of ‘taste’ that Gibbard deploys is the one associated with the *de gustibus* motto and is distinct from the sense of taste whose standard Hume sought to establish.) Matters are different with moral reasons. In accepting that abortion is wrong, Bernice accepts a reason not to have an abortion if pregnant. Moreover, Bernice believes that she would still have a reason not to have an abortion if instead she accepted that abortion was morally permissible. The moral reason not to have an abortion applies, if it does at all, independently of Bernice’s accepting it. Indeed, it applies, if it does at all, independently of *anyone’s* accepting it. Emma does not accept that abortion is wrong because she is unsure about the moral status of abortion. Though Emma does not accept that abortion is wrong, Bernice believes that Emma has a reason not to have an abortion if pregnant even though Emma does not accept that reason. In accepting a moral reason, a person takes that reason to apply independently of a person’s accepting it. The content of a moral reason is not linked to a person’s acceptance of it the way a reason of taste is. If an action is wrong in a given circumstance, then everyone who is in that circumstance has a reason not to perform that action, whether or not they accept that reason:

Noncontingency

In uttering a moral sentence that he understands, the existence of the reason conveyed is not contingent upon the speaker’s or anyone else’s accepting it.

In accepting the wrongness of abortion, not only does Bernice accept a reason that takes precedence over nonmoral reasons, that is not contingent upon her acceptance of it, but she also believes that she has good reason to accept that it is wrong and that this is a reason, not only for her, but for everyone else as well. Emma, unlike Bernice, does not accept that abortion is wrong, but from Bernice’s perspective Emma is thereby

unreasonable, if not indeed irrational. Bernice regards Emma as unreasonable in the sense that Emma is not responding to what reason there is to accept the wrongness of abortion. (She is not, however, irrational, at least not in the narrow sense of acting at variance with a reason she accepts.) The putative reason that Bernice has for accepting that abortion is wrong is a reason for everyone to accept that abortion is wrong, or would be if it were a genuine reason. We might describe this as *well-groundedness*: In uttering a moral sentence that he understands, the reason a competent speaker has, if sincere, for accepting the uttered moral sentence applies not only to the speaker but to everyone else as well.

There is a complication, however. Emma is unsure about the moral status of abortion. Suppose, however, she comes to regard Bernice's moral opinion as authoritative in this instance. Perhaps, Emma trusts Bernice's moral sensibility more than her own in the given circumstance. While Emma is sure that Bernice has a good reason to accept the wrongness of abortion, Emma herself remains unclear about that reason—she is unclear about which features of her circumstance count as a reason for the impermissibility of abortion, or even why these features should have this normative significance. Emma's reason for accepting the wrongness of abortion is that Bernice advises her that it is wrong. However, this is not a reason for someone who does not trust Bernice's moral sensibility the way Emma does. So Emma's reason for accepting the wrongness of abortion is not a reason for others to accept the wrongness of abortion. Thus, well-groundedness is false as presently formulated. Can this principle be reformulated to accommodate this complexity? Emma's reason for accepting the wrongness of abortion is that Bernice advises her that it is wrong. Bernice's advice is a reason, if it is, because Bernice has good reason for the wrongness of abortion. This reason, if genuine, is a *grounding* reason, since it is the grounds for the wrongness of abortion, and is a reason not

only for Bernice but for everyone else as well to accept the wrongness of abortion. Bernice's advice is reason to accept the wrongness of abortion because Emma believes that Bernice has reasons that ground the wrongness of abortion. Although Emma is unclear about the nature of these grounds, she is nevertheless sure of their existence. Moral testimony can provide access to grounding reasons even to persons who lack an adequate conception of those reasons. Moreover, it is these grounding reasons, if they exist, that are reasons for everyone to accept the wrongness of abortion and, hence, are reasons to accept the wrongness of abortion on behalf of others. The principle should be reformulated as follows:

Well-groundedness

In uttering a moral sentence that he understands, the grounding reason a competent speaker directly or indirectly has, if sincere, for accepting the uttered moral sentence applies not only to the speaker but to everyone else as well. So, in sincerely uttering a moral sentence that he understands, a competent speaker accepts that sentence on behalf of others.

There is one further feature of moral authority that is worth emphasizing. In accepting the wrongness of abortion, not only does Bernice accept a reason that takes precedence over non-moral reasons, that is not contingent upon her acceptance of it, for which there are grounds not only for her but for everyone to accept, but in uttering 'Abortion is wrong' Bernice is *demanding* that everyone accept that it is wrong. Stevenson (1937, 1944) highlights this feature of moral authority by the 'do so as well' component of his analysis. According to Stevenson, Bernice, in claiming that abortion is wrong, not only represents herself as disapproving of abortion, but also demands that others do so as well. In making a moral utterance, a competent speaker demands that his audience accept the uttered moral sentence and

so come to respond affectively in the relevant manner. The relevant response need not be the same as the speaker's: it may make sense for the hearer to feel guilty and the speaker to feel angry, say, but if it does this difference is grounded in their different relative positions in the circumstance and the normative appearance it presents:

Demand

In uttering a moral sentence that he understands, a competent speaker demands that his audience accept the uttered moral sentence.

While Stevenson emphasizes the demand conveyed by moral utterance, what is perhaps missing in his account is the recognition that the reasons that ground the acceptance of the uttered moral sentence are linked with this demand. Thus, MacIntyre writes:

Stevenson . . . understood very clearly that saying 'I disapprove of this; do so as well!' does not have the same force as 'This is bad!' He noted that a kind of prestige attaches to the latter, which does not attach to the former. What he did not note however—precisely because he viewed emotivism as a theory of meaning—is that the prestige derives from the fact that the use of 'That is bad!' implies an appeal to an objective and impersonal standard in a way in which 'I disapprove of this; do so as well!' does not. (MacIntyre, 1981: 19–20)

Stevenson's account of moral utterance fails to capture its authority, since he does not link the conversational demand it conveys with the reasons for accepting the uttered moral sentence. This is plausibly the source of the traditional criticism that moral utterance, as Stevenson conceives of it, is a form of manipulation—since competent speakers demand that others adopt the relevant emotional attitude without providing them with a reason for adopting that attitude. The conversational demand that others accept the uttered moral sentence is justified only if the speaker possesses a grounding reason for accepting

the uttered moral sentence that applies not only to himself but to everyone else as well. The conversational demand is justified only if the speaker has sufficient reason to accept the uttered moral sentence on behalf of others. Demanding that others accept the uttered moral sentence may require that the speaker possess a grounding reason that applies not only to himself but to everyone else as well, but the possession of a grounding reason only potentially justifies demanding that others accept the moral sentence. So, for example, a competent speaker would not be justified in demanding that another accept a moral sentence, even if he possessed a grounding reason, if, in that context, so demanding would humiliate the other.

Given the nature of moral authority, moral acceptance is always acceptance on behalf of others and so is subject to the appropriate norms. David Hume puts the point this way:

When a man denominates another his *enemy*, his *rival*, his *antagonist*, his *adversary*, he is understood to speak the language of self-love, and to express sentiments, peculiar to himself, and arising from his particular circumstances and situation. But when he bestows on any man the epithets of *vicious* or *odious* or *depraved*, he then speaks another language, and expresses sentiments, in which, he expects, all his audience are to concur with him. He must here, therefore, depart from his private and particular situation, and must chuse a point of view, common to him and others. (Hume 1751/1988: section 9)

Reacting to Moral Disagreement

Given the nature of moral authority, moral acceptance is always acceptance on behalf of others. Not only does well-groundedness entail that accepting a moral sentence involves accepting that sentence on behalf of others, but the conversational demand conveyed by moral utterance is intelligible only if sincerely uttering a moral sentence involves accepting that sentence on behalf of others. This is epistemically significant; for, taken

together with noncomplacency, it has the following important consequence. If moral acceptance is cognitive, then, in the context of a disagreement about reasons, a person is under a lax obligation to inquire further into the grounds of acceptance. By contraposition, it follows that if, in the context of a disagreement about reasons, a person is not under a lax obligation to inquire further into the grounds of moral acceptance, then moral acceptance is noncognitive. In this section I will argue for noncognitivism on just these grounds.

Edgar and Bernice disagree about the moral status of abortion. Whereas Bernice accepts the wrongness of abortion, Edgar is a complacent liberal moralist and accepts that abortion is permissible. Upon discussing the matter, Edgar and Bernice discover that their disagreement is a disagreement about reasons. Edgar reasons as follows (Edgar and Bernice's arguments are quoted from MacIntyre, 1981)

Everyone has a right over their own person and their own body. Given the nature of these rights, when an embryo is essentially part of the mother's body, the mother has the right to make her own uncoerced decision on whether she will have an abortion or not. Therefore, abortion is morally permissible. . . . (MacIntyre, 1981: 6–7)

So Edgar, implicitly at least, accepts a principle that counts a certain feature of the circumstance, the embryo being essentially part of the mother's body, as a reason to accept the permissibility of abortion. Bernice, however, rejects this principle. Bernice reasons instead as follows:

I cannot will that my mother should have had an abortion when she was pregnant with me, except perhaps if it had been certain that the embryo was dead or gravely damaged. But if I cannot will this in my own case, how can I consistently deny to others the right to life I claim for myself? I would break the so-called Golden Rule unless I denied that a mother has in general a right to an abortion. (MacIntyre, 1981: 6–7)

From Bernice's perspective, the embryo being essentially part of the mother's body has a different normative significance. As a consequence, Edgar and Bernice disagree about what would count as a reason for accepting or rejecting the sentence 'Abortion is permissible.' Edgar and Bernice's disagreement about the moral status of abortion is, at bottom, a disagreement about reasons.

In uttering 'Abortion is permissible' Edgar demands, implicitly at least, that his audience accept that sentence. So Edgar must accept that sentence on behalf of others if his utterance is sincere. Indeed, Edgar is sincere. He accepts 'Abortion is permissible' on behalf of others for he takes himself to have access to a grounding reason that is a reason to accept that sentence not only for himself, but for everyone else as well. Bernice, like Edgar, is motivated to accept on behalf of others a claim about the moral status of abortion. Supposing that she is an intelligent and articulate spokesperson, Bernice might strike Edgar as an otherwise rational and reasonable, informed human being who coherently accepts a reason that, if genuine, would undermine his acceptance of the permissibility of abortion. Nevertheless, Edgar feels no embarrassment about this. His persistence in his liberal morality is unflinching. Edgar is intransigent in the sense that he lacks a motivation to inquire further into the grounds of moral acceptance. Nor is Edgar alone in this. I suspect that we too would be unmoved by such a disagreement. Our own persistence in liberal morality would be unflinching as well. We too would be intransigent in the sense of lacking a motivation to inquire further into the grounds of moral acceptance. In normal circumstances, we are under no obligation to re-examine the foundations of moral claims that we accept as unproblematic even if they are disputed by otherwise rational and reasonable, informed, and interested people who coherently accept reasons that, if genuine, would undermine them.

I am not making an empirical claim about the actual extent of moral intransigence—that would require a sensitive interpretation of a moral sociology that has yet to be written; rather, I am making a conceptual claim about the norms that actually govern moral acceptance: Given the norms that we actually accept, it is intelligible to fail to be motivated to inquire further. If we can conceive of cases where such intransigence is intelligible, then it must be so at least by the norms governing moral acceptance that we actually accept and tacitly appeal to in so conceiving. Not only is it intelligible that one, as a matter of fact, takes no positive steps towards re-examining the grounds of moral acceptance—after all, one might have more compelling immediate concerns; but it is intelligible as well that one should lack this motivation altogether. And if the failure to adopt the end of further inquiry is intelligible, then we are under no rational obligation to adopt this end, at least by the norms of moral acceptance that we actually, if implicitly, accept.

As an illustration of this, consider Hilary Putnam's admirable description of the deep political disagreement between Nozick and himself:

But what of the fundamentals on which one cannot *agree*? It would be dishonest to pretend that one thinks that there are no better and worse views *here*. I don't think that it is just a matter of *taste* whether one thinks that the obligation of the community to treat its members with compassion takes precedence over property rights; nor does my co-disputant. Each of us regards the other as lacking, at this level, a certain kind of sensitivity and perception. 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. (Putnam, 1981: 165)

Putnam should be commended for his candour here. What Putnam holds in something akin to contempt is Nozick's moral sensibility ('a certain complex of emotions and judgments in the other')—a moral sensibility that privileges property rights over what Putnam regards as the compassionate treatment of the less well off. Nozick is a reasonable and interested person who accepts reasons that, if genuine, would undermine Putnam's commitment to liberal morality. But if Putnam holds Nozick's moral sensibility in something akin to contempt, what motivation would *Nozick's* accepting an undermining reason provide *Putnam* for inquiring further into the grounds of moral acceptance? None. Thus, the reaction that Putnam carefully describes is a manifestation of moral intransigence. The important point, however, is that Putnam's reaction is not obviously unintelligible. And if it isn't, then under such circumstances we would be under no obligation to re-examine the foundations of the liberal morality we accept, if we do.

Relativity and Error

The argument from intransigence can be summarized as follows. If acceptance were cognitive, then, in the context of a disagreement about reasons, a person would be under a lax obligation to inquire further into the grounds of acceptance. This, however, is plausible only if acceptance is understood as acceptance on behalf of others. Given the nature of its authority, however, moral acceptance is always acceptance on behalf of others. So if moral acceptance were cognitive, then in the context of a disagreement about reasons, a person would be under a lax obligation to inquire further into the grounds of moral acceptance. However, in the context of a disagreement about reasons, a person is under no obligation to inquire further into the grounds of moral acceptance. Therefore, moral acceptance must be noncognitive.

A moral relativist might object that the argument from intransigence is unsound. According to moral relativism, moral acceptance is belief in a moral proposition; it is just that the moral proposition has relative truth-conditions. Thus, if Edgar accepts that abortion is permissible, he believes that abortion is permissible, but that proposition is true only relative to a moral framework: abortion is represented as permissible only relative to the moral framework in which Edgar participates. Similarly, if Bernice accepts that abortion is wrong, she believes that abortion is wrong, but that proposition is true only relative to a moral framework: abortion is represented as wrong only relative to the moral framework in which Bernice participates. Suppose that Edgar and Bernice participate in distinct moral frameworks. A disagreement about reasons would not motivate Edgar to inquire further into the grounds of moral acceptance, any more than a disagreement about the perceived location of a rainbow would motivate him to inquire about the ‘true,’ or perspective-independent, location of the rainbow. Thus, according to the envisioned moral relativist, moral acceptance is cognitive, but in a disagreement about reasons a person is under no obligation to inquire further into the grounds of acceptance. The relativist provides an alternative explanation for the intelligibility of intransigence, and one that is consistent with cognitivism. He would thus object that the argument from intransigence is unsound because it involves a false premise:

If moral acceptance is cognitive, then, in the context of a disagreement about reasons, a person is under a lax obligation to inquire further into the grounds of moral acceptance.

The relativist makes a natural assumption—that, in a disagreement about reasons, the disputants participate in distinct moral frameworks. However, it is unclear why this assumption is invariably true. Why couldn’t a disagreement about reasons

arise within a single moral framework? After all, such moral frameworks, if they exist, are the contingent products of human culture and history. It is thus implausible to suppose that they are complete in the sense of partitioning the practical alternatives open to a person as permissible, forbidden, or obligatory in every possible circumstance. Unforeseen circumstances may give rise to practical alternatives unclassified by the moral framework. Indeed, there may be different ways of naturally extending the given framework to take these into account. But if this is the basis of the disagreement about reasons, then there is a clear sense in which the disputants participate in the same moral framework: they would be disagreeing about how to extend the incomplete moral framework that they share.

This is not the only way a disagreement about reasons can arise within a single moral framework. Consider a disagreement about the justice of an institutional policy. The disputants might disagree, for example, about the justice of an institutional policy, even though they agree about the description of the circumstance and the probable outcomes, and share substantially the same conception of justice. They might nevertheless reasonably disagree about how to apply the principles of justice they share to the circumstance as they commonly understand it. They would be engaged in a disagreement in reasons despite their common moral framework. If a disagreement about reasons can arise within a single moral framework, then, at least in these cases, the relativist loses his explanation for why the disputants are under no obligation to inquire further into the grounds of acceptance.

The relativist objection rests on a false, though natural, assumption—that in a disagreement about reasons the disputants participate in distinct moral frameworks. However, even if this difficulty were avoided, the relativist would face serious difficulties. The relativist denies the following:

If moral acceptance is cognitive, then, in the context of a disagreement about reasons, a person is under a lax obligation to inquire further into the grounds of moral acceptance.

However, this was a consequence of two claims:

Noncomplacency

If acceptance is cognitive and on behalf of others, then, in the context of a disagreement about reasons, a person is under a lax obligation to inquire further into the grounds of acceptance.

Authority

Given its authority, moral acceptance is always acceptance on behalf of others.

This is the basis of a dilemma: the relativist must deny either noncomplacency or authority. Recall that noncomplacency was urged on general grounds. So, if a relativist were to deny it, his relativism would no longer be confined to moral inquiry, but would instead be a form of epistemic relativism more generally. Epistemic relativism, however, is a controversial doctrine and should not be accepted merely by reflecting on the intelligibility of moral intransigence (but see Rosen, 2001, for a recent defence). If, on the other hand, the relativist were to deny authority—that moral acceptance is always acceptance on behalf of others—either by denying it outright or by restricting its scope, say, to those who are co-participants in a moral framework, he risks giving an unacceptably deflationary account of moral authority. These are serious difficulties, if not, perhaps, decisive ones. However, they suffice to cast doubt on the claim that relativist's alternative is clearly the better explanation of the intelligibility of moral intransigence. Indeed, the prevalence of relativist rhetoric in a moral culture might be a symptom of its noncognitive nature, since it might be a confused acknowledg-

ment of the intelligibility of moral intransigence that is best explained by noncognitivism.

An error theorist might also object to the soundness of the argument, though on different grounds.

Noncognitivists are committed to a nonmoral explanation of moral acceptance. But noncognitivists are not the only ones with this explanatory commitment. Consider the kind of error theory that Mackie (1977) espoused. Mackie held that moral facts are 'queer' (or would be if there were any) and best not believed in. So, according to Mackie, there are no moral facts. But if there are no moral facts, then a competent speaker's acceptance of a moral sentence cannot be subject to a moral explanation. Moral acceptance might be belief, but the explanation for moral acceptance cannot be that the moral facts are thus and so and that the speaker is justified in believing them to be, as the realist maintains. The explanation for why a speaker accepts a moral sentence must be a nonmoral explanation. Indeed, the nonmoral explanation might take the form that noncognitivists recommend. Sometimes belief can be explained in terms of the noncognitive attitudes of the believer. Wishful thinking is a case in point. Bernice believed that England would win the World Cup. However, her belief was not the product of an impartial assessment of comparative merit; rather, she believed that England would win because she desperately wanted them to. Similarly, the error theorist might claim that, in accepting a moral sentence, a speaker believes the proposition expressed not because the moral facts are as the proposition represents them to be, but because of the noncognitive attitudes of the speaker. So for example the error theorist might claim that Edgar believes that it is good to help those in need, not because helping those in need instantiates the property of goodness, but rather because of his compassion for the needy.

Suppose an error theorist maintains that moral acceptance is subject not only to a nonmoral explanation but to a noncognitive

explanation as well. This might seem to undermine the case for noncognitivism. Suppose moral sentences were best explained in terms of the noncognitive attitudes of the speaker. This is consistent with the possibility that a competent speaker, in accepting what he does on the basis of his noncognitive attitudes, is getting things systematically wrong.

The alleged difficulty is a product of not clearly distinguishing noncognitivism from noncognitive explanations of moral acceptance. It is one thing to claim that moral acceptance is subject to a noncognitive explanation; it is quite another to claim that moral acceptance is noncognitive. To claim that moral acceptance is noncognitive is to make a claim about the norms governing acceptance. If moral acceptance is in this sense noncognitive, then of course a competent speaker's acceptance of a moral claim will be subject to a noncognitive explanation. However, as our discussion of the error theory reveals, just because moral acceptance is subject to a noncognitive explanation, it does not follow that moral acceptance is noncognitive. Perhaps acceptance is belief. The fact that moral acceptance is best explained by the noncognitive attitudes of the speaker would then show only that the norms governing acceptance were being systematically violated. A noncognitive explanation of moral acceptance does not by itself establish the truth of noncognitivism. However, the conclusion of the argument from intransigence is not merely that moral acceptance is best explained by some attitude other than moral belief. Rather, as the intelligibility of moral intransigence reveals, the norms governing moral acceptance differ from the norms appropriate to belief. Thus, the error theorist's objection mistakes a conceptual claim about the norms that actually govern moral acceptance for an empirical claim about what actually explains moral acceptance. It ignores the way in which a potential explanatory difference between cognitivists and noncognitivists is a manifestation of a more fundamental normative difference.

The Argument from Aspect Shift

The argument from intransigence, if sound, only establishes the noncognitivist's distinctive denial. It remains silent, however, about the nature of the attitudes involved in moral acceptance. It is natural to ask: if moral acceptance is noncognitive, what kind of attitude is it? In this section we will consider a supplementary argument for noncognitivism, the argument from aspect shift, which has an informative conclusion about the nature of these attitudes. While not, even suitably elaborated, a full and substantive account of the nature of moral acceptance, the conclusion nevertheless provides positive information about the kind of attitudes involved in accepting a moral sentence.

So far, in discussing disagreements about reasons, we have naturally focused on interpersonal disagreement. However, corresponding to public moral conflict there is the possibility of private conflict: there could be intrapersonal disagreement corresponding to interpersonal disagreement. Specifically, in arguing for noncognitivism, the suggestion is that a disagreement about reasons might be approached from the deliberative perspective of a single practical reasoner.

There is an aspect of the phenomenology of intrapersonal conflict that is presently relevant. Emma is unsure about the moral status of abortion. She is genuinely undecided about the permissibility of abortion, even having considered Edgar and Bernice's explicit arguments. When Emma reflects on the rights people have over their own persons, certain features of her circumstance become salient and present a certain normative significance. Specifically, in rehearsing Edgar's argument, Emma has a tendency to focus on a material feature of the circumstance, the embryo being essentially a part of the mother, and a tendency to count this a reason for permitting abortion as well as a tendency to rule out other features of the circumstance, such as the inability to universalize the decision to abort, as a reason for

forbidding abortion. However, when she reflects on the value of her decisions, universalizing different features of her circumstance become salient and seem to have a different normative significance. Specifically, in rehearsing Bernice's argument, Emma has a tendency to focus on a formal feature of the circumstance, the inability to universalize a decision to abort, and a tendency to count that as a reason for forbidding abortion as well as a tendency to rule out other features of the circumstance, such as the embryo being essentially part of the mother, as a reason for permitting abortion. The rival arguments differently structure the reasons apparently available in the given circumstance. From the perspective of rights, certain features of the circumstance count as reasons and others are ruled out. From the perspective of universalizability, different features of the circumstance count as reasons and yet others are ruled out (compare Nagel, 1979: essay 9.) Since she cannot reconcile these arguments in a single coherent normative framework, Emma, in moving between these distinct normative perspectives, experiences what can only be described as a normative aspect shift: in moving between these distinct normative perspectives, different features of her circumstance become salient and seem to take on a different normative significance. Thus Ulrich, the protagonist of Musil's *The Man without Qualities* explains:

I maintained that a general who for strategic reasons sends his battalions to certain doom is a murderer, if you think of them as thousands of mothers' sons, but that he immediately becomes something else seen from another perspective, such as, for example, the necessity of sacrifice, or the insignificance of life's short span. (Musil, 1995: 295)

Noncognitivism provides the best explanation for the normative aspect shift involved in the phenomenology of intrapersonal conflict. Consider the way a normative perspective structures a person's moral consciousness. Adopting a normative perspective involves a tendency for certain features of the circumstance to

become salient in perception, thought, and imagination, and a tendency for these features to present a certain complex normative appearance. A normative perspective structures a person's moral consciousness in just the way that a certain kind of affect structures a person's consciousness.

Consider erotic desire. Edgar walks into a crowded room and is immediately struck by the presence of his beloved, Bernice. Bernice is immediately perceptually salient. Edgar experiences Bernice as quite literally standing out of the crowd. Bernice's salience is not merely confined to his perception. Ed has, as well, a tendency to focus on Bernice in thought and imagination. Not only is Bernice salient, but her desirability is also manifest in a phenomenologically vivid manner. Indeed in his vulnerable moments this phenomenologically vivid sense of Bernice's desirability can be unbearable and thus has a tendency to be shy-making. Not only does Ed have a tendency to see Bernice as desirable, but he also has a tendency to rule out from consideration certain features of Bernice that, to others at least, might count against her desirability. So, for example, he has a tendency to overlook certain annoying habits such as Bernice's penchant for chewing on pens when concentrating. It is not that Bernice's pen chewing is outweighed by her manifest desirability. At least in this instance, for Ed it is not even an issue. Indeed, Ed is so far gone that he sees Bernice's pen chewing as contributing to her unique charm. So Ed's desiring Bernice involves, among other things, a tendency for Bernice to become salient in perception, thought, and imagination as well as a tendency for her to present a certain normative appearance. Indeed, Scanlon (1998: chapter 1) characterizes a certain kind of affect, what he calls *desire in the directed attention sense*, precisely in terms of these effects. Specifically, according to Scanlon, a desire in the directed attention sense involves a tendency to focus on the object of desire as well as a tendency for the object of desire to appear in a favourable light. (There is a sense in which the label 'desire in the directed

attention' sense is inapt—it suggests a too narrow construal of the relevant kind of affect. Specifically, it suggests that the constituent normative appearance is invariably positive. However, whatever Scanlon's intention, I am not assuming that the object of the affect invariably appears in a *favorable* light, only that there is a tendency for the object of the affect to have a certain normative appearance, whether or not that appearance is positive.) So it seems that a normative perspective structures a person's moral consciousness in just the way a certain kind of affect, desire in the directed attention sense, structures a person's consciousness.

This might be so if a person's normative perspective were just their being disposed to respond affectively in the relevant manner. If that is right, then the normative aspect shift involved in the phenomenology of intrapersonal conflict is subject to a noncognitive explanation. From the perspective of rights, certain features of the circumstance count as reasons and others are excluded. From the perspective of universalizability, different features of the circumstance count as reasons and yet others are excluded. Emma, in moving between these distinct normative perspectives, experiences a normative aspect shift. This normative aspect shift is nothing other than Emma's vacillating between distinct and incompatible affective responses to her circumstance. In being unclear about the moral status of abortion, Emma quite literally does not know what to feel about it. (Thus Hume, 1740/2003: 3.1.2: 'Morality, therefore, is more properly felt than judg'd of'.)

This hypothesis receives some independent support by the rhetorical strategies deployed when reasons give out in basic moral disagreement. Thus, Edgar, in trying to persuade Bernice of the permissibility of abortion after she has listened to and rejected his explicit argument, might do any and all of the following: he might exhort Bernice to see it like this . . . , to consider certain pertinent analogies as well as certain cases

whose description has a narrative structure that expresses Edgar's feelings about abortion—he might even resort to brow-beating. These and other rhetorical strategies are essentially literary devices for focusing the audience's attention on certain features of the circumstance and presenting those features in a certain normative light. Such rhetorical strategies, when artfully deployed, get the audience to respond affectively in the relevant way to the given circumstance—they frame the perspective of the audience so as to induce the relevant affect. By such means Edgar might try to instill in Bernice what Putnam (1981: 165) describes as a 'certain complex of emotions and judgments.' In this way, a familiar, intuitionist rhetoric can have a noncognitive use. (See McIntyre, 1981: chapter 2; but see McDowell, 1998, essays 3, 10, for a different interpretation of this rhetoric.)

So far, the relevant kind of affect has been characterized in terms of its functional role, i.e. 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. But why do these effects hang together? What is it about the nature of this attitude that explains and renders intelligible that it should have this functional role?—Johnston (2001) forcefully presses this question. McDowell (1998: essays 3, 10) suggests that it is 'natural' for the noncognitivist to conceive of this attitude as a mixed state, a noncognitive refinement of sensing, where the noncognitive component is the source of the normative appearance. McDowell has done much to discredit the claim that the relevant kind of affect can be understood as a mixed state involving perceptual and noncognitive components that can be independently specified. Even if this kind of account were untenable and were the only substantive specification of the relevant kind of affect available to the noncognitivist, why should the noncognitivist provide a substantive account of desire in the directed attention sense? Why should desire in the directed attention sense be understood

as an attitude whose nature can be specified independently of its functional role, and can explain and render intelligible why this attitude has that functional role?

If the affect were conceived to be a particular event in a person's consciousness (a 'feeling,' in the philosopher's sense), then it would be reasonable to assume that its nature would be manifest in the way it structures a person's consciousness, and so reasonable to assume that the nature of the affect would explain and render intelligible the tendency of the object of the affect to become salient and the tendency for it to present a certain normative appearance. However, to assume at the outset that the affect is a particular conscious event is to overlook a metaphysical option available to the noncognitivist. Perhaps the affect is not some particular event in a person's consciousness, but the way in which events in the person's consciousness are structured. The suggestion is that there is nothing more to being an affect of the relevant kind than 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 *minimalist* denies that a person's being in an affective state consists in some further fact over and above the relevant way in which the person's consciousness is structured. Minimalism is thus the analogue in the philosophy of mind of T. S. Eliot's notion of the 'objective correlative':

The only way of expressing emotion in the form of art is by finding an 'objective correlative'; in other words, a set of objects, a situation, a chain of events which shall be the formula of that *particular* emotion; such that when the external facts, which must terminate in sensory experience, are given, the emotion is immediately evoked. (Eliot, 1932: 145)

This conception of emotional expression in art is the basis of Eliot's criticism of *Hamlet*: the emotion that has Hamlet in its grips is inexpressible precisely because it is a further fact over and

above the structure of events in the narrative, and for this reason Eliot reckons the play a failure:

If you examine any of Shakespeare's more successful tragedies, you will find this exact equivalence; you will find that the state of mind of Lady Macbeth walking in her sleep has been communicated to you by a skilful accumulation of imagined sensory impressions; the words of Macbeth on hearing of his wife's death strike us as if, given the sequence of events, these words were automatically released by the last event in the series. The artistic 'inevitability' lies in this complete adequacy of the external to the emotion; and this is precisely what is deficient in *Hamlet*. Hamlet (the man) is dominated by an emotion which is inexpressible, because it is in *excess* of the facts as they appear. (Eliot, 1932: 145)

Just as Eliot contends that it is the structure of the events in the narrative ('the facts as they appear'), and not some further fact, that constitutes the expression of emotion, the minimalist contends that it is the structure of the events in a person's consciousness, and not some further fact, that constitutes the relevant affect.

McDowell (1998: essays 3, 10) suggests that it is 'natural' for the noncognitivist to conceive of this attitude as a mixed state, a noncognitive refinement of sensing, where the cognitive and noncognitive components can be independently specified and where the nature of this mixed state explains and renders intelligible the way a person's consciousness is structured. According to the minimalist, desire in the directed attention sense is a mixed state: it is a noncognitive attitude that involves thoughts and perceptions about the morally salient features of the circumstance. However these attitudes are not distinct and so cannot be independently specified; nor can they explain the way they structure a person's consciousness. The thoughts and perceptions involved in moral acceptance are events in a person's consciousness whose structure constitutes the relevant affect. So, even if a substantive account of desire in the directed

attention sense were unavailable to the noncognitivist for the reasons McDowell describes, a noncognitivist might still claim that the relevant affect is nothing other than the tendency for certain features of the circumstance to become salient in perception, thought, and imagination and the tendency for them to present a certain complex normative appearance.

If one accepts the minimalist account of desire in the directed attention sense, then the case for noncognitivism is strengthened. It would no longer be a question of noncognitivism providing the best explanation for the normative aspect shift: minimalism and the claim that the affect is noncognitive would entail a noncognitive account of normative aspect shift.

Accepting a moral sentence will seem reasonable, given an appropriate background normative perspective. From the normative perspective of rights, accepting that abortion is permissible might seem reasonable. From the perspective of universalizability, accepting that abortion is wrong might seem reasonable. Given that the adoption of a normative perspective is just a matter of appropriately configuring one's affective sensibility, it is plausible that moral acceptance is itself noncognitive. Moreover, if it is, then it is no surprise that it should lack a cognitive virtue that genuine belief displays.

Conclusion

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, 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 reason 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 affects centrally involved in moral acceptance are in this way essentially other regarding.