

PHIL 260; Spring 2007  
 The Normativity of Ethics  
 David O. Brink  
 Handout #1: Moral Motivation and Expressivism

### A PUZZLE ABOUT MORAL MOTIVATION

Our puzzle about moral motivation can be seen as a tension that we encounter when we try to reconcile intellectual and practical aspects of morality. **Cognitivists** interpret moral judgments as expressing cognitive attitudes, such as belief. Moral judgments ascribe properties – axiological, deontic, and aretaic – to persons, actions, institutions, and policies. **Internalists** believe that moral judgments necessarily engage the will and motivate. We expect people to be motivated to act in accord with their moral judgments and would find it odd for people to be systematically indifferent to what they judge morally significant. It is also a common view that motivation involves **pro-attitudes**, such as desires. We explain intentional action as the product of informational states, such as beliefs, and practical states, such as desires. But beliefs and desires are logically independent states; no belief entails any particular desire. These assumptions are in tension.

1. Moral judgments express beliefs.
2. Moral judgments entail motivation.
3. Motivation involves a desire or pro-attitude.
4. There is no necessary connection between any belief and any desire or pro-attitude.

Different views in metaethics and moral psychology can be understood as responses, explicit or tacit, to the tension created by this quartet of claims, responses that avoid the tension by denying at least one of the constituent assumptions on the strength of others. **Noncognitivists** (e.g. Ayer, Stevenson, Blackburn, and Gibbard) reject (1), claiming that moral judgments express desires or other pro-attitudes, rather than beliefs. **Externalists** (e.g. Foot and Brink) reject (2), claiming that motivation is not internal or necessary to moral judgment, as such. **Rationalists about moral motivation** reject (3) or (4), claiming either that normative beliefs can motivate without the benefit of desires (e.g. McDowell) or that normative beliefs necessitate the desires that are necessary for moral motivation. (e.g. Nagel).

### SMITH'S MORAL PROBLEM

Michael Smith's book is focused on a related tension, which he calls "the moral problem". His problem consists of three claims (12).

1. Moral judgments of the form 'It is right that I  $\phi$ ' express a subject's beliefs about an objective matter of fact, a fact about what is right for her to do.
2. If someone judges that it is right that she  $\phi$ s then, ceteris paribus, she is motivated to  $\phi$ .
3. An agent is motivated to act in a certain way just in case she has an appropriate desire and a means-end belief, where belief and desire are, in Hume's terms, distinct existences.

It's worth noting some differences between Smith's problem and my puzzle.

- Smith's (3) combines the elements in my (3) and (4).
  - These two elements are independent. Because rejection of these two elements provides two distinct solutions to the puzzle, it is useful to separate them.
- Smith's version of internalism – his (2) -- is weaker than mine – my (2).
  - Several central figures (e.g. Price, Stevenson, Harman, Mackie, Blackburn) commit themselves to the stronger internalist thesis.
  - The trio/quartet is not inconsistent if we employ the weaker internalist thesis, which renders arguments that reject one element on the strength of others invalid.

## NONCOGNITIVIST SOLUTIONS

One way to motivate noncognitivism (e.g. emotivism, prescriptivism, expressivism) is as a solution to the puzzle about moral motivation that preserves (2)-(4). We can argue against cognitivism and for noncognitivism by appeal to internalism and belief-desire psychology. There are various possible noncognitivist views.

- Emotivism: moral judgments express (noncognitive) feelings.
- Emotivism: moral judgments express desires and aversions.
- Prescriptivism: moral judgments express (universal) prescriptions about how to feel or act.
- Expressivism: moral judgments express plans for action (actual or hypothetical).

This is reminiscent of Hume's argument in the Treatise that moral distinctions are not conclusions of reason.

Since morals, therefore, have an influence on the actions and affections, it follows, that they cannot be deriv'd from reason; and that because reason alone, as we have already prov'd, can never have any such influence. Morals excite passions, and produce or prevent actions. Reason of itself is utterly impotent in this particular. The rules of morality, therefore, are not conclusions of our reason [III.i.1/457].

Indeed, some writers have actually interpreted Hume as the father of contemporary noncognitivism.

1. Moral judgments entail motivation.
2. Motivation consists in pro-attitudes/desires.
3. By itself, no belief entails any pro-attitude/desire.
4. Hence, moral judgments express pro-attitudes/desires, rather than beliefs.

This argument for noncognitivism is worth considering. But it is not Hume's argument. Hume is a sentimentalist, not a noncognitivist. For Hume, moral judgments are true or false, and they have dispositional response-dependent truth conditions of the form

X is good just in case it is such as to elicit approbation in a suitable appraiser.

Or, to adapt ideas from Hume's Enquiry Concerning the Principles of Morals, we might specify a suitable appraiser in terms of the general point of view.

X is good just in case it is such as to elicit approbation from an appraiser adopting the general point of view.

The difference between Humean sentimentalism and noncognitivism is like the difference between subjectivism and noncognitivism. The emotivists (e.g. Ayer and Stevenson) were keen to distinguish their view from a more traditional form of subjectivism. Whereas subjectivists analyzed moral judgments as **reports** of the appraiser's noncognitive feelings or attitudes, emotivists analyzed moral judgments as **expressions** of these pro-attitudes. Moral judgments, on their view, express an attitude, pro or con (Hurrah! Or Boo!), toward people, actions, institutions, and policies; they do not ascribe any properties to those things. As such, moral judgments do not assert anything and are neither true nor false. They function more like exclamations or imperatives.

If moral judgments are, as the noncognitivist asserts, expressions of pro-attitudes, then there is a necessary connection between moral judgment and motivational states. Moral judgment may not be sufficient to produce associated behavior, because an agent may be subject to multiple and conflicting motivation that over-rides moral motivation. Thus, noncognitivism allows for the possibility of weakness of will. (Hare at one point committed himself to the idea that moral judgment expresses over-riding motivation, as a result of which he had to deny that genuine weakness of will is possible. But noncognitivism, as such, does not require this extreme commitment.) Noncognitivism does entail internalism. Because it also squares with belief-desire psychology, it might seem an attractive resolution of the puzzle about moral motivation.

### **WORRIES FOR NONCOGNITIVISM**

How well does noncognitivism fare in reconciling the intellectual and practical aspects of morality? Because it secures an internal connection between moral judgment and the will, it may seem to secure the practical dimension of morality. The natural question is whether noncognitivism can give a satisfactory account of the intellectual aspects of morality. Not surprisingly, this is where some find noncognitivism problematic.

Before getting to these intellectual worries, it's worth noting that one could doubt the practical adequacy of noncognitivism. It's true that noncognitivism delivers internalism and that internalism is one traditional way of interpreting the practical dimension of morality. But we can also ask whether it's the correct or best way to interpret this practical dimension. As we will see, some people (e.g. Foot, Brink, Railton, Sturgeon, Svavarsdottir) are externalists, who deny that there is a necessary connection between moral judgment and motivation. They owe us some plausible account of how there is a reliable or significant contingent connection between moral judgment and motivation or else some other quite different interpretation of morality's practical character. But if some form of externalism were a more plausible interpretation of the practical dimension of morality, then the fact that noncognitivism implies internalism would undermine, rather than vindicate, its account of this practical dimension.

The more obvious or traditional worry for noncognitivism concerns its account (or lack thereof) of the intellectual aspects of moral discourse and inquiry. Cognitivism would seem to be the more natural interpretation of moral discourse. The syntax of moral discourse is descriptive.

- It's wrong to torture babies for fun.
- She's a good parent.
- Justice is the first virtue of social institutions.

- The executive order allowing warrantless searches is unjust.

Moral judgments are expressed in the declarative mood and treat moral predicates as modifiers of noun and verb phrases. As such, they appear to express propositions ascribing moral properties to persons, actions, and institutions. Acceptance of these propositions seems to take the form of belief. By contrast, noncognitivism, at least in its more traditional forms, seems to be an error theory, or at least a revisionary theory. For it implies that moral judgments are not expressing beliefs, but are instead expressing noncognitive attitudes. If moral judgments are expressing noncognitive attitudes, then it seems they are not asserting propositions that might be true or false. But then the descriptive surface form of moral judgments hides a very different deep structure. This makes noncognitivism a sort of linguistic error theory. But error theorists assume the burden of proof. We ought to give up a cognitivist construal of moral discourse only if cognitivism can be shown to involve unacceptable commitments.

Rather than simply shift the burden of proof to the noncognitivist, the cognitivist might go on the offensive, arguing that there are significant intellectual aspects of moral discourse and inquiry that are only or best explained in cognitivist terms. Consider three such issues: disagreement, valid inference, and unasserted contexts.

## **DISAGREEMENT**

The cognitivist provides a natural account of the nature of moral agreement and disagreement. When A makes a moral judgment, she represents the world as being a certain way. B agrees with A when he makes a moral judgment that represents the world as being the same way. C disagrees with A when he makes a moral judgment that represents the world as not being that way. The cognitivist can explain the disagreement between A and C as consisting in the fact that their beliefs cannot both be true.

Historically, emotivists motivated their view by appeal to the problems that subjectivists had representing moral disagreement. According to the subjectivists, moral judgments report the appraiser's attitudes. That means that moral "disagreements" report the different feelings that different appraisers have toward one and the same situation. But, so the emotivists argued, there is no real disagreement here; different appraisers are just reporting different attitudes, as when one person says "I like Neapolitan ice cream" and another says "I don't". By contrast, the emotivists said that they could explain moral disagreement as a conflict in attitudes. On this view, moral disagreement is a practical conflict between two people with different attitudes or commitments. Those in moral disagreement are assuming rival practical postures.

But one may wonder whether this practical account of moral disagreement is plausible. One wants to agree that the subjectivist cannot represent moral disagreement, because there's nothing that prevents the different personal reports from both being true. Cognitive disagreement requires the assertion of mutually incompatible propositions. Mutual incompatibility is explained in truth-theoretic terms; incompatible propositions cannot both be true. But, of course, the traditional noncognitivist cannot understand disagreement as the commitment to mutually inconsistent propositions. There's nothing inconsistent about people assuming different practical stances toward one and the same action.

Moreover, there's a worry that if we understand moral agreement/disagreement as agreement/disagreement in noncognitive attitude, then we may be led to counter-intuitive verdicts about some moral agreements and disagreements.

- Moralists with different moral criteria (e.g. Bentham and Kant) might share a pro- or con-attitude toward a given action, yet not be in moral agreement about it.

But this argument may not be decisive. Though rival moralists may disagree about **why** an action has the moral valence it does, they need not disagree about what its valence **is**. They agree about its valence; their disagreement emerges about the valence of other cases, in which utility and the categorical imperative diverge.

- The moralist and amoralist may hold different attitudes toward an action, yet not be in moral disagreement.

But this worry takes sides in the internalist/externalist debate. Noncognitivists are internalists, and internalists deny the possibility of the amoralist (putative amoralists use moral language in inverted commas). Externalists recognize the possibility of amoralists. But then externalists already have reason to reject noncognitivism. But if we accept the possibility of the amoralist, then this worry about the inability to represent moral agreement among moralists and amoralists is a **further** way in which amoralism spells trouble for noncognitivism.

### VALID INFERENCE

There are valid inference forms involving moral claims, and our metaethical theories ought to be able to accommodate this fact.

1. It is wrong to murder innocent children.
2. Bart is an innocent child.
3. Hence, it is wrong to murder Bart.

The classical explanation of validity is truth-theoretic. In a valid argument, it is impossible for the premises to be true and the conclusion false. But the classical noncognitivist denies that moral judgments are either true or false; a fortiori, moral judgments cannot be inconsistent (in this sense). Perhaps there is some sort of pragmatic objection to affirming the premises of these arguments while denying their conclusions that it not itself parasitic on the truth-theoretic account of validity. But this needs to be shown, and, in any case, the noncognitivist apparently lacks this straightforward and compelling account of what makes these arguments good.

### UNASSERTED CONTEXTS

Another concern for the noncognitivist involves the phenomena of unasserted contexts. The noncognitivist construes moral assertion as the expression of the appraiser's attitudes, rather than as a description of the way the world is. But, as Peter Geach observed, the noncognitivist proposal fails to account for the meaning of moral predicates in **unasserted** contexts, such as the antecedents of conditional statements.

- If it is **wrong** to murder innocent children, then it is wrong to pay someone else to murder innocent children.

This conditional asserts something about and perhaps expresses an attitude toward paying someone to murder innocent children, but this is conditional on the wrongness of murder. The conditional

does not assert that murder is wrong. So it is hard to see how the appraiser could be expressing an attitude toward it. Because the first occurrence of the moral predicate 'wrong' is in an unasserted context, traditional forms of noncognitivism seem to have no account of its meaning. This is bad enough. But consider the following argument in which this conditional appears as a premise.

1. If it is **wrong** to murder innocent children, then it is wrong to pay someone else to murder innocent children.
2. It is wrong to murder innocent children.
3. Bart is a child.
4. Hence, it is wrong to pay someone else to murder Bart.

The argument certainly seems valid, and for this to be true all four occurrences of the predicate must have the same meaning. But then the predicate "wrong" must mean the same thing in both asserted and unasserted contexts. Geach attributes the general point about the univocity of terms across asserted and unasserted contexts to Frege. Whereas the cognitivist can accept Frege's point, the noncognitivist apparently cannot.

## **GIBBARD**

Someone might identify these as apparent **intellectual costs** of noncognitivism. Whereas traditional noncognitivists were typically either unaware of or indifferent to these costs, more recent noncognitivists, such as Gibbard and Blackburn, are concerned about these costs. However, they claim that these costs are avoidable. When properly articulated, noncognitivism has the resources to capture these and other intellectual aspects of moral discourse and inquiry every bit as well as cognitivist or realist views. Indeed, both Gibbard and Blackburn endorse **quasi-realism** – the view that all traditional realist claims about the logic, semantics, metaphysics, and epistemology of ethics can be vindicated within a view that restricts itself to expressivist essentials or at least that all such traditional realist claims have adequate surrogates within an expressivist conception.

A full-scale assessment of the adequacy of quasi-realism is not possible here. We might briefly consider Gibbard's expressivism and its resources for dealing with these intellectual costs. Gibbard's book *Thinking How to Live* is a sequel to his earlier book *Wise Choices, Apt Feelings*. His earlier book defended a conception of morality as involving distinctive moral feelings of guilt, shame, and resentment that it is rational for someone to have and went on to defend an expressivist conception of rationality according to which judgments of rationality involve acceptance of norms for behavior and feeling. Though Gibbard offered a novel conception of the semantics and logic of normative judgment, he spent much of that book discussing the evolution, psychology, and social dynamics of normative commitment. *Thinking How to Live* has a narrower focus. Though Gibbard regards morality as one form of normativity or practical reason, he largely avoids discussion of morality, focusing instead on normativity or practical reason. Here too, his focus is narrower, restricted primarily to issues about the semantics, logic, and epistemology of normative judgment. The new book reinterprets the earlier notion of accepting a norm in terms of **planning** or **commitment to a plan**. One consequence of this reinterpretation is to make the expressivist character of Gibbard's account of normative judgment clearer. He then proceeds to explain normative judgment in terms of commitment to a set of contingency plans, which allows him to defend an account of normative inference and reasoning. In fact, Gibbard argues that his expressivist conception of normative judgment allows him to

defend versions of various metaethical claims often thought to be the exclusive province of the practical realist.

I have been claiming triumph after triumph for a program of quasi-realism. Plans, and plan-laden judgments more generally, turn out in remarkable ways to mimic prosaic descriptive judgments. The predicates ‘is okay to do’ and ‘is the thing to do’ act much like ordinary, descriptive predicates, in a multitude of ways: judgments in terms of these predicates can be correct or incorrect. Standard logic applies. There is a natural property that constitutes being okay to do, and a naturalistic attribute that being okay consists in. Being okay to do can figure in causal explanations. We can even speak of a person’s “knowing” what to do, and of epistemic (or “quasi-epistemic”) virtues in planning and plan-laden judgment [251].

The result is a metaethically ecumenical form of expressivism that endorses many of the claims associated with the ethical intuitionist tradition and, more recently, with nonreductive forms of ethical naturalism.

If Gibbard can endorse such traditionally realist metaethical claims, in what sense is his view still expressivist? He regards his view as expressivist because of where it starts, not where it finishes (184). He thinks that he can work his way up to realist claims from resources that recognize only prosaic (nonnormative) belief and planning. Because Gibbard treats these resources as requiring nothing the expressivist cannot recognize, he thinks he can explain realist claims with expressivist resources (xii, 62-63).

To understand Gibbard’s expressivism and its ecumenical ambitions, it is necessary to say more about his account of normative inference and normative content. To make normative judgments is to make plans about actual or hypothetical situations. He appeals to the idealized notion of a **hyperplan**, which is a maximal contingency plan that prescribes a unique course of action and proscribes alternative courses of action for every conceivable situation. We might say someone has a **hyperbelief** if she is maximally opinionated about which situations she will find herself in. We might then describe a **hyperagent** as someone who combines a hyperbelief and a hyperplan. Of course, most of us have neither hyperplans nor hyperbeliefs. Instead we can represent our partial plans and beliefs as the sets of plans and beliefs consisting of disjunctions of plans and beliefs that are not ruled out by the plans and beliefs we accept. This allows Gibbard to characterize a state of mind that mixes fact and plan as “given by the hyperstates that it allows and the ones it rules out” (57). This picture of the content of an agent’s commitments appeals to the structure provided by possible world semantics to explain how there can be entailment relations among judgments that mix fact and plan. Gibbard applies this analysis to a disjunctive syllogism. Suppose Holmes has the contingency plan to pack now if it is not too late to catch the train.

1. Either packing now is what one ought to do [plan] or it is too late to catch the train [fact].
2. Even now it is not too late to catch the train [fact].
3. Hence, packing is what one ought to do now [plan].

Holmes’ contingency plan commits him to packing in all worlds in which it is not too late for him to catch the train. Once it is established that the actual world is one in which it is not too late to catch the train, it follows that the actual world is one in which Holmes must pack. In this

way, Gibbard claims that he can explain entailment relations among judgments that mix fact and plan. This, he claims, allows the expressivist to solve the so-called Frege-Geach problem about the content of normative judgments in various contexts.

### **WORRIES ABOUT GIBBARD'S ECUMENICAL EXPRESSIVISM**

There are various worries one might raise about the adequacy of Gibbard's expressivism for satisfying various cognitivist aspirations and claims.

For instance, consider Gibbard's account of expressivist inference. What kind of necessity can the expressivist attribute to the conclusion of valid arguments, such as these? The practical realist will offer a truth-theoretic explanation. It's not possible for the premises to be true and the conclusion false. But the expressivist must say that Holmes's plans commit him to the conclusion in the sense that they rule out any alternative; Holmes's contingency plan includes this practical conclusion. That means that Holmes is committed to the practical conclusion -- **unless he changes his plans**. But that means that the necessity of the practical conclusion is, for the expressivist, unlike the practical realist, **plan-relative**.

Indeed, plan-relativity seems to be a quite general feature of Gibbard's expressivist reconstructions of various realist claims. For instance, it seems to infect Gibbard's response to the Frege-Geach problem. Indeed, Gibbard responds to only one aspect of the Frege-Geach worry, the worry about valid inference. He doesn't reply to their more basic concern about the semantics of unasserted contexts. Consider our mixed conditional.

- If it is wrong to murder innocent children, then it is wrong to pay someone else to murder innocent children.

This conditional judgment may express an attitude toward paying someone to murder innocent children, although this attitude would be conditional on the wrongness of murdering innocent children. It expresses no attitude toward the wrongness of murdering innocent children. A related worry arises for Gibbard's analysis, I think. Perhaps the moral predicate 'wrong' in the consequent of the conditional picks out a conditional plan. But the natural way of understanding the conditional is that that plan is conditional on something that is not itself plan-relative, namely, whether murdering innocent children is wrong. Here, and in other unasserted contexts, moral predicates appear to have a meaning that is not plan-relative. If so, I don't think Gibbard plan-laden semantics for moral language has addressed this aspect of the Frege-Geach worry for noncognitivism.

We might also notice a worry about Gibbard's attempt to vindicate, within an expressivist framework, traditional realist claims about the way in which the normative supervenes upon and is constituted by the natural. Within any given hyperplan, there is a plan for every possible contingency. This allows us to say that the thing to do in any given situation supervenes upon a naturalistic description of the contingencies of that situation and that those contingencies are what makes the content of the contingency plan the thing to do in that situation. But this kind of plan-relativity is not built into the usual understanding of ethical naturalism. Consider the way most of us would understand how the injustice of apartheid supervened upon various laws, institutions, and practices of racial discrimination. We could also say that this naturalistic supervenience base of the injustice of apartheid constituted its injustice. In making these claims, we commit ourselves to two other claims – (a) that any social system with the same discriminatory laws, institutions, and social practices would also be unjust, and (b) that a social

system could not fail to be unjust unless its laws, institutions, and social practices were different from those of the apartheid system in some relevant way. The appraiser's plans do not figure in (a) and (b) or the supervenience base of moral properties, such as injustice. But then the sort of supervenience and constitution claims that ethical naturalists and other moral realists make appear not to be plan-relative, as Gibbard's analysis implies. Perhaps these are realist claims that the expressivist should eschew. That would leave Gibbard with a consistent position, but it would require him to forego his distinctive desire to defend a metaethically ecumenical version of expressivism.

It's also unclear whether Gibbard's account of moral agreement and disagreement is adequate. Gibbard must interpret moral disagreement as disagreement in plan; parties to a moral disagreement commit themselves to incompatible plans for a given situation. But there's no inconsistency in people adopting rival plans. We don't disagree when we form different plans about which college to attend in the way we seem to when we disagree about whether a certain sort of action is permissible or not.

Finally, I would note that there's something odd about identifying normative judgment with plans, as Gibbard does. Sometimes normative judgments do not affect one's plans, as when one makes normative judgments about other people, the past, or merely hypothetical circumstances. Gibbard represents such judgments as a form of contingency planning – planning what to do were one ever to find oneself in such circumstances. But that interpretation seems both strained and overly narcissistic. My judgments about the permissibility of slavery in ancient Greece or the antebellum South aren't contingency plans for what it would make sense for me to do were I to find myself in similar social circumstances. Even when normative judgments do affect one's plans, as when one decides upon a course of action after reaching a normative judgment, the planning appears to be **downstream** from the judgment. Plans are often made as the result of normative judgments about what one ought to do. It's easy for the practical realist to explain how planning is consequential on normative judgment, because she treats normative judgment as a belief that can and often does affect motivations and intentions. It's harder to see how Gibbard can explain how planning is consequential on normative judgment. Indeed, Gibbard's expressivism seems to have the same difficulty here that the behaviorist about mental states has explaining how there can be mental states that don't produce behavior and how even when mental states do cause behavior the behavior is downstream from the mental state.

## **REDUNDANCY ABOUT TRUTH**

Some think that the noncognitivist can address at least some these intellectual costs if she can avail herself of a conception of truth in ethics and that she can help herself to an account of truth by appealing to disquotational or redundancy accounts of truth or its ascription. However, we need to distinguish between two possibilities.

- Disquotational or redundancy accounts of **ascriptions** of truth
- Disquotational or redundancy accounts of **truth**

According to (a), to say that someone's moral judgment is true is to agree with that judgment.

'It is true that p' means the same as 'p'.

So, for example, an instance of this schema is:

'It is true that murder is wrong' means the same as 'murder is wrong'.

According to (b), truth for moral judgments involves no more than the following schema for moral sentences:

'p' is true iff p.

So, for example, an instance of this schema is:

'murder is wrong' is true iff murder is wrong.

I'm not sure how plausible (a) is. (a) makes a claim about the semantic equivalence of two different judgments or utterances. The second judgment or utterance is a first-order judgment or utterance, whereas the first is a second-order judgment or utterance that takes the first as its object. Perhaps the ground for asserting the two is the same, but it doesn't follow that their meaning is the same. In fact, their meaning would seem to be different insofar as the first judgment is about a judgment whereas the second is not. Whereas (a) strikes me as implausible, I see no special reason the noncognitivist cannot endorse it.

However, the situation with (b) seems quite different, and it is less clear how the noncognitivist can endorse it. In (b), the first half of the biconditional is a judgment or utterance, but the second half is neither. A disquotational account of truth for moral judgments seems to require moral states-of-affairs, or facts, or some such truth-maker for the second half of the biconditional. But the classical noncognitivist thinks that moral predicates are non-referring and so cannot recognize moral truth-makers. In fact, this worry is related to the worry about unasserted contexts. (b) involves moral predicates occurring in unasserted contexts, viz. the right-to-left conditional.

But if the noncognitivist is entitled at most to a disquotational account of ascriptions of truth and has no account of truth itself, then she lacks the resources necessary to explain disagreement or valid inference in truth-theoretic terms, for these explanations appeal to truth, not ascriptions of truth.

## **PROSPECTS**

I mention these problems for the noncognitivist, not because I am sure that they are insurmountable, but to point to costs or, at least, obstacles the noncognitivist solution to the puzzle about moral motivation faces. Before we abandon the default cognitivist construal of moral discourse and tackle semantic problems associated with noncognitivism, we should examine other solutions to the puzzle about moral motivation.