Philosophical Books Vol. 50 No. 2 April 2009 pp. 83–98

MORAL AND EPISTEMIC OPEN-QUESTION ARGUMENTS

CHRIS HEATHWOOD
University of Colorado at Boulder

As advertised in its subtitle, Terence Cuneo's rich new book, *The Normative Web*, puts forth an argument for moral realism.[1] Moral realism, Cuneo argues, is on a par with *epistemic realism*—roughly, the view that there are genuine epistemic facts, facts such as that it is reasonable to believe that astrology is false. These two realist doctrines stand or fall together. But epistemic realism is true. Therefore, so is moral realism. This is an important and widely held reason for thinking that moral realism might be true, and it deserves the book-length treatment Cuneo has provided. Readers will profit from studying it.

Cuneo puts his main argument roughly as follows:

P1. If moral facts do not irreducibly exist, then epistemic facts do not irreducibly exist.
P2. Epistemic facts irreducibly exist.
C1. So, moral facts irreducibly exist.
P3. If moral facts irreducibly exist, then moral realism is true.
C2. So, moral realism is true.

In a moment, I will review how Cuneo understands irreducibility, since others understand it differently. After I do this, I will move on to the main part of my comments. In the interests of full disclosure, let me say that I am attracted to moral realism of the sort Cuneo thinks is true. Nevertheless, I have my doubts about this particular kind of argument, the argument from epistemic realism. The purpose of this essay is to explain the main doubt I have about it, one that Cuneo does not explore much in the book.

My remarks will concern P1, the parity premise. My main goal is to investigate open-question-style arguments as they relate to moral and epistemic realism. I will attempt to demonstrate that the epistemic open-question argument (OQA) is less plausible than the moral OQA. Then I will explain how this makes trouble for the first premise of Cuneo's main argument. Along the way, I provide an additional argument for epistemic descriptivism, the thesis, to be introduced later, that makes trouble for the parity premise. First, I want to clarify what Cuneo means by "moral realism," since it is not exactly the thesis I was expecting.

1. T. Cuneo, *The Normative Web: An Argument for Moral Realism* (Oxford University Press, 2007, 263 pp. £37.50). This critical notice is a descendant of comments I presented in the author-meets-critics session on Cuneo's book at the Central Division Meeting of the American Philosophical Association in February of 2009.

What Cuneo Means by 'Moral Realism'

The core of the moral realist view, as Cuneo uses the term, is the thesis that moral facts irreducibly exist.[2] Cuneo is not here committing himself to the view that there are different kinds of existence or ways of existing. He just means that there are moral facts, and that these facts are not reducible to any other kind of fact.

If I were to say, "moral facts are not reducible" or "moral facts irreducibly exist," I would mean at least that they are not *identical* to facts of another kind—that is, not identical to facts that can be adequately expressed in non-moral language. But this is not what Cuneo means. He means that moral facts are not identical to facts of a kind that fail to satisfy the central platitudes of morality—that fail to make true our commonsense conception of morality. According to Cuneo, there are two such platitudes, a content platitude and an authority platitude (pp. 35–39). Simplifying somewhat, according to the content platitude, morality has fundamentally to do with human well-being. And according to the authority platitude, moral requirements are categorical, in that they give us reasons to act independent of our particular desires, interests, and goals (pp. 38–39). On Cuneo's usage, a meta-ethical view can be a form of nonreductionism about morality even if, according to it, moral facts are identical to facts of another kind. It will be a form of nonreductionism so long as the facts to which moral facts are identical do respect the central platitudes about morality.

So, for Cuneo, a reductive view of morality is one that fails to capture enough of our commonsense conception of morality. If someone held, for example, that it really is true that torture is morally wrong, but also held that the fact that torture is morally wrong just is the fact that torture is, say, illegal, then this would be a reductionist view of moral facts if the fact about legality did not capture enough of our shared commonsense beliefs about morality (perhaps such a view would run afoul of the authority platitude). But if someone held that moral facts are, say, just psychological facts (e.g., facts about what we are disposed to value), then this might *not* be a reductionist view of morality. It would not be reductionist so long as these psychological facts captured enough of our commonsense conception of morality.

Another thing worth mentioning about Cuneo's understanding of moral realism is that, as he defines it, it is no part of moral realism that moral facts are mind independent. If the moral fact that it is wrong to torture babies for fun is identical to the fact that the moral code of our society prohibits torturing babies for fun, moral realism can still be true, and it will be true, so long as this cultural relativistic theory preserves enough of the platitudes.

2. There are also linguistic and alethic components to moral realism, but they will not be relevant for our purposes. Also, Cuneo is, strictly speaking, defining 'moral realism of a paradigmatic sort,' rather than simply 'moral realism.' When I use 'moral realism,' we can understand this to be short for 'moral realism of a paradigmatic sort.'

Moral and Epistemic Descriptivism

Cuneo supports P1, the premise which requires a strong analogy between moral and epistemic realism, in two main ways. First, he identifies what he takes to be four important similarities between epistemic facts and moral facts (pp. 60–80). Second, he claims that the six main objections to the existence of irreducible moral facts apply just as strongly to epistemic facts, so that if we take these objections to give us reason to deny the existence of moral facts, we should also take them to give us reason to deny the existence of epistemic facts (pp. 89–112). These six objections are supposed to show that moral facts, were they to exist, would have certain *objectionable features*—features many think we should believe that probably nothing has. These are features such as supervening in a mysterious way, being epistemically inaccessible, being explanatorily idle, and others. I will discuss all of them a bit later. Cuneo claims that these features are no less true of epistemic facts, so that if we refuse to countenance moral facts on this basis, we should likewise refuse to countenance epistemic facts.

I want to explore a way to block this claim, by examining the extent to which it is more plausible that epistemic facts are identical to some class of descriptive facts (facts which will not display so many of the objectionable features) than it is that moral facts are identical to some class of descriptive facts.

The two relevant theses are

moral descriptivism: moral facts are identical to descriptive facts
and
epistemic descriptivism: epistemic facts are identical to descriptive facts.

Descriptive facts are facts that are expressed using descriptive terms. Some examples are the facts that snow is white, that Germany invaded Poland, that Bob is in pain, that I agreed to pay back the money I borrowed, and that people prefer happiness to misery. As some of these examples illustrate, all parties should agree that a fact can be a descriptive fact even if it is one that we think is morally *relevant*. The fact that Bob is in pain is, most people agree, a morally relevant fact, but it is nonetheless a mere descriptive fact about the world.[3]

In what follows, I will explain why it seems to me that epistemic descriptivism is more likely to be true than moral descriptivism. I will suggest a class of facts with which epistemic facts might plausibly be identified. I will try to show that the main obstacle to descriptivism—the OQA—makes trouble for moral descriptivism but not for epistemic descriptivism. Then I will explain why, if epistemic descriptivism is true, epistemic facts will not display most of the objectionable features that moral facts (about which descriptivism is not true) display.

3. The term 'descriptive' may not be ideal, for some who want to reject the spirit of moral descriptivism will be happy to say that moral statements *describe* reality. Some therefore prefer the term 'nonmoral fact' over 'descriptive fact,' and define a nonmoral fact as one that is expressed using nonmoral vocabulary.

The Moral OQA

I will begin with a sympathetic presentation of the OQA as applied to moral facts. Then I will describe an OQA as applied to epistemic facts and invite you to agree that this argument is far less compelling.

I believe that the best way to understand the idea of an *open question*—in particular, an open yes–no question like 'This is something we desire, but is it good?'—is in terms of whether one of the possible answers (yes or no) to the question would be self-contradictory, or incoherent. So I believe that, at bottom, Moore's OQA rests on intuitions about self-contradictoriness. So we may as well get down to brass tacks and state the argument directly in these terms.[4]

I will follow Moore in letting the focus of our moral OQA be the axiological notion of *goodness*, which appears in sentences of ordinary language such as 'Freedom is good,' 'The situation in Iraq is better than it was,' and 'The good things in life are free.' Consider some proposed definition or analysis of this notion, for example, 'good' (in the sense just specified) means the same as 'something we desire.' According to this analysis, if we say that freedom is good, what we are saying is that freedom is something we desire.

Now ask yourself whether the following claim is self-contradictory: *this is something we desire, but it's not good*. We are not asking whether such a claim is ever false. We are asking whether it is self-contradictory, or incoherent, in the way that the claim *this person is a bachelor, but he is married* is not merely false but self-contradictory. If we heard someone say it, we would presume that he was conceptually confused, and not a competent user of the relevant notions.

I believe that when we reflect on whether the statement 'this is something we desire, but it's not good' is self-contradictory, it seems to us—competent speakers of English, competent users of the term 'good'—that this statement, in all likelihood, is *not* self-contradictory. Thus we have:

M1. The statement 'this is something we desire, but it's not good' is not self-contradictory.
M2. If M1, then 'good' does not mean the same as 'something we desire.'
M3. Therefore, 'good' does not mean the same as 'something we desire.'

Thus, the analysis above is shown to be mistaken. Note that it might still be true—necessarily true—that all and only the things that we desire are good; it is just that it is not analytically true.

A.J. Ayer also states his version of the OQA in terms of self-contradictoriness. I will quote him at length, since doing so also serves to illustrate that the argument is meant to apply to *any* descriptive definition of any moral term:

4. I am not here concerned to be presenting any of Moore's actual arguments. For a recent interpretation of Moore's actual arguments, see F. Feldman, "The Open Question Argument: What It Isn't; and What It Is," *Philosophical Issues*, 15 (2005), pp. 22–43. My way of presenting the OQA is more similar to that of A.J. Ayer (see below).

> We reject the subjectivist view that to call an action right, or a thing good, is to say that it is generally approved of, because it is not self-contradictory to assert that some actions which are generally approved of are not right, or that some things which are generally approved of are not good. . . . And a similar argument is fatal to utilitarianism [put forth as an analysis, rather than a criterion, of moral rightness]. We cannot agree that to call an action right is to say that of all the actions possible in the circumstances it would cause, or be likely to cause, the greatest happiness, or the greatest balance of pleasure over pain, or the greatest balance of satisfied over unsatisfied desire, because we find that it is not self-contradictory to say that it is sometimes wrong to perform the action which would actually or probably cause the greatest happiness, or the greatest balance of pleasure over pain, or of satisfied over unsatisfied desire. And since it is not self-contradictory to say that some pleasant things are not good, or that some bad things are desired, it cannot be the case that the sentence "x is good" is equivalent to "x is pleasant," or to "x is desired."[5]

Upon observing this pattern, it is reasonable to conclude that no analysis of the moral in terms of the descriptive will be plausible. This is the conclusion of the moral OQA.

Importantly, the moral OQA is *not* intended to show that anything that has the descriptive feature in question (being generally approved of, being pleasant, etc.) might not have had the moral feature in question—that there is no necessary connection between the two features. Since all parties agree that the moral supervenes on the descriptive, we all agree that some moral term is necessarily coextensive with some descriptive term. The point is that these necessarily coextensive terms will not be synonymous. The descriptive term cannot serve as an adequate definition of the moral term. If someone claimed that the descriptive term applied to something that the moral term did not, we might correct this person's false moral belief, but we would not accuse this person of simply failing to understand the meaning of the terms involved.

Advocates of moral OQAs often explain why no descriptive definition of the moral can succeed by appealing to the fact that moral concepts are inherently normative or evaluative, while no merely descriptive concept is. Thus, any analysis of the moral in terms of the descriptive will leave something out. R.M. Hare, another advocate of the moral OQA, believes that no moral term means the same as any descriptive term because descriptive terms leave out the commendatory (or condemnatory) aspect of moral terms.[6] Moral assertions are inherently evaluative in a way that no mere description of things can be.[7]

5. A.J. Ayer, *Language, Truth, and Logic* (Dover Publications, Inc., 1946), pp. 104–5.
6. R.M. Hare, *The Language of Morals* (Oxford University Press, 1952), §5.3. See also A. Gibbard, *Wise Choices, Apt Feelings* (Harvard University Press, 1990), pp. 9–22.
7. This is not to deny that we cannot use descriptive terms to commend. A doctor who tells his/her patient that a certain procedure would be less painful might thereby be recommending it. But the remark functions as a recommendation only because it is common ground between the doctor and the patient that a less painful procedure is better. Descriptive terms are not *essentially* or *inherently* evaluative in the way moral terms are.

The moral OQA casts doubt on *analytic* descriptivism, the view that moral facts are identical to descriptive facts and that this identity can be discovered via conceptual analysis, or a priori reflection on the meaning of the sentences used to express these facts.[8]

Objections and Brief Replies to the Moral OQA

Some object that M1, the claim that 'this is something we desire, but it's not good' is not self-contradictory, begs the question, since it is an immediate implication of the target definition that M1 is false.

But M1 does not beg the question, because it is not just asserted; it is evidenced by the fact that it seems to us—ordinary, competent speakers of English—to be true. It is intuitively true. This is evidence that it is true. Note that the intuition that I am claiming has evidential value here is not even a moral intuition, which some may regard as less than reliable. It is a logico-linguistic (partly logical, partly linguistic) intuition, an intuition about whether a certain sentence of English is self-contradictory. Fewer philosophers are prepared to dismiss logico-linguistic intuitions, as it is not clear what other sources of evidence for claims of contradictoriness there could be.

Some object that the reasoning just stated behind M1 assumes that the full meaning of our words is transparently obvious to any competent English speaker. If that is what this reasoning assumes, that would be a problem; but it does not assume that. It assumes only that our intuitions about self-contradictoriness provide evidence of self-contradictoriness. It assumes only that we have *at least some* insight into the meaning of our words. Surely, no defender of analytic descriptivism, the view the OQA is attacking, could reject this.[9]

Some object that the OQA proves too much, since we could use it to refute any proposed definition of a term whose definition was not obvious, such as: 'knows' means the same as 'has a justified, true belief.' Imagine if Gettier had instead offered the following as the main premise in his argument: The statement 'this is a justified, true belief, but it's not knowledge' is not self-contradictory. This would have been a bad argument. But advocates of the moral OQA can agree that this is a bad argument without having to say that

8. This is in contrast to *synthetic* descriptivism, according to which, although moral facts are identical to descriptive facts, this can be known only a posteriori (just as, if water facts are identical to H_2O facts, this can only be known a posteriori). A different argument is therefore required to refute synthetic descriptivism. Since Cuneo rejects synthetic descriptivism about both the moral and the epistemic domains, I will not discuss these views much here. Incidentally, I reject these views as well, because I do not believe that moral and epistemic terms function the way that synthetic descriptivism about these domains would require them to function. See E. Gampel, "A Defense of the Autonomy of Ethics: Why Value Is Not Like Water," *Canadian Journal of Philosophy*, 26 (1996), pp. 191–210; D. Barnett, "Against A Posteriori Moral Naturalism," *Philosophical Studies*, 107 (2002), pp. 239–57.
9. Relatedly, we do not need to be able to produce the correct analysis of some term to know that certain proposed analyses could not be right. I cannot define "x causes y," but I am certain that it does not mean, say, "x is taller than y."

their own argument is a bad argument. For the main premise of this Gettier OQA is not intuitively compelling. Maybe the statement 'this is a justified, true belief, but it's not knowledge' *is* self-contradictory. Or at least this would have been reasonable to think before we were exposed to the Gettier counterexamples of this analysis. The point is that our linguistic and logical intuitions do not speak to this; we do not have the same sort of intuitions about this statement as we do about the moral statement. Since we know about Gettier's counterexamples, we could of course *reason* to the claim that 'this is a justified, true belief, but it's not knowledge' is not self-contradictory. But we do not have an immediate logico-linguistic intuition, prior to that reasoning, that it is.

Let me put the point another way. Consider the following definition of 'circle': 'circle' means the same as 'the set of all points equidistant from a single point.' Now consider this claim: The statement 'this is the set of all points equidistant from a single point, but it's not a circle' is not self-contradictory. This claim is not intuitively compelling. Maybe this statement is self-contradictory, maybe it is not. But our logico-linguistic intuitions do not speak to it. The moral OQA is different from OQAs about these other topics, I submit, because the evaluative and the descriptive just seem like two very different categories. The categories *sets of points* and *shapes* do not.

This third objection is especially tempting if the moral OQA is presented in the way Moore presented it, by asking us to compare two questions (or other sentences). This way of presenting it can seem to make the argument trade on the fact that, given the target analysis, a sentence like 'this is something we desire, but it's not good' should feel *just as* self-contradictory as the sentence 'this is good, but it's not good.' *This* sort of argument would prove too much. 'This is a brother, but it's not a brother' does not have the same feel as 'this is a male sibling, but it's not a brother,' even assuming that we know the corresponding analysis to be correct. But the OQA as I presented it did not make use of this sort of test.

The Epistemic OQA

For the epistemic OQA, I want us to focus on a certain kind of definition of epistemic normativity—of terms like 'justified belief,' 'reasonable belief,' 'has reason to believe,' etc. The basic idea is that these notions can be defined in terms of probability, or what is probably true. For reasons that will become clear shortly, we do not need to get hung up on the precise details of the definition. But one place to start is with a definition such as the following:

> 'p is reasonable for S to believe' means the same as 'p is likely, given S's information.'

So when we say that it is reasonable for us to believe that astrology is false, *all we mean*, on this view, is that this belief is likely to be true, given the information available to us. A closely related idea is that when we say, for example, that your having an experience as of a hand in front of your eyes *gives*

you reason to believe that there is a hand in front of your eyes, all we mean is something like that your having this experience *makes it likely* that there is a hand in front of your eyes.[10]

You might have all sorts of quibbles concerning the extensional adequacy of this definition. But, for now, suppose that the definition is extensionally adequate (i.e., that in all and only the cases in which it is reasonable for a person to believe something, that thing is likely to be true, given the information available to that person), and consider:

E1. The statement 'this is likely, given my evidence, but it's not reasonable for me to believe it' is not self-contradictory.
E2. If P1, then 'reasonable for me to believe' does not mean 'likely, given my evidence.'
E3. Therefore, 'reasonable for me to believe' does not mean 'likely, given my evidence.'

This argument, I hope you will agree, is not so compelling. The sentence 'this is likely, given my evidence, but it's not reasonable for me to believe it' does have an air of incoherence about it in a way that axiological statements—even such patently false ones like 'suffering is intrinsically good'—never do.[11] Perhaps my idea here can be made more vivid by considering a concrete case. Suppose I am having a visual experience as of a table in front of me. Suppose this in fact makes it very likely that there is a table in front of me. Now I say, 'Yes, I see that it's quite likely to be true that there is a table in front of me, but, still, I don't think it's reasonable for me to believe that there is a table in front of me.' This is a puzzling thing to hear. It seems to be grounds for thinking that I do not really understand what I am saying. At the very least, surely we *lack* the intuition that such a remark is *not* self-contradictory, even if we do not, in addition, positively *have* the intuition that it *is* self-contradictory. This is not at all how it is in the moral case, where we positively have the intuition that the axiological statement in question is not self-contradictory. This suggests that it might be true that all that reasonable belief amounts to is likely truth and that analytic descriptivism is true of epistemic normativity.

I can get the sentence in question (the one in E1 above) to seem coherent if I interpret 'reasonable' to stand not for epistemic reasonability, but for some other kind of reasonability, such as *prudential* reasonability. If my doctor tells me that there is a 90% chance that I have only 6 weeks to live, it is true that, given my information, it is likely that I have only 6 weeks to live; but still, it may be reasonable for me *not* to believe this (since if I refuse to believe it, I will be

10. The idea that epistemic facts are descriptive facts about probabilities is defended in C.S. Jenkins, "Epistemic Norms and Natural Facts," *American Philosophical Quarterly*, 44 (2007), pp. 259–72, though, importantly, Jenkins explicitly embraces the epistemic OQA, opting therefore for a form of synthetic epistemic descriptivism. Analytic epistemic descriptivism of the sort advocated here is endorsed in S. Street, "Evolution and the Normativity of Epistemic Reasons," *Oxford Studies in Metaethics*, 6 (forthcoming), at least for one sense of epistemic claims.
11. I mean axiological statements connecting an evaluative property to some descriptive property.

happier in my remaining days). But this is irrelevant to evaluating E1, since, of course, 'reasonable' in that premise stands for epistemic reasonability. It would still be *epistemically* irrational not to believe what the doctor has told me, even if it would be prudentially rational.

Here is a crucial point, to which I alluded earlier. It does not matter if the analysis of epistemic reasonability suggested above is exactly right. Maybe this definition is not quite extensionally adequate and needs some tinkering. The point is just that *the epistemic OQA itself* does not show the analysis to be inadequate. In the moral case, we could present a plausible OQA even against an analysis of 'good' *that we knew to be extensionally adequate*—in that, necessarily, all and only the things that are good have the descriptive feature specified in the analysis. If hedonism is the correct axiological theory, then the predicates 'is intrinsically good' and 'is an episode of pleasure' are necessarily coextensive, but it is still plausible that the statement 'this episode of pleasure is not intrinsically good,' though false, is not incoherent. Everyone agrees that there is some such equivalence between the descriptive and the evaluative. But we are still inclined to reject this equivalence *as an analysis*, because of the moral OQA.

With the epistemic OQA, things are different. Since we lack intuitions of self-contradictoriness in the epistemic case, we can draw a general conclusion: That descriptive analyses of epistemic notions do not leave anything out. Once we find—by doing substantive epistemology in the traditional way—the descriptive predicate that is necessarily coextensive with the epistemic predicate, we can plausibly hold that these predicates mean the same thing and express the same fact. No epistemic OQA will compel us to think, as one such argument does in the moral case, that we have two distinct, necessarily coextensive properties on our hands.[12]

This point is crucial because if it was otherwise, I would have to get into the business of defending a fully adequate analysis of epistemic reasonability. That is a worthwhile project, but would require at least its own paper. So if you think the analysis above is not quite right, that is okay, so long as you agree that its problem is *not* the same kind of problem the moral OQA is meant to bring to light concerning *its* target analysis. Here is yet another way to put the point. Ask yourself this: *On the assumption* that the predicates 'is reasonable to believe' and 'is likely, given the information available' are in fact necessarily coextensive, does an *analysis* of the former in terms of the latter work? Or, does it instead leave something out, and fall prey to the epistemic OQA. In the moral case, I believe we have a strong intuition that something is left out; in the epistemic case, I believe we do not.[13]

12. Well, that is not the only thing of which the moral OQA could convince someone. The failure of moral descriptivism might convince us, as it has others, to be noncognitivists or error theorists instead.
13. The picture for epistemology is therefore this. Do substantive 'normative' epistemology to figure out, for any epistemic term, which descriptive term is necessarily coextensive with it. Then put forth this descriptive term as a *definition* of the epistemic term. In this way, traditional (normative) epistemology shapes the semantic component of the metaepistemological view. Some philosophers hold similar views about the relation between normative ethics and meta-ethics, though these philosophers tend to be *synthetic* descriptivists, and so tend to offer the descriptive predicate (the one they discover, through substantive normative ethical theorizing, to be necessarily coextensive with some moral predicate) as a 'synthetic definition' of the moral predicate. See, for example, N. Sturgeon, "Moral Explanations," in D. Copp and D. Zimmerman (eds.), *Morality, Reason, and Truth* (Rowman and Allanheld, 1985), pp. 49–78, p. 61; D. Brink, *Moral Realism and the Foundations of Ethics* (Cambridge University Press, 1989), p. 176.

I conclude that there is an important disanalogy between moral and epistemic facts. Since the epistemic OQA fails against the kind of analysis above, we should be optimistic that epistemic facts just are facts about likelihoods. And since the moral OQA has force against any descriptive analysis of moral terms, we should be pessimistic about the possibility that moral facts, if they exist, are identical to descriptive facts.[14] In a way that will be explained shortly, this casts doubt on P1 of Cuneo's main argument, which requires a parity, or strong analogy, between the moral and the epistemic.

An Objection: Is Probability Itself Normative?

Perhaps the reason the epistemic OQA fails against its target analysis is that its target analysis is not in fact analyzing away the epistemic notion in the analysandum. Perhaps the key notion in that analysans—the notion of a proposition's being likely—is itself not merely descriptive, but inherently (epistemically) normative. If it is, then, although the epistemic OQA would fail, its failure would be no evidence for epistemic descriptivism, since the surviving target analysis would not be a descriptive analysis. It would be tantamount to analyzing the notion of goodness in terms of, say, the notion of what it is *fitting* or *appropriate* to want; if a normative analysis of goodness such as this fails, it is not because of open-question considerations (such an analysis, whatever its defects, is at least not leaving out the normative dimension of goodness).

How might probability be itself a normative notion? The most natural way is enshrined in a version of the subjective interpretation of probability: A proposition, p, is likely to be true given the available information just in case it is reasonable to believe p, given that information.[15] If this is the correct theory of the nature of probability and if the notion of epistemic reasonability contained therein is not amenable to some descriptive analysis, then this would undermine the main argument of this paper.

14. As has been indicated, there is room for synthetic descriptivism, the view that moral facts are identical to descriptive facts, but that moral terms do not mean the same as any descriptive terms. I have alluded (note 8 above) as to why I do not believe that this view will succeed. I should now note that Cuneo himself already agrees that moral facts are not identical to any descriptive facts (pp. 220–1). In the part in which he says this, he talks of the "grain of truth. . . in Moore's open question argument" (pp. 221, 56n). Thus, Cuneo already accepts not just that moral terms are not synonymous with any descriptive terms, but that moral facts are not identical to any descriptive facts. Cuneo would of course want to say exactly the same thing for epistemic facts. My point here has been to show that when we examine the open-question case for nondescriptivism about epistemic facts, it is much less strong.
15. Another version could be formulated in terms of what an epistemically rational person would believe if he/she had that information. A better formulation of either of these theories of the nature of probability would contain *degree* indexes. It would explain the degree of likelihood in terms of the degree of belief that it is reasonable to have.

But I do not believe that this is a plausible understanding of probability. Its main problem is that it founders on an epistemic variant of Plato's Euthyphro problem. Suppose we look out the window and see that the streets are wet. We conclude that it probably rained. We also think that it is reasonable to believe that it rained. But which seems right: That it is true that it probably rained in virtue of the fact that it is reasonable to believe that it rained, or that it is reasonable to believe that it rained in virtue of the fact that it probably rained? Surely, the latter is correct. Facts about what is probably true are more basic, and indeed, we use these facts in deciding what to believe (i.e., in deciding what is reasonable to believe). We explain why it would be reasonable to believe something by establishing that that thing is probably true, and not the other way around.[16]

What makes some probability claim (one concerning whether it has rained, say) true is intuitively not the attitudes of rational observers but objective patterns in the world. There is this connection in the world, independent of our beliefs, between streets being wet and its having rained. This is what makes it true that if the streets are wet, it has probably rained.

That probability claims are made true in this way is especially compelling when one considers a kind of limiting case of the makes-probable relation—that of entailment. Given that Xena is older than Yolanda and that Yolanda is older than Zooey, it is guaranteed to be true that Xena is older than Zooey. And if something is guaranteed to be true, then, a fortiori, it is probably true. Thus, we have the following claim:

> Given that Xena is older than Yolanda and that Yolanda is older than Zooey, Xena is probably older than Zooey.

What makes this true? The answer would seem to be the following: The claim above is true because of the fact that Xena's being older than Yolanda and Yolanda's being older than Zooey *entails* that Xena is older than Zooey (and that, if P entails Q, then Q is probable, given P). Thus, at least, when some evidence entails some conclusion, that the conclusion is probably true given the evidence is something that is made true by some objective feature of reality (viz., facts about entailment). It would be surprising if we said something radically different, and subjective, for cases in which some evidence does not entail, but merely strongly supports, some conclusion.

When we say that something is probably true, given some other thing, it is admittedly not obvious what we are saying. But intuitively, what we are saying is something *descriptive*. It is a mere matter of fact about the world. It is made true not by the rational attitudes of subjects, but by objective patterns in the world. It belongs on the fact side of the fact–value divide. When Hume said that you cannot derive an *ought* from an *is*, and this conjures in us an idea of the distinction between an *ought* and an *is*, we intuitively classify claims such as "It's unlikely to rain tomorrow" as belonging on the *is* side.

16. To make the presentation of this argument less clunky, I have omitted the modifier "given the available information."

Interlude: A Further Reason to Accept Epistemic Descriptivism

Before explaining what the upshot of the difference in plausibility between the moral and epistemic OQAs is for Cuneo's main argument, I want to take us away from open-question considerations and present an independent argument for epistemic descriptivism, and in particular, for the view that facts about epistemic reasonability just are facts about likelihoods. The argument has to do with the following interesting disanalogy between reasonable belief and reasonable action. If we believe that some belief of ours is not reasonable, we stop holding it. Even if there can be exceptions to this rule, they are rare. We can put the rule another way by saying that there is very little room between holding a belief and thinking it reasonable. But there is much more room between doing an act and thinking it rational. It is depressingly common that I see that some act of mine is not, say, in my long-term interest, but I cannot help myself and I do it anyway. What explains this? Why is there a much tighter connection between believing that some belief of ours is not reasonable and ceasing to hold it than there is between believing that some available act would not be reasonable and refraining from doing it?

The view that to believe that some belief is rational just is to believe that it is likely to be true can explain this. For if believing that some belief of mine is irrational amounts to believing that it is unlikely to be true, then believing that some belief of mine is irrational is the same as believing that the object of the belief is unlikely to be true. And that is the same as disbelieving it. This is why belief about the irrationality of a belief that p entails, or at least is very likely to accompany, disbelief in p.

The failure of the epistemic OQA suggests that epistemic descriptivism is true. The argument in this section provides further reason to think that epistemic descriptivism is true. Neither of these reasons apply to moral facts.

The Upshot of this Disanalogy

If epistemic facts are identical to descriptive facts, but moral facts are not, this has serious implications for the popular idea that we can argue for moral realism on analogy with epistemic realism. In particular, if epistemic descriptivism is true, then epistemic facts exhibit far fewer of the objectionable features that Cuneo and others claim will plague epistemic facts to whatever extent they plague moral facts. Let us go through five such features:[17]

- Supervenience. Cuneo describes the problem concerning supervenience as follows:

  > . . . the problem is that, although realists freely speak of mere descriptive facts 'determining' moral facts, they offer no plausible account of what

17. I leave out the discussion of "intrinsic motivation" (pp. 92–98), since I do not believe that epistemic descriptivism affects whether this objectionable feature would attach to epistemic facts.

> this determination relation is. The relation cannot be a causal one, for that would be too weak. . . . Nor could the relation in question be one of identity, for that would be too strong. It would leave out the normative or prescriptive dimension of moral facts by identifying them with merely descriptive ones. At best, the relation appears to consist in moral facts being bound to mere descriptive facts, in Hilary Putnam's phrase, by some sort of mysterious 'metaphysical glue.' But, on the plausible assumption that there are no such mysterious determination relations, it follows that there are no moral facts. (p. 91)

I agree that this is an important problem concerning moral facts. But if, unlike moral facts, epistemic facts *are* identical to descriptive facts, there is no corresponding problem for epistemic facts.

Cuneo says that for moral facts, the relation in question cannot be identity, because that would leave out the normative dimension of moral facts by identifying them with descriptive facts. I agree that this is what would happen in the moral case. But Cuneo holds the same for the epistemic case:

> To identify the epistemic with the merely descriptive would be to omit the normative or prescriptive dimension of epistemic facts. (p. 92)

But I think that our exploration of the epistemic OQA above shows that this is not the case. Epistemic normativity *is* definable in terms of the descriptive in a way that moral or axiological value is not. As the failure of the epistemic OQA is made clear, identifying the epistemic with the descriptive seems not to leave anything out.

- Authoritativeness. Attributing this antirealist objection to J.L. Mackie, Cuneo writes,

> by claiming that [moral] facts are authoritative, Mackie meant to say that they would have to be the sort of thing that implies, first, that there are moral reasons to conduct ourselves in certain ways and, second, that these reasons inescapably govern our conduct or are categorical. . . . whether I have a moral reason to attempt to stop the occurrence of a wicked deed doesn't depend on whether I have certain cares that would be satisfied by attempting this, or whether so attempting would further my interests, or whether I enjoy a certain type of social standing. But nothing—so the argument goes—answers to this description of a moral fact. If there are any reasons, then they are hypothetical in character; they apply to agents only insofar as they have cares, interests, projects, or social standings of certain kinds. It follows that moral facts don't exist. (p. 99)

Epistemic descriptivism's effect on whether this objectionable feature—authoritativeness—would attach to epistemic facts is not completely obvious. But here is how things seem to me. First, let me say that I agree that, despite the

truth of epistemic descriptivism, epistemic facts are still authoritative, or categorically reason giving. But if epistemic descriptivism is true, this categoricity may be less objectionable than the categoricity of moral reasons, about which descriptivism is false. If epistemic normativity is characterizable in purely descriptive terms, this suggests that, in some sense, epistemic normativity is not *real* normativity. If some belief is unreasonable, we do not *really have* to stop holding it. True, it would be epistemically irrational not to, but this is to say no more than that the belief is not likely to be true; it is not to say that we *really must* stop believing it. In other words, epistemic reasons are *faux* reasons (they are on a par with, say, "etiquettical reasons"). But still, these faux reasons that epistemic facts imply are *categorical* faux reasons. When it is true that it is epistemically reasonable for a person to believe something (in other words, that thing is likely to be true given his/her evidence), the reason here holds independently of his/her desires, interests, goals, etc. Epistemic reasons are in this way categorical. So the platitude (if it is a platitude) that epistemic reasons are categorical is not violated. Cuneo might insist that it is violated, since these are not *real* categorical reasons. But since, if epistemic reasons are faux, it is the failure of the epistemic OQA that reveals this, it cannot be a platitude that epistemic reasons are both categorical *and* real. If this was a platitude, the epistemic OQA would be more compelling.

Here is one way to try to flesh this out: The terms 'reason,' 'reasonable,' etc., are ambiguous. On one reading, they involve genuine normativity: If you have reason to do something, and no other reasons are in play, then you *really have* to do it. On another reading, such as the epistemic reading, these expressions are definable in non-normative terms; if you have an epistemic reason to believe something, you do not *really have* to believe it. All we are saying rather is that, that thing is probably true, given the available information. My contention is that the failure of the epistemic OQA suggests that ordinary ascriptions of epistemic reasons take the latter reading. Epistemic reasons really are reasons, if only because it is appropriate to use the word 'reason' to stand for them. But these reasons are not really binding, or normative. They are, nevertheless, categorical, in that whether a person has an epistemic reason to believe something does not depend upon the person's desires, interests, and goals. That they are categorical in this way simply follows from the fact that whether a proposition is likely to be true, given some person's available information, does not depend upon the person's desires, interests, and goals.

- Epistemic inaccessibility. Here is how Cuneo describes the problem:

> . . . we don't have any explanatorily informative story about how we could gain epistemic access to moral facts. That is, we do not have any informative account of how facts about what morally ought to be the case impinge on our cognitive faculties so as to produce the corresponding states of knowledge. . . . In light of this failure, it is best to conclude that there is no explanation available. But on the assumption that, if moral facts were to exist, then some such explanation would be available, it follows that moral facts do not exist.

> I think it should be admitted that this argument poses a serious challenge to moral realism. (pp. 100–1)

I agree. Cuneo goes on to claim the following, however:

> But I take it to be fairly plain that the argument poses an equally serious challenge to epistemic realism. (p. 101)

I agree that *if* descriptivism was as untenable in the epistemic realm as it is in the moral realm, then this problem would pose an equally serious challenge to epistemic realism. But if, instead, epistemic facts are identical to descriptive facts, as the failure of the epistemic OQA suggests, then the problem vanishes, because we can know epistemic facts by knowing the descriptive facts to which they are identical. We can know what we have reason to believe by coming to know what is likely to be true. Moreover, we can also come to know *which* descriptive facts epistemic facts are identical to (such as facts about likelihoods). We can know this via conceptual analysis.

- Explanatory idleness. Cuneo describes this problem as follows:

> . . . were moral facts to exist, at least the following must be true: Such facts would have to explain the existence of non-moral facts of a sufficiently wide array of types. . . . an ordinary descriptive fact such as *that the rocks are wet* can explain the existence of facts of many different sorts. . . . By contrast, . . . a putative moral fact such as *that an act is wrong* does not directly explain any states of affairs in the world that do not involve propositional attitudes; arguably, it doesn't do anything at all. If this is right, however, then the putative moral fact *that an act is wrong*—and presumably any other putative moral fact—fails to satisfy [this explanatory requirement]. From which it follows that moral facts do not exist. (p. 104)

I believe that considerations in this neighborhood provide an important challenge to moral realism, if descriptivism is false about moral facts. Cuneo claims that:

> when suitably modified, the argument we're now considering (if sound) yields a similar conclusion regarding epistemic facts. (p. 105)

But not if descriptivism is true of epistemic facts, as the failure of the epistemic OQA suggests. If epistemic descriptivism is true, epistemic facts will play explanatory roles, because the descriptive facts to which they are identical will.

- Lack of agreement. Cuneo describes the familiar idea that:

> . . . there is massive and seemingly ineradicable disagreement among otherwise competent parties concerning certain kinds of moral issues.

He goes on to advance the original and provocative idea that there is also ineradicable disagreement over epistemic matters, particularly over the extent to which religious belief is epistemically justified.

Epistemic descriptivism has less obvious of a role to play in showing epistemic facts to be disanalogous with respect to this allegedly objectionable feature. Nonetheless, I will close this section by mentioning one way in which epistemic descriptivism might be thought to help eradicate seemingly ineradicable disagreement over epistemic matters. One could resolve the disagreement by first engaging in conceptual analysis on the epistemic normative notions (as mentioned in the previous subsection). Then we could apply the resulting analysis to the problem cases. This would be to do in the epistemic domain exactly what G.E. Moore was at such pains to establish, through his open-question argument, that one could *not* do in the moral domain:

> . . . if I am right, then nobody can foist upon us such an axiom as that 'Pleasure is the only good' or that 'The good is the desired' on the pretence that this is 'the very meaning of the word.'[18]

Since the epistemic OQA fails, perhaps one can foist epistemic axioms such as that "What's reasonable to believe is what's probable" on the grounds that this is the very meaning of the word.[19]

Conclusion

Cuneo's aim is to establish moral realism by appeal to a putative analogy with epistemic realism. My goal has been to raise a problem for this analogy. I have tried to show, via consideration of OQAs, that it is more likely that epistemic facts are identical to descriptive facts than it is that moral facts are identical to descriptive facts. I have also tried to show why, if in fact epistemic but not moral descriptivism is true, most of the main objections that turn people away from moral realism will not apply to epistemic realism. I have thus tried to call into question P1 of Cuneo's main argument, the idea that if you are an epistemic realist (as most people are), you should also be a moral realist.[20]

18. G.E. Moore, *Principia Ethica* (Cambridge University Press, 1903), §6.
19. Admittedly, it is not clear how powerful this strategy would be in resolving disputes over what is reasonable to believe, for such disputes would surely just be replaced by disputes over what is probably true. Still, perhaps this is some progress.
20. For helping me improve this paper by giving me comments on earlier drafts, I am grateful to Matt Bedke, Eric Chwang, Bob Hanna, Tom Metcalf, and Michael Rubin.