



Christos Kyriacou • Robin McKenna  
Editors

# Metaepistemology

Realism and Anti-Realism

## Epistemic Reductionism and the Moral- Epistemic Disparity

Chris Heathwood

Reductionist doctrines about normative and evaluative phenomena enjoy serious advantages, such as in explaining how we can come to know about normative reality, in explaining why the normative depends on the non-normative, and in avoiding the spectre of Ockham's razor.<sup>1</sup> Unfortunately, some evaluative phenomena resist reduction. This is true, in my view, of moral and axiological facts. When we say that people ought to be more kind, or that things would be better if they were, it does not appear that we could report these same facts using non-normative, non-evaluative language. But things are different, I believe, when it comes to epistemic facts. When we say that someone is justified in believing something, we can report that same fact using non-normative language. Reductionism in metaepistemology is more plausible than reductionism in metaethics.

This moral-epistemic disparity is not only interesting in its own right, it shows that an important and popular line of argument for a robust

---

C. Heathwood (✉)  
Department of Philosophy, University of Colorado Boulder,  
Boulder, CO, USA  
e-mail: [heathwood@colorado.edu](mailto:heathwood@colorado.edu)

palgrave  
macmillan

© The Author(s) 2018  
C. Kyriacou, R. McKenna (eds.), *Metaepistemology*, Palgrave Innovations  
in Philosophy, [https://doi.org/10.1007/978-3-319-93369-6\\_3](https://doi.org/10.1007/978-3-319-93369-6_3)

moral realism does not succeed. According to this argument, there is a kind of moral-epistemic parity: epistemic facts and moral facts share deep similarities, so that whatever account is most plausible of the nature of one will be most plausible of the nature of the other. Furthermore, the argument continues, a robust realist account of epistemic facts is most plausible. Consequently, we have good reason to accept robust moral realism. In my view, this argument founders on the fact that whereas we are forced to reject a reductive naturalist account of moral facts, we are not so forced when it comes to epistemic facts. Committed as I am to robust moral realism, I don't welcome this result. But I believe that this is where the arguments lead.

I have defended epistemic reductionism and the moral-epistemic disparity in earlier work (Heathwood 2009). The purpose of the present essay is to clarify and refine this position and to defend it against recent criticisms by Richard Rowland (2013) and Terence Cuneo and Christos Kyriacou (2018).<sup>2</sup>

## 1 Preliminaries

I feel quite confident in certain ethical claims, some to the effect that some things that people do are not ok, others to the effect that some states of affairs are good in themselves. Consequently, I believe that there are such properties as those of being morally wrong and being intrinsically good. But what sort of properties are they?

According to *reductionism* in metaethics, moral facts and properties are identical to facts and properties from some other domain—facts and properties that can be adequately captured, expressed, or stated using wholly non-moral vocabulary. Most reductionist theories in ethics reduce ethical facts and properties to natural—that is, empirical—facts and properties. Psychological properties are the most common. Hume's famous remark that “when you pronounce any action or character to be vicious, you mean nothing, but that from the constitution of your nature you have a feeling or sentiment of blame from the contemplation of it” (1975, 469) is suggestive of a reductive naturalist account of the moral property of viciousness. A theory according to which the property of

being intrinsically good just is the property of being a state of pleasure is another example of reductive naturalism in metaethics.

*Non-naturalists* in ethics, by contrast, hold that moral properties are their own kind of thing, that they cannot be reduced to or identified with any natural or descriptive property.<sup>3</sup> Each moral property is what it is and not another thing. Crucially for our purposes, this is not to deny that there are necessary connections, and indeed explanatory ones, between moral and non-moral facts. Virtually everyone believes in these. So, for example, some non-naturalists about intrinsic value are hedonists. They hold that whenever a thing is intrinsically good, this is in virtue of the fact that that thing is a state of pleasure. They hold that the property of being intrinsically good is necessarily coextensive with the property of being a state of pleasure.<sup>4</sup> But they hold that this connection is a synthesis of two distinct phenomena rather than, as reductionists hold, a single phenomenon with two different, equally legitimate characterisations, one evaluative and the other descriptive.<sup>5</sup>

I should also make clear how I understand the moral realism/anti-realism distinction. As I use these terms, moral realism is the doctrine that there are some moral facts and that at least some of the moral properties involved in these facts are “stance-independent.”<sup>6</sup> A property is stance-independent just in case whether a thing has the property does not depend upon the attitudes or practices that certain (possibly hypothetical) observers take up towards that thing. Thus, the Humean view above is not a realist view, since it makes viciousness stance-dependent, while the hedonist view above is a form of realism.<sup>7</sup> Although the moral realism/anti-realism distinction cuts across the moral reductionism/non-reductionism distinction, virtually all non-reductionists are realists. I use the term “robust moral realism” to refer non-reductionist moral realism of a certain sort: non-naturalist moral realism.

Corresponding to these categories in metaethics are analogous categories in metaepistemology. Just as there are properties like that of an act's being morally wrong or a state of affairs' being intrinsically good, there are the properties of having a reason to believe something or of a belief's being epistemically justified or reasonable. About these epistemic properties, the same questions arise. Can they be reduced to properties from some other domain? Are they stance-independent?

Some of us are occasionally gripped by moral scepticism. We are visited by genuine doubts about whether anything anyone does really is forbidden or whether any states of the universe really are any better or worse than any others. Such doubts are rare, but even rarer are doubts about whether any *beliefs* are any more or less *reasonable* than any others. For this reason, if there is moral-epistemic parity, the near irresistibility of the idea that there are epistemic facts can be used to quell doubts over the existence of moral facts.

## 2 The Moral Open Question Argument

It seems to me that moral properties and facts resist analysis in terms of natural facts and properties, or any other facts and properties that can be adequately captured in non-moral terms. My belief is based on open-question-type considerations. The Open Question Argument has been discussed, defended, criticised, reformulated, and re-criticised for generations; I will content myself with laying out how I understand the argument in the moral case, which will put us in a position to compare it to an analogous argument in the epistemic case.<sup>8</sup>

Reductive naturalists about some normative or evaluative property identify it with some natural property. This view can take one of two forms. On *analytic reductionism*, the natural-normative identity claim is analytic; on *synthetic reductionism*, it is synthetic. The Open Question Argument makes trouble for analytic reductionism.<sup>9</sup> Analytic reductionism will be our focus here since it is, in my view, the crux of the moral-epistemic disparity: an analytic version of epistemic reductionism does not fall prey to the Open Question Argument, whereas all versions of analytic ethical reductionism do fall prey.

The trouble with analytic reductionism in ethics is that whatever natural property is put forth in the reductionist theory as the one that is identical to a certain ethical property, it will be an “open question” whether the things that have the natural property also have the evaluative property. That is to say, of a thing that is known to have the natural property, neither answer (“Yes” or “No”) to the question of whether it also has the evaluative property would suggest that the person giving that answer fails

to grasp the concept of that evaluative property; nor would either answer suggest that the person simply means something peculiar when they use the corresponding evaluative term. Relatedly, neither answer seems incoherent, even if one of them is known to be false.

To illustrate, suppose that, after working on the topic for some time, considering all the many known theories and arguments and counterexamples, you come to the view that hedonism is the correct axiological theory. You believe that

all states of pleasure are good in themselves and only states of pleasure are good in themselves.

Thus, you believe that the properties of being a state of pleasure and being intrinsically good are necessarily coextensive. You then encounter a fellow philosopher who, having scrutinised the same data, disagrees. This philosopher thinks that malicious pleasures are of no intrinsic value. You tell him that of course you agree that a malicious pleasure may have bad effects, and that malicious pleasures are a sign of a vicious character. Is he sure that he isn't thinking of those things when he refrains from evaluating malicious pleasure positively? Yes, he says, he is sure; he has *intrinsic* value in mind. And on the question of the intrinsic value of malicious pleasure, it continues to seem to him that malicious pleasure is of no intrinsic value.

Committed hedonist that you are, you have no choice but to think that your colleague has simply got this one wrong. His view is mistaken. But will you think, further, that his view betrays a conceptual confusion, or a failure to grasp the very question of whether something is intrinsically good, or is an incoherent view? I don't think so. As much as we might enjoy being able to accuse a staunch philosophical opponent of these things, such an accusation is not plausible in this case. Nor would you think that your colleague just must mean something different by “intrinsically good” than you do. Instead, how things seem is not that you and your colleague are talking past each other but that you simply have different beliefs about a single topic.

This finding has metaethical implications for you. You are convinced that *being intrinsically good* and *being a state of pleasure* are two properties

that necessarily go together. But since it does not betray a conceptual confusion to say, of a thing that is known to be a state of pleasure, that it is not intrinsically good, you conclude that you cannot analytically reduce intrinsic goodness to the property of being a state of pleasure, or regard “intrinsic goodness” and “being a state of pleasure” as two names for a single notion. Rather, to say of a thing that it is an episode of pleasure is one thing, and to say of a thing that it is intrinsically good is another. The Open Question Argument thus convinces you that analytical reductionism in metaethics is mistaken.

The point here is not just one about the particular identification of the concepts of being intrinsically good and being a state of pleasure. The same line of reasoning would go through no matter which natural or non-evaluative concepts we thought to be necessarily co-extensive with the concept of intrinsic goodness. Those who affirmed evaluative judgments in conflict with this necessary connection would not thereby seem to us to be conceptually confused or to be talking about some other topic. We would simply take them to be speaking falsely about our shared topic.

Some people object that this sort of argument proves too much. They think that it could be used to refute reductive analyses on any topic, or at least any reductive analysis that is less than obvious. I agree that there are some non-obvious analytical reductions, but I don't think this argument conflicts with that fact.

Before considering a case of a non-obvious reduction, it will be useful to begin with an example of an obvious one. Consider unclehood. After briefly reflecting on what it is to be an uncle, you come to the view that

to be an uncle is to be a male who has a sibling who has an offspring.

You then encounter a fellow inquirer who, having considered the matter as well, disagrees. He insists, let's suppose, that the offspring has to be a male for the person to be an uncle. Before us is Joe. Joe, we all agree, has a sibling with an offspring. But is Joe an uncle? Our interlocutor answers, “No” (because the offspring is a girl). Committed as you are to the view that to be an uncle is to be a male who has a sibling who has an offspring, you have no choice but to think that your fellow inquirer has got this one wrong. His view about Joe is mistaken. Will you think, further, that his

view betrays a conceptual confusion, or a failure to grasp the very question of whether someone is an uncle? *Yes*. Your fellow inquirer simply fails to grasp unclehood. He means something different from the rest of us when he uses the term “uncle” (he also means something different from what the term actually means). Consequently, the question of whether Joe is an uncle is *not* an open question. Thus, the Open Question Argument, when applied to our analysis of unclehood, does not generate a false positive in this case of an obviously correct reduction.

Nor does it in the case of non-obvious reductions. Consider bachelorhood. Suppose that you are thinking about the extension of this property. After considering the known theories and arguments and counterexamples, you come to the conclusion that what we can call “the simple view of bachelorhood” is right:

all unmarried men are bachelors and all bachelors are unmarried men.

You then encounter a fellow inquirer who, having scrutinised the same data, disagrees. She thinks that the pope is not a bachelor. You tell her that of course you agree that the pope is not an *eligible* bachelor, and you note that often when we are interested to know whether some man is a bachelor, we are interested in his eligibility for marriage. Is she sure that she isn't just thinking of eligibility for marriage? Yes, she says, she is sure; she is talking about plain old bachelorhood, and it seems to her that the pope simply does not qualify as a bachelor.

Committed defender of the simple view that you are, you have no choice but to think that your colleague has got this one wrong. Her view is mistaken. Will you think, further, that her view betrays a failure to understand the concept of bachelorhood? Yes; you *should* think this. You should think that she just does not get what it is to be a bachelor. Though it is more understandable in this case than in the unclehood case, since it is less obvious what the correct analysis of bachelorhood is. In the hedonism case, you don't think that your interlocutor fails to get *what it is* to be intrinsically good; he simply fails to see which things have this property. But in the bachelorhood case, if you are convinced of the simple view of bachelorhood, you should think that your interlocutor just fails to get what it is to be a bachelor. She must mean something different

when she uses the term “bachelor,” something different from what you think the term means. Since you think that *in fact* the right property to associate with the term “bachelor” is given by the simple view, you will think that she is linguistically or conceptually confused. Thus, an Open Question Argument designed to show that we should be non-reductionists about bachelorhood would fail. And it would fail even though it is not obvious what the correct analysis of bachelorhood is.

When we encounter someone who holds a belief that conflicts with our hedonistic theory of intrinsic value, we judge, perhaps after some initial questioning, that they are working with the same concept as us—the concept of intrinsic value. We just think that they hold different beliefs about when the concept applies. This is the normal case of disagreement, as when people disagree about, say, whether alligators eat their young. The parties to such disagreements share the same concepts. But when we encounter someone who holds a belief that conflicts with our simple view of bachelorhood, we conclude that they in fact have a different *concept* of bachelorhood and that the meaning that they associate with the word “bachelor” is different from what the word actually means. The Open Question Argument is a good argument against analytic reductionism in metaethics, but it would not be a good argument against any analytic reductionist view about bachelorhood.

This should go without saying, but, as with any philosophical argument, it is always possible for a stubborn advocate of the view being attacked simply to reject the key premise on which the argument against their view is based. Thus, the committed analytic reductive hedonist can respond to the colleague who is convinced that malicious pleasure is not good by saying that the colleague just must not get the concept of intrinsic goodness; they must not get what we are saying about something when we say that it is intrinsically good; they must just mean something different by “intrinsically good” from what the word actually means. If someone is not prepared to reconsider their theory, then this is probably what they should say. Our job, then, is to evaluate the plausibility of saying this. For my part, this imagined reply strikes me as implausible, as something that only someone in the grip of a theory would say. Thus, it seems to me that, all things considered, the Open Question Argument

casts doubt on analytic reductionism in ethics, though an analogous argument does not cast doubt on analytic reductionism about unclehood or bachelorhood.

### 3 The Epistemic Open Question Argument

So it is with ethics and uncles and bachelors. What about epistemology? I believe that when it comes to analytic reductionism about epistemic properties, the Open Question Argument is not compelling. It is like the Open Question Argument against analytic reductionism about bachelorhood or unclehood. Since the Open Question Argument is the main obstacle standing in the way of analytic reductionist doctrines about normative topics, and since it fails in the epistemic case, we should accept analytic epistemic reductionism. And, interestingly, we should do so even if we aren't sure just what the correct reductive analyses of the epistemic notions are.

To explain why I think the epistemic Open Question Argument fails, we can begin with a sample analysis of epistemic reasons or reasonableness in non-normative terms. But I would like to emphasize at the outset that it does not matter whether the sample analysis is perfectly extensionally adequate. For our question is the question of whether epistemic reductionism as a general thesis about the nature of epistemic facts, abstracted away from any particular proposed reduction, is true. My method for answering this abstract question is as follows. First, we identify an analysis that is at least not obviously extensionally hopeless; *prima facie*, it at least has a shot at being extensionally adequate. Next, we *suppose* that this analysis is in fact extensionally adequate (we can do this even if we suspect or even know that it is not). Then we ask, within the scope of this supposition, how plausible it would be to hold that not only does the condition on the right-hand side of the analysis necessarily covary with that on the left-hand side, but they are in fact just one condition under two labels. And the method for answering this is the open-question test (concerning conceptual confusion, idiosyncratic meanings, and incoherence) that we have been working with.

Here is a kind of first-pass theory of the concept of reasonable belief that is at least not obviously extensionally inadequate on its face:

a claim is reasonable for a subject to believe just in case the claim is likely to be true, given the subject's evidence.

To illustrate, suppose Sally looks out her window first thing in the morning and sees that the streets are wet. Relative to Sally's prior background beliefs together with this new piece of information, it is very likely *that it recently rained*. Moreover, intuitively, given the case as described, it is in fact reasonable for Sally to believe that it recently rained. Thus, we have a confirming instance of our account of epistemic reasonability in terms of probable truth.

Suppose we find that the theory in question is in fact extensionally adequate.<sup>10</sup> There are no counterexamples, or cases in which a belief is, intuitively, epistemically reasonable for some subject but not probable relative to their evidence (or vice versa). Our first-order epistemological investigation into the conditions under which a belief is reasonable would be complete (at least in the absence of new data on the matter). But a metaepistemological question of central importance would remain. When we say that a claim is reasonable for a subject to believe, *is all we are saying* that that fact is likely to be true given the subject's evidence? Or, alternatively, are a claim's being reasonable for a person to believe and its being probable relative to their evidence two distinct states of affairs—one irreducibly normative and one natural—that necessarily go together? In other words, is epistemic reductionism true or is epistemic non-reductionism true?

To answer this question, we can apply our open-question reasoning as earlier described. We can suppose that, although we are convinced of and committed to the extensional adequacy of our epistemological theory (the biconditional on display above), we encounter, just as in the earlier cases, someone who holds a belief in conflict with it. Let's suppose that it is Sally, about her own case above. Sally agrees that, given her information, it is likely that it recently rained. But Sally, to our surprise, denies that it is reasonable for her to believe that it rained. Here is what she says: "I see that the streets are wet. And I realise that, given this, and given

what else I know, it is very likely that it rained recently. Nonetheless, I don't think that it would be reasonable for me to believe that it rained recently."

If you are like me, these remarks of Sally's will strike you as very odd indeed. She is saying roughly this: "Yes, that claim is, from my perspective, very probably true, but it is not reasonable for me to believe it." These remarks seem more similar to the words of the person who thinks that, for you to qualify as an uncle, your sibling needs a male child than they are to the remarks of the person who thinks that malicious pleasure does not qualify as intrinsically good. Sally sounds conceptually confused, like she doesn't understand what she is saying.

Note that her claim is worse than Moore-paradoxical. Moore-paradoxical claims, such as of the form, "that claim is true, but I don't believe it," are at least logically consistent, even if seriously infelicitous. But Sally's remark sounds incoherent. At a minimum, we do not have a positive intuition that the claim *is* self-consistent, as we do in the case of the claim, relative to the supposition that hedonism is true, that malicious pleasure is not intrinsically good. To build a case for metaepistemological non-naturalism on the basis of the idea that Sally's view is coherent would be to place quite a bold thesis on quite a frail reed.

It is possible to get Sally's claim to sound positively coherent, but only if we interpret the term "reasonable" in the wrong way, as signifying some other sort of normative evaluation, such as a prudential one. Pascal's position was something like this: "given all our evidence, theism is just as likely to be false as it is to be true, but, still, it is reasonable for us to believe that it is true." A starker variation on Pascal is this: "given all our evidence, theism is very unlikely to be true but, still, it is reasonable for us to believe that it is true." These Pascalian claims are indeed coherent. But that is only because we are interpreting "reasonable" to mean *prudentially* reasonable. That such claims are coherent is of no help in supporting the sort of claim that defenders of the epistemic Open Question Argument need. They need claims like the following to be coherent: "given all our evidence, theism is very unlikely to be true but, still, it is *epistemically* reasonable for us to believe that it is true."

I conclude that, unlike the moral Open Question Argument, the epistemic Open Question Argument is not convincing. Consequently, in the

epistemic realm, the door is open to embrace analytic reductionism and to enjoy its attendant theoretical advantages. The only disadvantage is to those of us who were hoping to be able to use the premise of moral-epistemic parity to support robust moral realism. Unfortunately for us robust moral realists, epistemic normativity appears not, after all, to be a partner-in-guilt of irreducible moral normativity.

## 4 What If the Proposed Analysis of Epistemic Reasonability Is Not Extensionally Adequate?

### 4.1 Richard Rowland's Counterexample

In a recent paper, Richard Rowland argues against the moral error theory on the grounds that it implies that there are no epistemic reasons (Rowland 2013). In doing so, he considers the position that I favour, “that a conceptual analysis of epistemic reasons in nonnormative terms is more plausible than a conceptual analysis of moral reasons in nonnormative terms” (Rowland 2013, 7). But he rejects this moral-epistemic disparity on the grounds that the particular analysis that he considers—one that reduces facts about epistemic reasons to facts about probability—is not extensionally adequate.

I don't believe that Rowland's arguments should make us worry that epistemic facts are not reducible after all. First, as a general matter, since epistemic reductionism is a doctrine in metaepistemology, we don't have to worry whether our working analysis of epistemic facts in terms of facts about likelihoods in perfectly extensionally adequate. Our central question is this: supposing (perhaps contrary to fact) that our theory is extensionally correct—in other words, that the corresponding biconditional is necessarily true—should we conclude (a) that, as in the unclehood and bachelorhood cases, the theory is analytic and what the right-hand side says is just another way of saying what the left-hand side says, or should we conclude (b) that, as in the moral case, the extensionally correct biconditional is a synthetic necessary truth rather than an analytic one?

In the previous section, I argued for (a). That argument suggests that epistemic reductionism is correct even if our supposition (that the particular biconditional under consideration is true) is mistaken. Showing via first-order epistemological considerations, as Rowland tries to do, that the theory is in fact not extensionally correct does not undermine this metaepistemological position.

These abstract points can be illustrated by examining Rowland's particular counterexample. Although Rowland's specific target is my (2009) paper, the analysis he criticises is a little different from the main one discussed in that paper.<sup>11</sup> According to the analysis Rowland focuses on,

the epistemic predicate “is an epistemic reason to believe  $p$ ” means the same as the naturalistic predicate “is a fact that raises the probability that  $p$ .” (Rowland 2013, 8)

Rowland thinks that this equivalence fails, but not for open-question-type reasons. It fails, according to Rowland, for first-order epistemological reasons: it isn't even extensionally correct.

Consider the fact that I am in my office today. This raises the probability that I will die in my office today. That is to say, if I were at no point in my office today, the probability that I die in my office today would have been less than it actually is. “But,” Rowland claims, “it does not seem that the fact that I am in my office today is a reason to believe that I will die in my office today; I am perfectly healthy, and probably safer in my office than anywhere else” (2013, 8).

Rowland's counterexample is interesting, but I don't think that it is successful in the end. Moreover, even if it were, it wouldn't matter for our central concern, which is not the precise extensional adequacy of any particular analysis but the prospects of going reductive once an extensionally adequate theory has been settled on.

About Rowland's particular counterexample, my considered judgment is that the fact that I am in my office today *is* a reason to believe that I will die in my office today. Intuitively, it *is* a consideration in favour of that belief. After all, and as Rowland agrees, that I am in my office today *makes it more likely* that I will die there today. How could it not then *support the belief* that I will die in my office today? Notice how weak the claim in

question is. Saying that the fact that I am in my office today is *a* reason to believe that I will die in my office today is just to say that it adds *some* support for the belief, or that it counts in its favour *at least somewhat*. It does not imply that one *should* believe that I will die in my office today. Of course one shouldn't believe that (given the other details of Rowland's case: that I am perfectly healthy and that my office is very safe). But there can be reasons in favour of beliefs even when those beliefs are, all things considered, unreasonable to have.

That the fact in question does provide a reason in this case can be seen by comparing two irrational believers. The first knows that I am in my office and believes, irrationally and despite overwhelming evidence to the contrary, that I will die there today. In the second scenario, at no point today am I in my office. Our second believer knows this and yet believes, also irrationally and despite even more overwhelming evidence to the contrary, that I will die there today. Both believers are irrational, but the second one is more irrational. Why? The best explanation of this will involve the claim that the first believer's belief had at least one thing going for it that the second believer's belief did not: namely, that I was, in that first case, at least in my office, a fact that at least made my dying there more likely. The first believer is less irrational because he had at least this reason for his belief. Thus, that I am in my office today is indeed a reason to believe that I will die in my office today.

But even if you agree with Rowland's counterexample, that doesn't affect our larger question, the question of epistemic reductionism in general, abstracted away from the extensional adequacy of any particular reduction. To answer the question that concerns us, we can first *suppose* that in fact Rowland is wrong and that the analysis that he discusses is extensionally adequate. Then, within the scope of this supposition, we consider whether it is more plausible to go reductive or to go non-reductive. When we do, we find that there are no open-question-type pressures to go non-reductive. Supposing that all and only cases of a fact giving a reason to believe are cases of probability raising, the thought "I see that this fact raises the probability that *p*, but I don't think it gives me any reason to believe that *p*" isn't just false but is plausibly regarded as incoherent.

Another way to see why Rowland's putative counterexample should not worry the epistemic reductionist, even if it is a successful counterexample,

is to notice that once, via first-order epistemological theorising, we repair our analysis to avoid the counterexample, we can apply the open-question test again, and again it will support reductionism over non-reductionism. This is not the place to get lost in the salt mines of first-order epistemology, but here is a quick illustration. Perhaps we learn from Rowland's example that our analysis should be formulated as follows:

for a fact to be a reason to *increase one's credence* in *p* just is for that fact to raise the probability that *p*, given one's evidence.

This theory avoids Rowland's counterexample. The fact that I am in my office today clearly is a reason to *increase one's credence* in the proposition that I will die in my office today (presumably from a very low credence to a slightly higher but still very low credence), even if it isn't a reason to *positively believe* that I will die in my office today. Now imagine someone saying, "I see that the fact that you are in your office today increases the probability that you will die in your office today, but, still, I don't think it gives me any reason to increase my credence that you will die in your office today." That has the same air of incoherence about it that our earlier examples did. This is further confirmation for epistemic reductionism. I conclude that Rowland has not given us good reason to doubt this doctrine.

## 4.2 Cuneo and Kyriacou's Counterexample

Terence Cuneo and Christos Kyriacou's recent paper, "Defending the Moral/Epistemic Parity," is a sustained attack of my (2009) attempt to establish moral-epistemic disparity (Cuneo and Kyriacou 2018). They advance a number of interesting criticisms, one of which begins as a kind of counterexample to my working analysis of epistemic reasonability, the theory that for a belief to be epistemically reasonable for some subject is for it to be probable given the subject's evidence.<sup>12</sup>

In Cuneo and Kyriacou's counterexample, you walk into a room and have a visual experience as of a table in front of you. On this basis, you form the belief that there is a table in front of you. Just before you step into the room, however, there is a small placard that enters your visual



field but that you do not read or even notice. The placard indicates that you are about to encounter a masterful *trumpe l'oeil* mural containing an image of a table. Cuneo and Kyriacou write,

Given all your available evidence, it follows that the proposition *that there is a table in front of me* is not likely and, thus, according to Heathwood's proposal, it is not reasonable for you to believe that there is a table before you. But, by all appearances, it is reasonable for you to believe that there is a table before you. By no fault of your own or your eyes, you simply missed taking into account some available but not easily detectable information. (Cuneo and Kyriacou 2018)

I think that Cuneo and Kyriacou are too quick to assume that my proposal implies that it is not reasonable for you to believe that there is a table before you. For it is not clear that the information on the placard is a part of your evidence. Indeed, intuitively, I would say, what is written on the placard is not a part of your evidence. And note that some prominent accounts of evidence possession agree. According to Timothy Williamson's view, for example, a person's evidence consists of everything they know (Williamson 2000, 184–85). Since, in the case above, you don't know what the placard says, this is not a part of your evidence on Williamson's account. And according to Richard Feldman's account, a person's evidence at a time consists of everything that they are thinking about or are aware of at that time (R. Feldman 1988). The information on the placard is not a part of your evidence on this account either. My theory leaves it unsettled what it is for some proposition to be a part of a subject's evidence. This is by design, to avoid getting mired in first-order epistemological controversies. But I suspect that the correct view of evidence possession will agree with the views of Williamson and Feldman that, in the imagined case, the information on the sign is not a part of your evidence.<sup>13</sup>

In concluding that, on my proposal, it is not reasonable for you to believe that there is a table before you, Cuneo and Kyriacou cite your "available" evidence. I concede that there is a reading of "available evidence" on which the information on the placard is *available* to you. And in my previous paper I slid incautiously between the notions of evidence *in a subject's possession* and evidence *available to a subject*. I consider the

present reply to Cuneo and Kyriacou, then, as an opportunity to lay out my position more clearly and carefully. Though I hasten to add that my main argument should still go through despite my incautiousness. That argument asks us only to *suppose* that the theory in question is extensionally adequate, or immune from counterexample, and then to consider whether, given that, one should prefer a reductive or a non-reductive version of the theory.

## 5 Cuneo and Kyriacou on "Relevant Evidence"

A theme throughout Cuneo and Kyriacou's paper is that, for the purposes of my analysis, which makes use of the notion of a person's evidence, it will not do to appeal to the person's *total* evidence. That would be too indiscriminate, as the *trumpe l'oeil* case is meant to bring out. In Cuneo and Kyriacou's view, the best alternative is to appeal the subject's *relevant* evidence, where this amounts to the evidence that the subject *ought* to take into account (Cuneo and Kyriacou 2018). They worry that while this, an account stated in terms of relevant evidence, might make for an extensionally correct analysis, it will no longer be a descriptive analysis of the epistemic notion—and thus my case for moral-epistemic disparity would be undermined.

But now that we have clarified that the analysis is appealing to evidence that the subject *possesses*, and not evidence that is in some wider sense *available* to the subject, it is plausible for the view to appeal to the total evidence that the subject possesses. One worry of Cuneo and Kyriacou that goes away when we restrict the theory to possessed evidence is the worry that it is not possible for a subject to base a belief on their total evidence. For perhaps Feldman is correct and a subject's evidence—their total evidence—consists merely of everything that they are currently thinking. If so, Cuneo and Kyriacou's worry disappears, since there is presumably no difficulty in basing a belief on everything that one is currently thinking. Or perhaps Williamson's more capacious view is correct and a subject's evidence is everything that they know. Again, there seems to be no great worry here; a person can base a belief on everything

they know. And if there is a worry, that may simply be a reason to favour a more restrictive view of evidence possession, such as Feldman's, over a more capacious one. We, or, better, first-order epistemologists, can in fact construct their theory of what it is to possess evidence with the aim of making the theory of epistemic reasonability more plausible.<sup>14</sup> It is not our place to settle such matters. But once they are settled by first-order epistemological investigation, there is no reason to think that the epistemic Open Question Argument will be any more effective in undermining epistemic reductionism.

Let me make one final point, concerning the idea that the notion of relevant evidence should be used and understood normatively, as Cuneo and Kyriacou suggest. I have just explained why we don't need to take the route of appealing to relevant evidence. But even if we were to take this route and were convinced that the notion of relevant evidence should be understood normatively—as meaning *evidence that the subject ought to take into account*—this would not thereby undermine the case for epistemic reductionism. Because the normative supervenes on the descriptive, we know in advance that there is some (possibly complex and disjunctive) descriptive property on which the property of *being evidence that some subject ought to take into account* supervenes. If we were to take Cuneo and Kyriacou's suggested route, then, given the general failure of the Open Question Argument for epistemic notions, once we have identified this descriptive property, we would be entitled to reduce the normative notion of *evidence that a subject ought to take into account* to it.<sup>15</sup>

## 6 Cuneo and Kyriacou's Argument for Moral and Epistemic Parity as Regards the Prospects for Reduction

In addition to airing the above suspicions about the kind of reductive analysis of epistemic normativity that I favour, Cuneo and Kyriacou present a positive argument for the view that morality is on a par with epistemic rationality. They call attention to a challenging kind of dynamic for formulating different sorts of reductive analyses of normative notions. On the one hand, if the *analysans* is conceptually rather similar to the

*analysandum*, then the analysis is more likely to seem plausible but less likely to seem genuinely reductive (that is, the *analysans* will include, or at least might seem to include, normative concepts, just like *analysandum*). If, on the other hand, the *analysans* is conceptually very dissimilar from the *analysandum*, then the analysis is more likely to seem genuinely reductive but less likely to seem true. A successful reductive analysis must steer a middle ground between these two pitfalls.

Cuneo and Kyriacou think that, when looked at in light of this challenge, the moral and epistemic cases are closely analogous. We can formulate a spectrum of analyses in each case, and the prospects of success for the various positions on the spectra will seem similar in each case. This supports, Cuneo and Kyriacou believe, moral-epistemic parity as regards the prospects for reduction.

More specifically, Cuneo and Kyriacou think that my analysis of epistemic reasonableness and an imagined analysis of a moral notion (moral goodness) occupy similar positions on their respective spectra such that, in each case, it is both hard to say whether the analysis is true and, even if it is true, unclear whether the *analysans* contains a normative notion. Their imagined analysis of moral goodness is this:

a thing is morally good just in case it would be prized by an ideal agent in idealised conditions. (Cuneo and Kyriacou 2018)

Because, they suggest, this theory is perhaps somewhat plausible yet not clearly true and also not so clearly genuinely reductive, they think the moral case is relevantly similar to the epistemic case. For in the epistemic case they think that my analysis of epistemic reasonableness in terms of probability is perhaps somewhat plausible yet not clearly true and also not so clearly genuinely reductive. They conclude that “we have as much reason to believe that we can furnish a reductive analysis of epistemic concepts as we have to believe that we can furnish a reductive analysis of moral concepts”. (Cuneo and Kyriacou 2018)

I continue to think, however, that crucial disparities remain. One difference is that Cuneo and Kyriacou's imagined analysis of moral goodness, as an ideal-observer analysis, admits of two familiar interpretations, one of which is clearly not reductive and the other of which clearly is. The

covertly normative, non-reductive version is one in which the definition of “ideal” or “idealised” contains normative terms; the truly reductive version contains only psychological and other naturalistic notions in these definitions. That the *analysans* in the moral case admits of these two familiar interpretations, one clearly normative and the other clearly not, is one difference between the moral and epistemic cases. The other, more crucial difference becomes manifest once this first difference has been brought to light: if the analysis of moral goodness is interpreted as a genuinely reductive one, it will fall prey to the Open Question Argument.<sup>16</sup> This is in contrast to our epistemic analysis, which is of course intended to be understood as a genuinely reductive one. Even when this analysis is understood as intended, as genuinely reductive, an Open Question Argument against it fails. Thus, the crucial moral-epistemic disparity remains.

A truly reductive version of Cuneo and Kyriacou’s imagined theory might hold, as Roderick Firth’s well-known reductive version of the theory held, that the ideal agent in idealised conditions has these features: this person is omniscient with respect to non-ethical facts, omniperceptive, disinterested, dispassionate, consistent, and in other respects normal (Firth 1952, 333–45). It is a very difficult empirical question what things such a being would prize, but let’s suppose that it turns out that such a being would prize (intrinsically, or for its own sake) even malicious pleasure. Even so, the following remains an open question: I see that a being who is omniscient with respect to non-ethical facts, omniperceptive, disinterested, dispassionate, consistent, and in other respects normal would prize malicious pleasure for its own sake, but is malicious pleasure really good in itself? A person who held that the answer to this question is “No” would not thereby indicate that she meant something different by “good in itself,” would not thereby manifest conceptual confusion, and would not appear to be holding an incoherent position. Thus, a genuinely reductive version of Cuneo and Kyriacou’s imagined theory falls prey to the Open Question Argument.

The other main version of the view, the version that is covertly normative, arguably does not fall prey to the Open Question Argument. One simple (if simplistic) way for the view to have an ultimately normative *analysans* would be for the characterisation of the ideal observer to include

omniscience about *ethical* facts together with a disposition to prize things in accord with one’s ethical beliefs. The following is not an open question: I see that a being who is omniscient with respect to ethical facts and who prizes only what they judge to be good would prize malicious pleasure, but is malicious pleasure really good? A “No” answer to this question is plainly incoherent.

But my proposed version of epistemic reductionism is not covertly normative. Or at least that is and has been my position, and the argument of Cuneo and Kyriacou described in this section has given us no reason to think that my version of epistemic reductionism has a covertly normative *analysans*.<sup>17</sup> But Cuneo and Kyriacou have attempted to give such reasons elsewhere in their paper. To these I now turn.

## 7 Is Probability Normative?

Cuneo and Kyriacou argue that if my analysis of epistemic reasonability is to be plausible, the notion of probability at work in it needs to be covertly normative. They present a dilemma:

Suppose, on the one hand, that we attempt to analyze the concept being reasonable by appeal to an *objective intrinsic* account of probability, as Heathwood suggests. If we do, then we will not be able to fashion an adequate analysis of the concept being reasonable. Suppose, on the other hand, we attempt to analyze the concept being reasonable by appeal to a *conditional epistemic* account of probability. If we do, then we might arrive at an adequate analysis of this concept. That analysis, however, would provide no reason to believe that probability facts are descriptive; to the contrary, we’ll suggest, it would provide reason to believe that such facts are normative, in which case Heathwood’s case for the disparity thesis would collapse. (Cuneo and Kyriacou 2018)

I agree that neither of these accounts of probability is suitable for the role needed by my reductive analysis of epistemic reasonableness. Objective intrinsic probability, as Cuneo and Kyriacou understand this expression, won’t do because a subject may have no evidence pertaining to an event’s objective intrinsic probability. To take a variant of their

example, suppose you have a biased coin in your hand—both sides are “tails”—but I have every reason to think that it is a normal, fair coin. The objective intrinsic probability of its landing on “tails” on each of the next ten tosses is 1 (or very nearly 1), but it would not be reasonable for me to believe that it will land on “tails” on each of the next ten tosses.

Conditional epistemic probability, as Cuneo and Kyriacou understand it, won't do either. For as they understand it, it is a normatively loaded notion, amounting either to a normatively loaded kind of ideal epistemic agent account or to an account that appeals, again, to relevant evidence, which they believe must be understood normatively.

I reject both of these understandings of the relevant sort of probability. The notion that is playing a role in my analysis, the notion of *a claim being likely, given some subject's evidence* is at once *subject-relative* and *objective*. It is subject-relative in that the likelihood is relative to the subject's evidence. And it is objective in that whether that body of evidence makes the claim likely is an objective fact, made true by objective relations between that evidence and the claim in question, and not, say, by the agent's wishes or whims. This kind of objective subject-relativity is familiar. It appears, for example, in the unremarkable notion of *a claim being entailed by some subject's beliefs*. Such a notion is obviously subject-relative: what is entailed is entailed by this particular subject's beliefs. It is also objective: whether some proposition is entailed by some specified set of propositions is an objective fact about the world.

## 8 Concluding Remarks

I believe that, though the Open Question Argument shows that reductionism in ethics is not plausible, it does not show the same about reductionism in epistemology. I have here tried to lay out more fully and clearly why I believe this, and to defend this position against a number of recent objections. Since the Open Question Argument is the main obstacle to reductive views about normative phenomena, the door is open to be epistemic reductionists. This, unfortunately, is bad news for robust moral realists like me who might have hoped to find support from the epistemic realm, in the form of a partner-in-guilt, for robust moral realism.

## Notes

1. For brief explanations of these advantages, see Heathwood (2013).
2. Another interesting defense of moral-epistemic parity can be found in Greco (2015). I lack the space to discuss Greco's arguments here, but see Côte-Bouchard (2017).
3. Furthermore, they are not themselves natural properties. This addition is required to distinguish non-naturalism from non-reductive naturalism.
4. More exactly, they should hold that the property of *basic* intrinsic goodness is necessarily coextensive with the property of being a state of pleasure (see Harman 1967 and Feldman 2000). I ignore this complication in what follows.
5. A terminological note: in what follows, since it won't matter, I won't fuss over the differences between natural and descriptive properties, and I will also sometimes use the expression “non-moral property” for the same thing. I also won't fuss over the difference between reductionism and naturalism.
6. This phrase is due to Milo (1995, 182, 190–93). See Shafer-Landau (2003, 15–17) for the same sort of definition of moral realism.
7. Provided that pleasure is stance-independent, which is in fact a controversial view. See Bramble (2013) for a stance-independent account of pleasure and Heathwood (2007) for a stance-dependent theory.
8. The Open Question Argument is due to Moore (1903). For a helpful interpretation, see Feldman (2005).
9. Incidentally, I believe that synthetic reductionism fails for other reasons: moral terms don't function the way that synthetic reductionism requires them to function. See Barnett (2002) and Gampel (1996).
10. To emphasise again: it is not my claim that this biconditional in fact has no counterexamples. Nor am I claiming that it is in no need of clarification or refinement. For example, does “likely to be true” here mean merely a likelihood of greater than 0.5, or does it mean something greater than that, or perhaps something both greater than that and vague? For another example, might we want to include something to the effect that the belief must be *based on* the evidence? Or perhaps we are talking here about propositional rather than doxastic justification, so that appeals to the basing relation are not needed? Also, what is it for something to be a part of a person's evidence (this will we discuss briefly below (Sect. 4.2)). These are all important issues, but since my project is

the metaepistemological one, I don't want to—and don't much need to—get bogged down in the controversies of first-order epistemology.

11. The main one discussed in my (2009) paper is the same analysis that is on display in Sect. 3 above. In that paper, I do allude in passing to the closely related analysis that Rowland focuses on (Heathwood 2009, 89–90).
12. I don't respond to every one of Cuneo and Kyriacou's concerns in the main text. One of these concerns is that even if I am right that *epistemic reasonability* is reducible, this does not show that all epistemic notions are reducible. This is a fair point. My reductive analysis of reasonable belief is only a first step towards showing that the entire epistemic realm can be explained naturalistically. But epistemic reasonability (in other words, epistemic justification) is, it is fair to say, *the* central normative notion in epistemology. Cuneo and Kyriacou also object that my position commits non-naturalists in metaethics to the view “that there are not one but two types of reason properties: one that is wholly descriptive and one that is not” (Cuneo and Kyriacou 2018). This is true. But I just think that it is a straightforward consequence of the fact that epistemic facts are naturalistically reducible while moral facts are not, a position that, I have been arguing, we have good reason to accept.
13. And if it really is some epistemologist's view that, in the example, you have, as a part of your evidence, the information that you are looking at a *trumpet l'oeil* mural, then I take it that the epistemologist should say that it is in fact not reasonable for you to believe that there is a table before you. After all, you had decisive evidence to the contrary.
14. Cf. Thomas Kelly's remarks about “theoriz[ing] in the opposite direction”: “to the extent that one has independent intuitions about what an individual would be justified in believing in a given scenario, such intuitions will shape one's views about what evidence must be available to an individual so situated—and therefore, one's views about the more general theoretical issue about what evidence is, or what sorts of things can and cannot qualify as evidence” (Kelly 2014, sec. 1).
15. This argument assumes that the “ought” in *evidence that a subject ought to take into account* is an epistemic “ought.” If this were some other kind of “ought,” such as a moral “ought,” the strategy described above would not succeed (at least given, what I believe, that moral notions are not reducible to descriptive notions). In general, if the moral and the epistemic are in this way entangled, in that whether a belief is reasonable depends on moral factors, my argument for moral-epistemic disparity

would fail. For a recent argument for this kind of moral-epistemic entanglement, see Case (forthcoming).

16. The idea that an ideal observer analysis will fall prey to the Open Question Argument if and only if it is understood reductively is recognised by advocates of the Open Question Argument in ethics, such as, for example, Michael Huemer: “I think that there are two ways of understanding this sort of theory [the ideal observer theory], one that makes it non-reductionist and immune to the Open Question Argument, and another that makes it reductionist but vulnerable to the Open Question Argument” (Huemer 2005, 68).
17. I should emphasise that Cuneo and Kyriacou don't mean the argument that I discuss in this section to stand on its own. It is part of a larger line of argument, which includes the part that I address in the next section.

## References

- Barnett, David. 2002. Against a Posteriori Moral Naturalism. *Philosophical Studies* 107 (3): 239–257.
- Bramble, Ben. 2013. The Distinctive Feeling Theory of Pleasure. *Philosophical Studies* 162 (2): 201–217.
- Case, Spencer. forthcoming. From Epistemic to Moral Realism. *Journal of Moral Philosophy*.
- Côte-Bouchard, Charles. 2017. Epistemological Closed Questions: A Reply to Greco. *Manuscrito* 40 (4): 97–111.
- Cuneo, Terence, and Christos Kyriacou. 2018. Defending the Moral/Epistemic Parity. In *Metaepistemology*, ed. Conor McHugh, Jonathan Way, and Daniel Whiting. Oxford: Oxford University Press.
- Feldman, Fred. 2000. Basic Intrinsic Value. *Philosophical Studies* 99 (3): 319–346.
- . 2005. The Open Question Argument: What It Isn't; and What It Is. *Philosophical Issues* 15 (1): 22–43.
- Feldman, Richard. 1988. Having Evidence. In *Philosophical Analysis*, ed. D.F. Austin, 83–104. Dordrecht: Kluwer Academic Publishers.
- Firth, Roderick. 1952. Ethical Absolutism and the Ideal Observer. *Philosophy and Phenomenological Research* 12 (3): 317–145.
- Gampel, Eric H. 1996. A Defense of the Autonomy of Ethics: Why Value is Not like Water. *Canadian Journal of Philosophy* 26 (2): 191–209.

- Greco, Daniel. 2015. Epistemological Open Questions. *Australasian Journal of Philosophy* 93 (3): 509–523.
- Harman, Gilbert. 1967. Toward a Theory of Intrinsic Value. *Journal of Philosophy* 64 (23): 792–804.
- Heathwood, Chris. 2007. The Reduction of Sensory Pleasure to Desire. *Philosophical Studies* 133 (1): 23–44.
- . 2009. Moral and Epistemic Open-Question Arguments. *Philosophical Books* 50 (2): 83–98.
- . 2013. Reductionism in Ethics. In *International Encyclopedia of Ethics*. Oxford: Blackwell.
- Huemer, Michael. 2005. *Ethical Intuitionism*. 2005 ed. Basingstoke: Palgrave Macmillan.
- Hume, David. (1738) 1975. *A Treatise of Human Nature*. Edited by L.A. Selby-Bigge and P.H. Nidditch. 2nd ed. Oxford: Clarendon.
- Kelly, Thomas. 2014. Evidence. Edited by Edward N. Zalta. *The Stanford Encyclopedia of Philosophy (Fall 2014 Edition)*. <http://plato.stanford.edu/archives/fall2014/entries/evidence/>.
- Milo, Ronald. 1995. Contractarian Constructivism. *The Journal of Philosophy* 92 (4): 181–204.
- Moore, G.E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.
- Rowland, Richard. 2013. Moral Error Theory and the Argument from Epistemic Reasons. *Journal of Ethics and Social Philosophy* 7 (1): 1–24.
- Shafer-Landau, Russ. 2003. *Moral Realism: A Defence*. Oxford: Clarendon.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.