

## Knowledge, Reasons, and Errors about Error Theory

Charles Côté-Bouchard, Rutgers University  
Clayton Littlejohn, Kings College London

Forthcoming in *Metaepistemology: Realism and Anti-Realism*. Eds. Christos Kyriacou & Robin McKenna. Palgrave Macmillan.

*Penultimate version – please do not cite without permission*

**Abstract:** According to moral error theorists, moral claims necessarily represent categorically or robustly normative facts. But since there are no such facts, moral thought and discourse are systematically mistaken. One widely discussed objection to the moral error theory is that it cannot be true because it leads to an epistemic error theory. We argue that this objection is mistaken. Objectors may be right that the epistemic error theory is untenable. We also agree with epistemic realists that our epistemological claims are not systematically in error. However, this is not because there are robustly normative facts, but rather because the truth of our epistemic claims doesn't turn on whether there are such facts. Epistemic facts, we argue, are not robustly or categorically normative. Moral error theorists should therefore respond to the objection that their view does not commit them to the epistemic error theory.

**Keywords:** Epistemic Error Theory, Epistemic Realism, Moral Error Theory, Moral Realism, Robust Normativity, Epistemic Reasons

### 1. Introduction

Readers are probably familiar with this objection to the moral error theory. According to the moral error theorist, our moral claims are systematically in error. If our moral claims are systematically in error, we should expect the same would be true of our epistemic claims. The moral error theory, in other words, entails the epistemic error theory. This, however, is a step too far. The epistemic error theory cannot be true. Thus, the moral error theory cannot be true, either.

This line of objection treats the epistemic error theory as the absurdum of a reductio.<sup>1</sup> Although we don't accept the epistemic error theory, we think that this objection to the moral error theory rests on a mistake. The objection paints the moral error theorist as someone who denies the existence of a certain kind of reason and who is therefore committed to denying the existence of epistemic facts. We think that denying the existence of the relevant kind of reasons does not commit us to the epistemic error theory. Epistemic realism doesn't require the relevant kind of reason.

---

<sup>1</sup> For defences of companion in guilt arguments, see Case (forthcoming), Cuneo (2007, 2016), Nobis (2005), and Rowland (2013). For critical discussion, see Cowie (2014) and Olson (2011, 2014).

In §2, we shall present the companions in guilt objection to the moral error theory. In §3, we shall flesh out the kind of normativity and reasons that concern error theorists and realists. In §4, we shall argue that we should not think of epistemic norms as robustly normative. In §5, we shall explain why this does not undermine the error theorists' claim that morality is robustly normative.

## 2. Error theories and companions in guilt

Error theories about some domain of thought and discourse D standardly take the following form:

Conceptual Claim: It is conceptually necessary that D-claims represent facts with some feature F (the truth of D-claims turns on whether there are F-facts; they would be systematically in error if there are no F-facts, but not otherwise).

Ontological Claim: There are no F-facts; nothing has feature F.

Conclusion: D-claims are systematically in error.

In metaethics, error theorists argue that *moral* claims are systematically mistaken because they necessarily represent facts that would have to have a feature that nothing in our universe has (Mackie (1977), Garner (1990), Joyce (2001), (2006), Olson (2014), Streumer (2017)). According to most error theorists following Mackie (1977), that feature has to do with the kind of *normativity, authority, prescriptivity, or to-be-doneness* that moral facts would have to have (see especially Garner (1990), Joyce (2001), and Olson (2011), (2014)). As Garner explains:

It is hard to deny that the recognition of a moral fact (say the wrongness of an action), if such a thing were possible, would give us a justifying reason for refraining from performing that action. [...] If we recognize a moral fact, then we recognize it as a moral fact, which is to say that we recognize it as directing us to be or act one way rather than another--and we recognize this whether or not we are moved to obey the directions. It is the peculiar combination of objectivity and prescriptivity [...] that makes moral facts and properties queer [...] (Garner 1990: 143).

The problematic commitment of our moral discourse and thought, then, is a commitment to the existence of a certain kind of normativity and normative reason. It's difficult to state precisely what the problematic reasons are like and we shall try to get clear on this in the next section.

In response to this kind of moral error theory, Cuneo offers a “companions in guilt” argument based on the apparent parity between ethics and *epistemology*. Epistemic facts, he argues, are no less normative and reason-implicating than moral facts. The error theorists’ conceptual claim about moral thought and discourse is also true of epistemic thought and discourse: they essentially represent facts with the same relevant kind of normativity. But if, as error theorists maintain, there are no facts with that kind of normativity, then there are no epistemic facts either. The moral error theory therefore entails the epistemic error theory. This is what Cuneo calls the Parity Principle.

The problem, the objection goes, is that the epistemic error theory cannot be true. This is because it faces a fatal dilemma:

[E]pistemic error theory is either (i) self-defeating insofar as it presupposes the very sort of entities it says do not exist or (ii) implies wholesale epistemological scepticism, according to which we have no epistemic reason to believe anything, including the error theory itself. An implication of this second horn is that epistemic error theory would be polemically toothless if true, as no one could make a rational mistake in rejecting the position. Since neither of these options are palatable [...] we have powerful reason to reject the epistemic error theory. (Cuneo 2016: 78).

Therefore, since it entails the epistemic error theory, we should reject the moral error theory. This objection can be summarised as follows:

*The Companions in Guilt Argument*

1. If the moral error theory is true, then the epistemic error theory is true (Parity Principle)
2. The epistemic error theory is false.
3. Therefore, the moral error theory is false.

Error theorists have two main options in response to this objection. The first is to concede that moral and epistemic error theories are “companions”, but deny their “guilt.” It is, in other words, to attack the second premise and defend the epistemic error theory against Cuneo’s dilemma. The second is to deny companionship and reject the Parity Principle. Whether or not Cuneo is correct about the “guilt” of the epistemic error theory, such an error theory is not

entailed by the moral error theory in the first place. In what follows, we offer a version of the second strategy.

### 3. Normativity thick and thin

What error theorists reject is a particular kind of problematic normativity, one that involves a certain kind of categorical reason. Getting clear on the distinction between the unproblematic and the contentious is the business of this section.

Cuneo represents the disagreement between the moral error theorists and the moral realists as a disagreement about “robust” or “categorical” reasons. Some might treat these notions as interchangeable, but since there are different ways of drawing the distinction between the categorical and non-categorical, we think it will pay dividends to proceed carefully.

To start, let’s look at the commitments of moral and epistemic realism, as Cuneo conceives of them. He thinks that the moral realist should accept some platitudes about morality. Some have to do with the content of morality (for example, that it cannot be wrong to watch Foot’s hedgehogs in the moonlight). The platitude that interests us is what Cuneo calls the *Authority* platitude:

[M]oral [...] facts are or imply excellent reasons for us to behave in certain ways, regardless of our desires, goals, or the social roles we may occupy. In terminology that has become popular, these facts are or imply ‘robust reasons.’ (Cuneo 2016: 74).

Cuneo thinks that this feature of moral facts is what the moral realist and error theorist disagree about. He sees them as agreeing that the existence of moral facts turns on whether there are features of the world that fit these descriptions and then disagreeing about whether there are such things.

As Cuneo sees things, the error theorist’s problem with moral realism is that the moral realist is committed to the idea that there could be features of the world that are authoritative in the way suggested above. To show that there can be such features, he turns to the epistemic domain. He thinks that the epistemic realist has similar commitments. The epistemic realist, he thinks, accepts an epistemic authority platitude as with our conception of moral facts; it’s partly constitutive of our conception of epistemic facts and claims that they have robust normative authority. If, as the epistemic realist believes, there truly are such features of the world, the moral error theorist’s view must be mistaken.

We think that the connection between epistemic realism and the thesis that epistemic facts are robustly normative isn't as clear as Cuneo believes. Just to put our cards on the table, we think our epistemic speech can correctly pick out epistemic facts even if there are no robustly normative facts and reasons. The authoritativeness that interests Cuneo and the error theorists is connected to the existence of certain kinds of categorical reasons. They account for the inescapability that the moral realist associates with moral facts and the reasons that they provide, so in assessing the Parity Principle we should remember that the issue that concerns us is whether we should believe that the epistemic realist is committed to the existence of such reasons. The commitment isn't trivial, recall, because the issue is whether the existence of such reasons is entailed by epistemic realism.

One way of distinguishing the categorical from the non-categorical reasons is in terms of a connection to an individual's contingent desires or ends. The categorical reasons apply to everyone irrespective of their particular contingent ends or interests. If it's true that Agnes has a reason to work on her dribbling or practice her scales, it's probably only true because she wants to become a better midfielder or improve as a cellist. Moral reasons aren't supposed to be like this. If she loses her love of football, it might no longer be true that she has this reason. If she loses her love for her fellow persons it's still supposed to be true, the moral realist says, that she has good reason to promote their interests, keep her promises, and so on.

This isn't a terribly interesting understanding of categoricity. The reasons of etiquette and the law are like this, but they aren't supposed to have the kind of authority that the moral realist thinks moral reasons have. So, we need a more robust notion of robustness and categoricity to capture what's at issue. Brink (1992) distinguishes this first superficial kind of categoricity from more interesting notions. Let's focus on his understanding of authority. We might think that what makes moral reasons interesting is not that they apply to everyone irrespective of their contingent ends or desires but that, in some sense, it is irrational for an agent not to be moved by them even if heeding them wouldn't suit the agent's contingent ends or desires. This is how Brink understands the authoritativeness of moral reasons.

In this respect, it might seem that the realists' epistemic and moral reasons would be similar. The epistemic realists agree, presumably, that a thinker's beliefs might be irrational for reasons that have nothing to do with the thinker's contingent ends or inclinations. The moral realists agree, presumably, that a thinker's actions might be irrational for reasons that have nothing to do with the thinker's contingent ends or inclinations.

The similarities here might be purely superficial. Amongst those who seem to think that morality is a source of normative pressure that's different to, say, etiquette and accept something like the moral authority platitude are philosophers who caution us against using the term "irrational" too cavalierly. We shouldn't just use it so that we have something that we can say about the wicked to "hit him where it will hurt," as Foot (1978: 152) puts it. There is a special kind of unreasonableness that's supposed to be manifest when a wicked person does some unspeakable thing for no good reason. If we let such things give us the relevant sense of "irrational," it's hard to believe that anyone has taken seriously the idea that *this* is the kind of irrationality a person can be charged with if, say, they don't proportion their beliefs to the evidence. Let's consider a slightly different case, the case of instrumental rationality.

The standards of instrumental rationality are different to those of, for example, etiquette because while both are categorical in our first superficial sense, the standards of etiquette lack authority. But the standards of instrumental rationality are also supposed to differ from those of morality. This difference isn't helpfully brought out by talking about standards that we'd be irrational to violate because that's true of the standards of instrumental rationality. When moral realists tell us that they think moral reasons are robustly normative, we think they want to distinguish the requirements of morality from a larger set of requirements of practical rationality. Consider the enkratic requirement and the requirements of instrumental reason:

- (a) If an agent judges that she ought to  $\phi$  but doesn't intend to  $\phi$ , she's irrational;
- (b) If agent intends to  $\phi$  and believes that it's impossible for her to  $\phi$  unless she  $\phi_2$ s, she's irrational if she doesn't intend to  $\phi_2$ .

These requirements display a problematic kind of generality. They seem to apply to all individuals whatever specific intentions and beliefs they have. If they hold with sufficient generality, they tell us that people shouldn't be akratic and shouldn't be instrumentally irrational but such requirements are non-substantive in an important sense. However silly or trivial our ends might be, we can violate them. In this way, they seem to differ in an important respect from the violation of moral standards.

What's missing, we think, is something about a standard that cannot be recovered from the observations about who or when the standard applies. It has to capture the thing that makes morality and its requirements important in the way that formal requirements of rationality of the kind just mentioned might not be. One suggestion might be that there are values associated with

the moral reasons that call for certain responses and make us criticisable if we're not invested or committed to serving these values well. Such values might be lacking in the case of the rational requirements just mentioned. Let's label a difference here between *thin* and *robust* normative standards and reasons to conform to them.<sup>2</sup> We think that epistemic norms, like the requirements of rationality just mentioned, are thinly normative, not robustly so. And we think that this is sufficient for epistemic realism even though moral realism requires robust normativity.

When it comes to these thinly normative standards we think that it's important to appreciate that their form, shape, or content isn't a good guide to understanding the normative force or pressure that thinkers are under. We think that pairs of agents, say, might be subject to different kinds of criticism for failing to meet some single standard of practical rationality because, say, there are differences in the substantive reasons that apply to them. Compare the case where Agnes violates the instrumental principle while playing pushpin to a case where she violates this trying to find a way to pay for her daughter's surgery. If we proceed in this way, we can make progress by asking whether epistemic realism requires reasons more robust than those associated with instrumental rationality or if our epistemic discourse commits us only to reasons that are thinly normative in a way moral reasons aren't supposed to be.

#### **4. The thin epistemic normativity strategy**

##### **4.1 Thin epistemic normativity**

We think moral error theorists should use the fundamental distinction between thin and robust normativity to reject Cuneo's parity premise. Instead of defending the epistemic error theory, they should reject the idea that epistemic facts are robustly normative in the first place. Instead, they should maintain that epistemic thoughts and claims are only thinly normative. Although they essentially imply norms or standards, they do not imply robust reasons. To put the point another way, although the authority platitude is true of morality, it is not true of epistemology. It is not an essential part of our conception of epistemic facts that they are robustly normative.

---

<sup>2</sup>The labels "thin" and "robust" are from Cuneo (2016). Other labels used in the literature include *strong* versus *weak categoricity* (Joyce, 2001), normativity versus mere *norm-relativity* (Hattiangadi, 2007), *reason-implying* versus mere *rule-implying* normativity (Parfit, 2011), *robust* versus merely *formal* normativity (McPherson, 2011), *normative requirements* versus mere *requirements* (Broome, 2013), and *irreducible* versus mere *reducible* normativity (Olson, 2014).

We can think of two reasons to suggest that epistemic norms are thin in a way that moral norms are not. The first has to do with the possibility of knowledge without accountability. The second has to do with some surprising features of the epistemic “ought”.

Animals are not subject to moral standards. Your dog’s actions cannot be morally right or wrong. More generally, since animals lack the rational capacities needed for agency and attributions of responsibility, they are exempt from robustly normative standards. They cannot have robust reasons that apply to them and the “oughts” that require such reasons do not apply to them.

At the same time, non-human animals can have knowledge. Their beliefs can therefore be epistemically successful or defective even if we wouldn’t say that such attitudes could be justified or that these animals have met their intellectual responsibilities. That they can meet or fail to meet certain epistemic standards would be puzzling if epistemic standards were robustly normative. Since animals are exempt from robustly normative standards, they could not have knowledge if knowledge were like moral rightness in requiring conformity to robust reasons. So, another reason why thin epistemic normativity is more plausible than thin moral normativity is that unlike moral standards, epistemic standards can apply to animals even though they cannot be held accountable.<sup>3</sup>

There is a point, often credited to Harman (1986), which is also important to consider. He thinks that we aren’t subject to criticism for failing to fill our heads with as many of the obvious consequences of uninteresting claims as we can. If you would see, for example, with just a moment’s reflection that something is logical consequence of the fact that the state beverage of Delaware is milk, it’s unlikely that we’d think that this is something you ought to believe. Notice that this isn’t because your epistemic position with respect to this claim might be weak. Far from it. It might be better than many of the things that you may believe. It might be better than many

---

<sup>3</sup> Maria Alvarez pressed the importance of animal knowers on one of us when one of us was too quick to say that justification was required for knowledge. On many standard accounts of knowledge, a belief constitutes knowledge only if the thinker, *inter alia*, is justified in holding her beliefs. We think that this conception of the connection between justification and knowledge is problematic. One idea, worth exploring further, is to think of knowledge as a natural and non-normative property, one that figures in the application conditions of the norms that determine whether a thinker’s beliefs are justified. These norms apply only to thinkers that can be held accountable. That’s why, for example, if Agnes and I know that dinner is ready but I’m a human and Agnes is a dog, only my belief can be justified. My belief’s being justified is down to two things. The first is that it is knowledge or, if you prefer, down to the fact that it stands in some interesting relation to knowledge. The second is that I can be held accountable for meeting a standard like this. (This is where the dog and I differ.) For a defence of this approach to knowledge and justification, see Littlejohn (forthcoming).



of the things that you ought to believe. Still, if you don't believe it, it's unlikely that you can be subjected to criticism for this.

This observation seems to indicate a strange difference between the moral "ought" and the epistemic "ought." In any plausible moral outlook, there will be some things that the agent ought to do that have nothing to do with furthering this agent's non-moral ends but it isn't clear that there is any purely epistemic relation between a thinker and a proposition in virtue of which the thinker ought to believe this thing. In this regard, the connection between belief and epistemic standards seems more akin to the connection between arrangements of cutlery and the standards of etiquette than the connection between actions and moral standards. Think about the kinds of things that it seems a thinker ought to believe. If someone who is generally reliable confides in you that they've been abused, it's probably the case that you ought to believe them. And now compare the strength of epistemic position here to the strength of epistemic position in the case of the logical consequences of the boring belief about Delaware's beverage. We doubt that the reason that one belief is in the ought to be believed pile and the other is not is down to a difference in strength of position. The difference seems to be one of importance.

How should importance be understood? We doubt very much that the importance can be understood epistemically. The beliefs are, for all that's been said, equally likely, equally true, and equally good candidates for knowledge. The difference, we submit, is that the most compelling case of the ought to be believed are cases where two things happen. First, the relevant belief is epistemically kosher. Second, there are non-epistemic reasons to have this belief. They might be non-epistemic reasons to enquire, to settle a question, to trust your informants. We think that the force of the criticism that we level against a thinker for failing to believe what she ought to believe comes from non-epistemic reasons—reasons to form beliefs that meet epistemic standards.

In general, it's hard to see that epistemic standards generate reasons that require us to get into the belief-game. It's interesting to note, however, that epistemic standards generate directives for us even if we're not interested in settling some question correctly. If all the evidence indicates that the state beverage of Delaware is gin, you shouldn't believe that it is milk. There appears, then, to be a weird asymmetry here that we don't find in the moral case. There doesn't seem to be any value associated with this standard that calls for us to get in the belief-game. At best, the standards flag potential moves in that game as incorrect, unwarranted, et cetera. When we

criticise people for failing to form beliefs that meet a standard, it seems to be that that's because of some non-epistemic reason for them to try to form a belief on some issue.<sup>4</sup>

This suggests that although the standard that we use to evaluate a belief might be epistemic and although it might apply to all thinkers at all times, its application doesn't tell us whether a thinker should care about whether they have the sanctioned or right attitudes. We think that when a thinker should care about having such attitudes, it's not purely down to meeting an epistemic standard; rather, it's partially down to non-epistemic factors.

In this regard, epistemic reasons are quite different to moral reasons. Moral reasons are the kinds of things that if you're not committed to serving them to the best of your abilities you are criticisable for this reason. Perhaps this is because, as Raz (2011) has suggested, practical reasons are connected to a good that calls for a response of some kind and epistemic reasons are not inherently connected to any such good. Once we see this, we can see that epistemic reasons must be thinly normative in a perfectly good sense—a thinker's indifference with respect to meeting certain standards isn't itself grounds for criticism.

So, unlike moral facts, epistemic facts are only thinly normative. Unlike moral norms, and just like etiquette, fashion, and the law, epistemic norms are such that there is not necessarily a robust reason to conform to them. Consequently, for the same reason that it does not entail a legal error theory or a fashion error theory, the moral error theory does not entail an epistemic error theory. Moral claims are systematically mistaken, Garner, Joyce, and Olson tell us, because they essentially aim at representing robustly normative facts, which are non-existent. Etiquette, fashion, and the law, in contrast, do not embody the same mistake. Merely thinly normative facts about what is correct or incorrect relative to some standard are not objectionably strange. So, facts about what is what is legal, fashionable, decorous, polite, or grammatical are not strange. Legal claims and fashion claims are thus often true, even in the error theorist's normativity-free universe. We think error theorists should say the same about the epistemic domain. Facts about what is known, epistemically justified, epistemically rational, and the like are only thinly normative, and therefore not strange. Epistemic claims are thus often true, even in the error theorist's ontology. The moral error theory does not entail the epistemic error theory.

---

<sup>4</sup> See Thomson (2008) for a discussion of a hybrid view of the epistemic "ought" according to which a thinker ought to believe only when the belief is epistemically kosher and furthers something non-epistemic.

Many will recoil at the thought that epistemic normativity might be thin like legal or etiquette normativity. What about the authority platitude? Isn't it just part of our ordinary conception of epistemic facts that they imply robust reasons? To view epistemology as only thinly normative is to reject that authority platitude. Consequently, doesn't our proposal fail to capture our conception of epistemic facts and claims? Isn't it, at best, revisionist?

We do not think so. As far as epistemology is concerned, the authority platitude is not, in fact, a platitude. Robust normativity is *not* part of our conception of epistemic facts and claims. A thin conception of epistemic normativity suffices to capture our conception of the epistemic domain. This is because, first, although thin epistemic normativity is incompatible with the authority platitude in epistemology, it still accommodates a number of closely related platitudes—call this set of closely related platitudes the *epistemic norms* platitudes. Second, although they do not include the authority platitude, the epistemic norms platitudes suffice to capture our conception of epistemic facts and claims. Therefore, thin epistemic normativity suffices to capture our conception of the epistemic domain. Error theorists can thus plausibly deny the authority platitude in epistemology.

#### **4.2 The epistemic norms platitudes**

*The epistemic evaluations platitude.* Essential to the epistemic domain are attributions of positive and negative epistemic evaluations or, as Cuneo calls them, epistemic merits and demerits. Epistemic thought and discourse essentially involve, for example, evaluations of beliefs as knowledge, as epistemically justified, epistemically rational, and the like. It essentially involves, in other words, viewing things like beliefs as epistemically good or bad.

This feature does not require robust normativity, however. Evaluations, merits, and demerits can very well be thinly normative. The ascription of merits and demerits is also essential to legal discourse, etiquette discourse, and other paradigmatic thinly normative domains. Essential to those domains are evaluations of acts as lawful or unlawful, legal or illegal, decorous or indecorous, polite or impolite, stylish or unfashionable, etc. Part of making legal claims, etiquette claims, and fashion claims, in other words, is to evaluate them as good or bad from a legal, etiquette, or fashion standpoint. A thin conception of epistemic normativity has therefore no problem capturing the epistemic evaluations platitude.

*The importance platitude.* Epistemic norms, values, and goals seem to matter. Getting things right and avoiding error is undoubtedly important for human beings. Epistemic standards thus seem

to be a very important kind of standards. But this is compatible with such standards being only thinly normative. Thinly normative standards can very well be highly important. After all, we care deeply about many norms and merits that are also paradigm cases of thinly normative standards. The law, etiquette, and many other conventions are highly important domains. Yet they are not robustly normative.

*The deontology platitude.* The epistemic domain seems, at least in part, deontological. Epistemology, many think, is essentially about what doxastic attitudes we are epistemically permitted, prohibited, or obligated to have. Talk of permission, prohibition, and obligation appears essential to epistemic discourse.<sup>5</sup> This third platitude is also compatible with our proposal. Just like values and merits, not all permissions and prohibitions are robustly normative. Many thinly normative domains are deontological. The law, etiquette, grammar, games, and many other sets of norms are best put in terms of permissions, prohibitions, and requirements.

*The inescapability platitude.* You cannot opt out of epistemic evaluations and requirements by not caring about them or by not caring about getting things right. Whether your belief counts as knowledge, as epistemically justified, as epistemically rational, and the like does not depend on what you want or care about. But thinly normative standards can very well be inescapable or desire-independent in that way. Your acts are legal or illegal, decorous or indecorous, and the like whether or not you care about being law-abiding or decorous. Epistemic norms are no different according to our proposal. Even though they are only thinly normative, they apply to you whether or not conforming to them would promote your desires, interests, values, and the like.

*The epistemic reasons platitude.* Epistemic facts plausibly imply epistemic reasons. If believing that P is epistemically justified or unjustified, then there is an epistemic reason to believe that P or not to believe that P. Moreover, epistemic reasons seem inescapable. Whether there is an epistemic reason for you to form a particular doxastic attitude does not depend on what you care about.

---

<sup>5</sup> This platitude is, admittedly, more controversial. Some (e.g. Alston (1988) and Levy (2007)) have argued that deontology has no place in epistemology because belief is involuntary. We set this issue aside, in what follows, and assume that deontology does have a place in epistemology. It is worth noting, however, that our proposal might help resolve the tension between epistemic deontologism and doxastic involuntarism. If deontic 'oughts' imply voluntariness, this will only be true of robustly normative 'oughts'. Thinly normative deontological claims, in contrast, do not imply 'can'. The law forbids you to steal, for instance, even if you are an unwilling kleptomaniac. For more on this point, see Côté-Bouchard (forthcoming).

Whether or not you care about getting things right or avoiding error, there are epistemic reasons for you not to, for example, engage in wishful thinking.

But this fifth feature of the epistemic domain is also compatible with a thin conception of epistemic normativity. This is because, although they do not entail *robust* reasons, thinly normative facts still imply thin, domain-specific reasons. As Olson points out:

‘[O]ught’ and ‘reason’ are both polysemous terms. [...] For example, it seems perfectly fine to say that according to Catholicism there is *reason* to abstain from eating meat on Fridays and to mean by this simply that it is incorrect according to the rules of Catholicism to eat meat on Fridays. [...] The claim that there are reasons for writers in English not to split the infinitive might simply mean that splitting the infinitive is inappropriate according to (some) rules of grammar; the claim that there are reasons for male guests to wear a tie at formal dinners might simply mean that this is required by the rules of etiquette. (Olson 2014: 120)

Even when there is no robustly normative reason to do as, for example, the law or etiquette says, it remains true that there are legal reasons to do what the law requires and etiquette reasons to avoid what etiquette proscribes. It is just that in these situations such reasons lack robust normative authority. Our proposal is that epistemic reasons are just like that. While epistemic norms do imply epistemic reasons, they do so in the same, thinly normative sense that, for example, legal norms imply legal reasons. Just like legal reasons, epistemic reasons are not essentially robust reasons.

This explanation also accommodates the inescapability of epistemic reasons. Thin reasons can very well be inescapable and desire-independent. As Olson writes:

A soldier might not *desire* to comply with the general’s order and he might have no other desires whose satisfaction would be promoted by his compliance, but he can still be said to have reasons to comply with the general’s order since complying with the orders of those of superior military rank is part of the role of being a soldier. The same goes for chess players and football players; they might not desire to play by the rules and they need not even desire to win. Agents can occupy roles they have no desire to fulfil and engage in activities they have no desire to succeed in. (Olson 2014: 121)

Similarly, there are legal reasons and etiquette reasons for you to avoid illegal and improper behaviour, whether or not you care about the law or etiquette. In the same way, there are

epistemic reasons for you to conform to epistemic norms regardless of what you care about. It is just that, just like legal reasons, epistemic reasons are not essentially robust reasons.

*The universality platitude.* We have been comparing epistemic norms with norms that are clearly not universal or absolute, but rather conventional and relative. Obviously, there is not one framework or code of law or etiquette that is absolutely “correct” or “valid.” Things are not illegal or indecorous *period* or absolutely, but only relative to a specific code of law or etiquette. Moreover, facts about what is legal, stylish, or polite are—or are grounded in—facts about social conventions.

Epistemic norms, in contrast, do not seem to be relative and conventional in this way. Rather, much like moral claims, epistemic claims seem to be true universally and absolutely. For most epistemologists, there is not a plurality of equally valid and conventional epistemic frameworks, but rather one absolute and universal system of epistemic standards. It *does* make sense to claim that a belief is epistemically justified, rational, or correct *period*. This is because—again, according to most philosophers—these epistemic norms are plausibly not just a matter of arbitrary conventions that we came up with. Rather, they are discoverable facts that exist independently of such conventions.<sup>6</sup>

Once again, this feature of epistemic facts is compatible with a thin conception of epistemic normativity. Norms can be universal and nonconventional without having robust normative authority. Thinly normative claims, in other words, do not necessarily imply relative and conventional norms. One way in which this can be the case is if the norms in question are grounded in or derived from some fundamental end or standard that is itself absolute and universal. Importantly, this is precisely how many epistemologists construe the epistemic domain. Many of them argue that epistemic norms are those that derive, somehow, from the fundamental goal or norm of *truth* or *knowledge*. This is in line with what Cuneo calls the *substance* platitudes. It is a platitude, Cuneo points out, that

Propositional attitudes have epistemic merits such as being justified [...] in virtue of their being such as to represent reality aright, being likely to represent reality aright, or being such that an agent has done what she ought to in attempting to represent reality aright. (Cuneo 2016, 74).

---

<sup>6</sup> This is, of course, controversial. Many epistemologists defend forms of epistemic relativism and conventionalism. See for instance, Kusch (2002), Neta (2007), Hazlett (2013), and Kerr and Carter (2016). See, for example, Boghossian (2006), White (2007), and Carter (2016) for criticism.

It is part of our conception of the epistemic domain, in other words, that it has to do with getting things right and not getting things wrong.

If something like this is correct and if truth is universal and not just a matter of convention, then there might very well be only one true system of epistemic norms. Crucially however, none of this entails that this one epistemic system or framework has robust normative authority. It does not mean that there is necessarily a robust reason to achieve the absolute goal or norm of knowledge or truth. The thin conception of epistemic normativity is therefore compatible with epistemic facts being universal, and not relative or conventional.

### 4.3 Is something missing?

A thin conception of epistemic normativity therefore suffices to capture the set of platitudes, which we have labelled the *epistemic norms* platitudes. But do those platitudes suffice to capture our ordinary conception of epistemic facts and claims? Do we *also* need to capture the authority platitude to really capture the epistemic domain? Does our ordinary conception of epistemic facts and claims include the authority platitude *in addition* to the epistemic norms platitudes? Are we missing something platitudinous if we only say that epistemic facts are highly important, inescapable, universal, deontological, and that they imply epistemic reasons and epistemic evaluations?

We do not think so. Our ordinary conception of epistemic facts and claims does not plausibly include the idea that epistemic facts are not only norm-implying, important, inescapable, universal, and so on, but *also* robustly normative. Perhaps that claim would be plausible if it were also platitudinous that there cannot be situations where we have no robust reason to do as epistemic norms require. That is, if it were obvious and uncontroversial that we necessarily have a robust reason to conform to epistemic norms, then perhaps the epistemic norms platitudes would be insufficient and we would also need to capture the authority platitude.

But this is plainly not the case. It is not difficult to imagine cases where, many would agree, conforming to epistemic norms would be so trivial, reckless, or harmful that there would be simply no robust reason to do so.<sup>7</sup> Granted, proponents of robust epistemic normativity have a ready response to such cases: there is still a *pro tanto* robust reason to conform to epistemic

---

<sup>7</sup> Of course, there might still be *epistemic* reasons to do so and doing so might still be *epistemically* valuable in these cases. But as we have seen, this is compatible with those epistemic reasons and values lacking robust normativity. As an analogy, in cases where there is no robust reason to conform to the law, there are still *legal* reasons to do so; that act still has positive legal value. See Côté-Bouchard (2017a) for a more extensive discussion of this point.

norms in those cases. It's just that it has almost no weight in comparison to reasons against doing so. But while this response is not obviously mistaken, it is not obviously sound either. After all, someone could make the same claim about many norms, including paradigmatic examples of thinly normative ones. For example, in response to cases where the demands of etiquette are completely trivial or reckless, someone could easily say that there is still a robust, *pro tanto* reason to conform to etiquette, but that it has almost no weight. Yet few would be persuaded by such a defence of the robust normativity of etiquette. In the same way, one could legitimately reject that defence in the epistemic case (see e.g. Whiting (2013), Hazlett (2013), Wrenn (2017), and Côté-Bouchard (2017a)). It is therefore not platitudinous that there must always be a *robust* reason to conform to epistemic norms, including in cases where doing so is completely trivial or reckless.

More generally, it is not a platitude that epistemic facts and claims are robustly normative *in addition* to being important, inescapable, universal, deontological, and related to epistemic reasons and epistemic value. Capturing what we called the epistemic norms platitudes is sufficient for capturing our conception of epistemic facts and claims.

## 5. Why not say the same about morality?

If we are right that the epistemic domain includes the epistemic norms platitudes, but not the authority platitude, then our strategy avoids the revisionism objection; thin epistemic normativity suffices to capture our conception of epistemic facts and claims.

But could this prove too much? Could this be a pyrrhic victory for error theorists who want to avoid the epistemic error theory, but keep the *moral* error theory? After all, one might claim, what we have said against robust epistemic normativity could also be said against the authority of moral facts. That is, one might claim that the same case could be made against robust *moral* normativity and in favor of *thin* moral normativity.

The worry is that our strategy might face a dilemma. Either our case against robust epistemic normativity is persuasive or it is not. If it is not, then our strategy is not successful, and Cuneo's argument remains intact. If it is, then it may very well undermine the error theorist's thesis that *moral* thoughts and claims essentially aim at representing robustly normative facts. This would be



unacceptable for moral error theorists. If moral claims are only about *thinly* normative facts, then their truth does not require the existence of strange normative facts.<sup>8</sup>

We do not think that our case against robust epistemic normativity applies or extends to moral normativity, however. That is, we do not think that the foregoing undermines the moral error theorist's conceptual claim. This is because there remain key differences between the moral and epistemic that make thin epistemic normativity more plausible than thin moral normativity (and robust moral normativity harder to deny than robust epistemic normativity). Put differently, there are differences that make the authority platitude more plausible in ethics than in epistemology.

First, as we have already seen, many would agree that conforming to epistemic norms is not necessarily robustly good and that violating epistemic norms is not necessarily robustly bad. Once again, it is not hard to imagine situations where knowing the truth, avoiding error, and following the evidence would be so trivial, detrimental, or pointless that there would plausibly be no robust value in doing so. Similarly, we can easily conceive of cases where there is plausibly nothing robustly bad in violating epistemic norms. Granted, some have maintained that epistemic successes like true belief have intrinsic robust value.<sup>9</sup> Nevertheless, many epistemologists and value theorists remain unpersuaded and agree that, in some cases, there is simply no robust value in getting things right. But the same idea is much less plausible in the moral case. It is much harder to deny that moral value necessarily constitutes robust value. It seems more plausible, in other words, that there is necessarily something robustly good in doing as morality requires.

Second, recall the point about animals. Non-human animals can meet or fail to meet certain epistemic standards, but they are not subject to moral standards. Moreover, animals are exempt from robustly normative standards. Therefore, a plausible explanation for the exemption of animals from moral standards is that we see those standards as robustly normative. Conversely, a plausible explanation for the fact that animals are not exempt from epistemic standards is that they are *not* robustly normative. Finally, one feature that is commonly attributed to morality – and to robust normativity more generally—is an *autonomy* from the non-normative. Many think we cannot derive robustly normative conclusions from non-normative premises alone.

---

<sup>8</sup> See Finlay (2008), (2014) for a critique along the lines of the latter branch of the dilemma.

<sup>9</sup> See e.g. Zagzebski (2003), Lynch (2004), Horwich (2006), and Kvanvig (2008). For criticism, see Littlejohn (forthcoming), Maguire and Woods (MS), and Thomson (2008).

One important way in which the autonomy of the robustly normative manifests itself is in the phenomenon of first-person deliberation. From the first-person point of view, it seems that non-normative considerations alone cannot settle for you the robustly normative questions you ask yourself in deliberation. Enoch explains:

Because only normative truths can answer the normative questions I ask myself in deliberation, nothing less than a normative truth suffices for deliberation. And because the kind of normative facts that are indispensable for deliberation are just so different from naturalist, not-obviously-normative facts and truths, the chances of a naturalist reduction seem rather grim. [...] The gap between the normative and the natural, considered from the point of view of a deliberating agent, seems unbridgeable. (Enoch 2011: 80).<sup>10</sup>

Robustly normative deliberation seems autonomous from the non-normative. Yet the same does not seem true of deliberation that is distinctly epistemic. To deliberate is to ask oneself questions like “what should I do?” or “should I  $\phi$ ?” But suppose you ask yourself whether you should believe, e.g., that your partner is cheating on you. Even though you can conceivably answer this question by determining whether it would be *desirable* or *beneficial* to have that belief, you can also answer it by settling for yourself the question *whether they really are cheating on you*. Importantly, the former, kind of doxastic deliberation is not strictly speaking relevant for epistemology and epistemic facts. Answers to the question whether believing that p would be desirable (regardless of the truth of P) are not epistemic claims. Only the latter—the kind of doxastic deliberation that can be answered by settling the question “is p true?”—results in a distinctly epistemic conclusion.

Crucially however, this latter kind of doxastic deliberation is *not* autonomous from the non-normative. It can be settled by non-normative facts alone from the first-person perspective.<sup>11</sup> Once again, suppose you ask yourself “should I believe that my partner is having an affair?” On the epistemic interpretation, this question is settled by answering the question “is my partner really having an affair?” But non-normative premises can suffice to settle that question for you. Therefore, unlike moral deliberation, non-normative facts alone can suffice to settle epistemic doxastic deliberation from the first-person perspective. All you need to settle the question

---

<sup>10</sup> See also van Roojen (2015: 258).

<sup>11</sup> See Hieronymi (2005, 2006) for a helpful discussion on the relationship between non-normative considerations and the closure of doxastic deliberation.

whether you should *epistemically* believe that p is to settle the question whether p is true and non-normative facts can plausibly suffice to settle this kind of question for you. Epistemic claims, then, do not seem to have the autonomy from the non-normative that is distinctive of robustly normative claims like moral claims.<sup>12</sup>

For these reasons, we do not think that our defence of thin epistemic normativity extends to moral normativity. There remain differences between the epistemic and the moral that make robust moral normativity much harder to deny than robust epistemic normativity. Our strategy does not, therefore, undermine the moral error theory.

## 6. Conclusion

We have argued that a certain interesting and influential line of objection to the moral error theory rests on a mistake about the commitments of the epistemic realist. It would appear that epistemic standards are like moral standards in some interesting respects, such as their apparent *a priori* status and their independence from human conventions. They are different to moral standards in an important respect, which is that there do not appear to be robust reasons associated with these standards. In this regard, epistemic standards are closer to the standards of etiquette. If it really matters whether we meet them or not, it's not because of the standards themselves but because of some extrinsic reason to care about meeting them. Once we see that epistemic standards are thin, we can see that the chance that the companions in guilt strategy will succeed is thin, too.

## References

- Alston, W. P. (1988). The deontological conception of epistemic justification. *Philosophical Perspectives* 2: 257-299.
- Boghossian, P. (2006). *Fear of Knowledge: Against Relativism and Constructivism*. Oxford University Press.
- Broome, J. (2013). *Rationality through reasoning*. Chichester: Wiley Blackwell.
- Carter, J.A. (2016). *Metaepistemology and Relativism*. Basingstoke: Palgrave Macmillan.
- Case, S. (Forthcoming). From Epistemic to Moral Realism. *Journal of Moral Philosophy*.
- Côté-Bouchard, C. (2017a). Is Epistemic Normativity Value-Based? *Dialogue*, 56(03), 407-430.
- Côté-Bouchard, C. (2017b). Epistemological closed questions: A reply to Greco. *Manuscrito*, 40(4).
- Côté-Bouchard, C. (Forthcoming). 'Ought' implies 'can' against epistemic deontologism: beyond doxastic involuntarism. *Synthese*.
- Cowie, Christopher. Why Companions in Guilt Arguments Won't Work. *The Philosophical Quarterly* 256: 407-22.

---

<sup>12</sup> For more on the same theme, see Jenkins (2007), Heathwood (2009), and Côté-Bouchard (2017b).

- Cuneo, T. (2007). *The Normative Web: An Argument for Moral Realism*. Oxford University Press.
- Cuneo, T. (2016). Destabilizing the Error Theory. In T. Cuneo, P. Schmechtig, & M. Grajner (Eds.), *Epistemic Reasons, Norms and Goals*. De Gruyter.
- Derek, P. (2011). *On What Matters*. Oxford University Press.
- Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press.
- Finlay, S. (2008). The error in the error theory. *Australasian Journal of Philosophy*, 86(3).
- Finlay, S. (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University Press.
- Foot, P. 1978. *Virtues and Vices*. Oxford University Press.
- Harman, G. (1986). *Change in View*. MIT Press.
- Hattiangadi, A. (2007). *Oughts and Thoughts: Rule-Following and the Normativity of Content*. Oxford University Press.
- Hazlett, A. (2013). *A Luxury of the Understanding: On the Value of True Belief*. Oxford University Press.
- Heathwood, C. (2009). Moral and epistemic open-question arguments. *Philosophical Books*, 50(2).
- Hieronymi, P. 2005. The Wrong Kind of Reason. *Journal of Philosophy* 102: 437-57.
- Hieronymi, P. 2006. Controlling Attitudes. *Pacific Philosophical Quarterly* 87: 45-74.
- Jenkins, C. (2007). Epistemic Norms and Natural Facts. *American Philosophical Quarterly*, 44(3).
- Joyce, R. (2001). *The Myth of Morality*. Cambridge University Press.
- Joyce, R. (2005). *The Evolution of Morality*. Bradford.
- Kerr, E., & Carter, J.A. (2016). Richard Rorty and Epistemic Normativity. *Social Epistemology*, 30(1).
- Kusch, Martin (2002). *Knowledge by Agreement: The Programme of Communitarian Epistemology*. Oxford University Press.
- Levy, N. (2007). Doxastic Responsibility. *Synthese* 155 (1):127-155.
- Littlejohn, C. Forthcoming. The Right in the Good: A Defense of Teleological Non-Consequentialism in Epistemology. In K. Ahlstrom Vij and J. Dunn (ed.), *Epistemic Consequentialism*. Oxford University Press.
- Mackie, J. (1977). *Ethics: Inventing Right and Wrong*. Penguin.
- Maguire, B., & J. Woods. MS. Explaining Epistemic Normativity.
- McPherson, T. (2011). Against Quietist Normative Realism. *Philosophical Studies*, 154(2), 223-240.
- Neta, R. (2007). In defense of epistemic relativism. *Episteme*, 4(1).
- Nobis, Nathan. 2005. *Truth in Ethics and Epistemology: A Defense of Normative Realism*. (Dissertation, University of Rochester).
- Olson, J. (2011). In Defense of Moral Error Theory. *New Waves in Metaethics*, 8(770), 62-84.
- Olson, J. (2014). *Moral Error Theory: History, Critique, Defence*. Oup Oxford.
- Raz, J. (2011). *From Normativity to Responsibility*. OUP Oxford.
- Garner, R. T. (1990). On the genuine queerness of moral properties and facts. *Australasian Journal of Philosophy*.
- Rowland, R. (2013). Moral Error Theory and the Argument From Epistemic Reasons. *Journal of Ethics and Social Philosophy*, 7: 1-24.
- Streumer, B. (2017). *Unbelievable Errors: An Error Theory About All Normative Judgments*. Oxford University Press.
- Thomson, J.J. 2008. *Normativity*. Open Court.
- van Roojen, M. (2015). *Metaethics: A Contemporary Introduction*. Routledge.
- White, R. (2007). Epistemic subjectivism. *Episteme*, 4(1).

Wrenn, C. (2017). Truth is not (Very) Intrinsically Valuable. *Pacific Philosophical Quarterly* 98 (1):108-128.

Zagzebski, L. (2003). The Search for the Source of Epistemic Good. *Metaphilosophy* 34 (1-2):12-28.