

Crytonormative judgments

Alex Worsnip

Forthcoming in the *European Journal of Philosophy*; please cite published version

Rex, the secret agent. Rex is a secret agent for a major Western nation. He has captured a terrorist suspect whom he believes to know the location of a ticking bomb. Short of time, and unable to extract information from the suspect by conventional means, Rex turns to less conventional methods that border on torture. Rex's partner agent protests, "you shouldn't do this – it's morally wrong!" Rex replies, "I don't give a damn about the morality of it. The most important thing is that we save the millions of innocent lives that are in danger right now."

Roland, the sexist headhunter. Roland is a headhunter who recruits people for major financial firms. Anyone who observes Roland's headhunting practices can tell that he tends to favor hiring men over women, other things being equal. Confronted about this by a client, Roland frankly admits that he prefers not to hire women. His client asks him, "how can you justify such a sexist attitude?" Roland shakes his head knowingly, having had this conversation hundreds of times before, and replies, "you have to understand – I'm not making any value judgment here. It's just how things are – the rough-and-tumble of the financial world is no place for a woman. It's just a biological fact that they're not suited to it. I wish it were different – but we have to be realistic."

Christine, the modernizing politician. Christine is a politician in a Western democratic state. She belongs to a traditionally left-of-center party which used to be quite radical, but went through some rough times with the electorate, and she is part of a small cabal of senior figures within the party who are intent on modernizing it and bringing it toward the center-ground of the nation's politics. One of the things she particularly wants to avoid is what she calls the 'ideological' character of the party's previous leadership. She states that she believes that modern politics has moved on from ideological dispute and arcane philosophical debates about justice and equality, and that the party now has to simply focus on being pragmatic and doing what works.

In my view, the cases of Rex, Roland and Christine all exemplify a single phenomenon which I call *cryptonormative judgment*. Very roughly, a cryptonormative judgment is a judgment which is presented by the agent as non-normative (either generally or in some particular respect), but which is in fact normative (either generally or in that particular respect).

The idea of cryptonormativity is familiar from debates in social theory, social psychology, and continental political philosophy,¹ but it has to my knowledge never been treated in analytic metaethics, moral psychology or epistemology except in passing. This is somewhat surprising, since

For helpful discussion and comments, I'm grateful to audiences at the University of Notre Dame and at conferences hosted by the Northwestern University Society for the Theory of Ethics and Politics (NUSTEP) and by the British Society for Ethical Theory (BSET), and in particular to Raff Donelson, Tamar Gendler, and Jonathan Leader Maynard.

¹ See, e.g., Habermas (1987); Anderson (1992); Kolodny (1996); Kamolnick (1998); King (2009); Maes *et al* (2012).

cryptonormative judgments are familiar and pervasive features of ordinary life. In this paper, I hope to show that cryptonormative judgments are not only philosophically interesting in and of themselves, but that they shed light on further phenomena of both philosophical and political interest.

In part 1, I will give a definition of ‘cryptonormative judgment’ and makes some remarks about it. In part 2, I will give my basic diagnosis of the cases we began with as cryptonormative judgments, and suggest that such cases are pervasive. In part 3, I consider several possible attempts to reject the diagnosis, and reject those attempts. In part 4, I argue that the phenomenon of cryptonormative judgment shows us that normative judgment is a mental state that can be quite deeply non-transparent to its bearer, in a way that is not, for example, assimilable to the phenomenon of self-deception. In part 5, I draw some lessons from this conclusion for existing debates in metaethics and moral psychology. In particular, I argue that the phenomenon of cryptonormative judgment sheds light on debates over amoralism and lends some support to a picture of normative psychology that links normative judgment constitutively to motivation. In the conclusion, I make some remarks about the social and political importance of cryptonormativity and its insidiousness, looking forward to future work.

1. Defining ‘cryptonormative judgment’

Although cryptonormative judgments are, I have just said, pervasive, ‘cryptonormative judgment’ is far from a folk term. I am thus using the term stipulatively here. First, let me be clear about how I am using ‘judgment’. Sometimes ‘judgment’ is used in philosophy in a way that is supposed to signal an occurrent act of the mind, as opposed to a settled and permanent state. However, in metaethics, ‘judgment’ is not generally used in this way. Rather, ‘judgment’ is used as a kind of placeholder term for a normative mental state, intended to be neutral as to whether such mental states are ordinary beliefs, some special kind of belief, pro-attitudes, some hybrid of the above, or some other mental state.² I am using ‘judgment’ in this latter way. This should make it clear that I am also not using ‘judgment’ to refer to any kind of utterance, or sentence, or proposition. Though at times it will be important for me to consider what someone’s utterances do or don’t tell us about their judgments, the judgment itself is a mental state, not an utterance.

That clarified, here is my stipulative definition of a cryptonormative judgment:

Cryptonormative judgment. A cryptonormative judgment is a judgment an agent makes such that

- (i) The judgment’s content is at least partially normative in character
- (ii) Given a presentation of the judgment that makes the particular normative content of the judgment fully explicit, the agent denies (or would deny) that she makes the judgment
- (iii) The agent presents the judgment such that, if her description were accurate, the judgment would either (a) lack normative content entirely or (b) lack a particular kind of normative content which the judgment in fact has.

² See, e.g., Sturgeon (2007).

A few remarks on this definition.

First, let me say something about what I mean by ‘presents’. To present a judgment as non-normative may in some cases involve doing something in language. One simple kind of example is that where one makes an explicit claim, for example, that the judgment in question is “not a value judgment”, or (in the case of a judgment that is in fact moral) “not a moral judgment”. A second example is that where one is asked to say what one thinks of Φ -ing, and one says that the *only* stance one takes towards Φ -ing is that of believing that it has some descriptive property. This can also be communicated by phrases like “all I’m saying is...” A third is an explicit conjunction of the non-normative claim with the denial that one makes the normative claim: for example, “it’s not that I think Φ -ing isn’t the right thing to do – it’s just that Φ -ing isn’t possible in these circumstances.”

However, there are also subtler cases of presenting a judgment as non-normative which are harder to detect merely by their surface linguistic features. Sometimes, the mere utterance of a non-normative claim in the stead of a normative one can count as presenting one’s normative judgment as non-normative. Of course, not every utterance of a non-normative claim counts as doing this. What matters is whether the agent’s non-normative claim is an attempted articulation of *the judgment that is (in fact) normative*. There may often be no easy way of telling simply from the words that the agent utters whether this is so. As I have already said, what I am primarily concerned with here is the agent’s mental states, and what she does in language is only a guide to this.³ Indeed, I think that an agent counts as satisfying condition (iii) above if she merely presents her judgment *to herself* as non-normative, without any utterance whatsoever being involved. Though I have not been able to give a perfectly precise analysis of what I mean by “presents” here, I hope I have said enough to allow the reader to at least latch on to central cases of what I mean to pick out.

Second, note that the definition does not require that the agent have the concept *normative judgment*. This would be the case if we restricted ‘cryptonormative judgment’ just to the cases in which the agent explicitly denies *that she makes a normative judgment*. We don’t want this result, since *normative judgment* may be a concept that many ordinary people lack. Fortunately, the simplest cases of cryptonormative judgment clearly do not require the agent to have this concept. Suppose someone judges that one ought to Φ , but denies that she judges that one ought to Φ . When asked to explain her views with respect to Φ -ing, she simply presents her judgment as the belief that Φ -ing is what most people do. Her judgment then meets all the above conditions, but clearly we have not required that she possess the concept *normative judgment*.

That said, the definition will also cover more sophisticated cases where *normative judgment* or a closely related concept (such as the more common *value judgment*) is directly employed by the agent. Suppose an agent affirms that Φ -ing is common sense, but denies that he makes any ‘value judgment’ thereby. Unlike judging that something is what most people do, there is a case to be made that judging

³ So there may be cases in which the normative content of her judgment is expressed semantically by what is said, others where it is pragmatically implicated, and still others where what is said is merely evidence for the judgment. These could all be cases of cryptonormative judgment if (and only if – in particular, the agent must also deny that she makes the judgment, regardless of how its content is related to what she says) our conditions are satisfied: we do not need to decide between them for the purposes of our analysis.

something to be common sense *is* a normative judgment.⁴ But even though the agent says that Φ -ing is common sense, the fact that he explicitly says that his judgment is not a value judgment means that if his description were accurate, the judgment would lack normative content. So the case meets condition (iii). And once the normative content implicit in the claim that Φ -ing is common sense is made explicit, he would deny that he makes any such judgment, so he also meets condition (ii). So this kind of case is also covered by the definition.

Clearly the simpler and more sophisticated cases have some differences, but they both exhibit the same phenomenon of interest for our purposes. So I have no compunctions about covering them both under the same definition.

Notice also that the definition allows for a judgment to be cryptonormative if it is presented by the agent as having some kind of normative content, but a different kind than it actually has. So, for example, an agent might present a moral judgment as being a judgment of prudence, and this could count as a cryptonormative judgment. It might be proposed that such a case is better described as a ‘cryptomoral judgment’, with ‘cryptonormative’ reserved for cases where the judgment is presented as entirely non-normative. But I think that the cryptomoral case is interesting in a similar way that the more narrowly cryptonormative case is, and so once again I am happy to include both under the definition of ‘cryptonormative’, which is in any case a term of art.

Finally, it is helpful to distinguish between *sincere* cryptonormative judgments and *insincere* cryptonormative judgments. In the sincere case, both the agent’s denial that she makes the normative judgment and her presentation of the judgment as non-normative are sincere. In the insincere case, at least one of the two is not. In this paper I will be concerned primarily with *sincere* cryptonormative judgments. Insincere cryptonormative judgments are less interesting from the point of view of thinking about mental states, since the possibility of deliberately misreporting or lying about your mental states is not in any way a puzzling one.⁵ So for our purposes today, sincere cryptonormative judgments are our primary topic, and I will often omit the word ‘sincere’ while tacitly restricting myself to such cases. However, from the point of view of social theory, insincere cryptonormative judgments are also of great interest, so I do not want to exclude them from the definition entirely.

2. The pervasiveness of cryptonormative judgments

I hope that it will be initially plausible to most readers that cryptonormative judgments are possible – I do not intend this to be a striking or controversial claim, but rather to draw attention to a particular human phenomenon that has gone undertheorized in analytic philosophy to date. Beyond the mere claim that cryptonormative judgments are possible, however, I also want to argue that they are

⁴ A different interpretation of the case has it that *Φ -ing is common sense* is not itself a normative judgment, but that under many conditions those who rely on the claim that Φ -ing is common sense are usually also making a further normative judgment that Φ -ing ought to be done. In that case, this example may still meet the definition of cryptonormative judgment in the same way that original simple case did – though with respect to the judgment that Φ -ing is to be done, not with respect to the judgment that Φ -ing is common sense. Nevertheless, the more sophisticated structure seems possible in at least some cases.

⁵ As Jonathan Leader Maynard pointed out to me, however, an insincere representation of yourself as merely making a non-normative judgment can play on the same confusion in your *audience* that the agent herself is subject to in the cases I am interested in.

pervasive. A good place to start is with the claim that cryptonormative judgments take place in the recognizable and archetypical cases of Rex, Roland, and Christine.

Someone like Rex, the secret agent, thinks of himself as unconcerned with the niceties of morality. Perhaps this is a self-image he has built up, and even become attached to, over many years. Moreover, his idea of someone concerned with morality is one of someone who is weak, like his partner agent – unable to do what must be done because of emotional attachment. It may be of someone who is out of touch with the realities of the world, or who thinks too theoretically or abstractly. Perhaps Rex even implicitly thinks that to the extent one is concerned with compliance with moral judgments, one is necessarily committed to a kind of absolutism that is indifferent to consequences or outcomes. But what Rex does not realize is that his own judgment – that saving the millions of innocent lives at risk is more important than refraining from torture – is in fact itself a moral one. The judgment concerns the relative moral priority of two considerations. So Rex misdescribes his own mental state. It is a moral judgment, even though he sincerely thinks that it isn't. As such, it fits our definition of a cryptonormative judgment.

Roland, the headhunter, goes further than Rex: he denies making any kind of value judgment whatsoever. According to Roland, all he has is a descriptive belief that women are not suited to the rough-and-tumble of the financial world: no kind of normative judgment comes into it. But again, Roland is wrong. First, the notion of being 'suited' in this context has normative content. There is no straightforward set of criteria for what it is to be 'suited' to a workplace that can be given without saying what a *suitable* worker in that workplace would be like. Second, even if we were to concede to Roland that 'women are not suited to the financial world' is a purely descriptive judgment, Roland would still be making a normative judgment in taking this purported descriptive fact to have an (apparently decisive) upshot for who should be hired. Roland, too misdescribes his own mental state. He makes a cryptonormative judgment.

Christine, the politician, is similar to Roland, but in an explicitly political context. She assumes that there is a value-neutral characterization of 'doing what works' that avoids the need for any kind of philosophical or ideological framework. But what she describes as 'working' will smuggle in her own value judgments about what the goals of policy should be – what counts as a policy 'working' – which likely reflect particular normative assumptions that she takes for granted, without realizing that she is making them. Again, she makes a cryptonormative judgment.⁶

I have picked these three characters to be recognizable, common archetypes. Most moral philosophers have probably found themselves at some point in conversation with someone like Rex, asking them what they do for a living. Such people are fond of making it clear how 'impractical' they think moral thinking is, and profess not to be guided by such judgments. But of course, such people do make many moral judgments. They make moral judgments about the fairness of being required to work overtime, about their duties toward their children and partners, about their entitlement not be overtaxed, and so on in countless other ways. Often they simply do not reflect on these myriad

⁶ Note that her judgment that we ought to do what works may not be cryptonormative – she may not be inclined to deny that she makes that judgment, even when its normative content is made explicit. What are cryptonormative, rather, are her judgments about which policies "work".

judgments; other times they have explicit rationalizations for how such judgments aren't really moral in nature. Either way, these are all cryptonormative judgments.

Likewise, Roland and Christine should be recognizable. In many professions and in politics, value-disagreements are often stifled by appeals to 'facing facts' or 'facing reality' that encode cryptonormative assumptions about what the normative import of those facts or reality is. Newspaper op-ed columns frequently oppose 'moral' or 'ethical' considerations to the promotion of practically beneficial consequences without explanation.⁷ Many people think of rationality or what 'makes sense' as non-normative concepts. And some politicians, policy-makers and economists frequently use the notion of 'what works' (and related notions) in precisely the way that Christine does in our example.⁸

Moreover, philosophical debates themselves are not immune from cryptonormativity. The recent 'realist turn' in political philosophy often appeals to a need to eschew "moralizing" political thought for a kind of "pragmatism" that is, in my view, cryptonormative.⁹ And for many years (though to a far lesser extent now), some philosophers underappreciated the normative character of various non-moral judgments such as those concerning rationality, prudence, and epistemic justification – leading to cryptonormative claims in epistemology and ethics.¹⁰

3. Arguing against alternative diagnoses

But some may wonder if there are alternative diagnoses available for many of these cases. To argue that these cases are not cryptonormative judgments, one will need to hold that they ultimately either (i) don't really count as normative judgments at all or (ii) are not really disavowed by the agent. I will consider several strategies for each route, building from the least challenging up to the most challenging, as I see it. (This takes me in an order that is formally odd, but materially better for the exposition.)

Strategy 1 for resisting (ii): the agent is insincere

One option would be to deny (ii) by denying that the agents in our examples are sincere in their disavowal that they make a normative judgment. This would still count them as cryptonormative judgments on my definition, but not as *sincere* cryptonormative judgments – and I do want to argue that *sincere* cryptonormative judgments are pervasive.

In at least many of these cases, however, is strategy is just too *ad hoc*. Consider Rex, for example. I do not see any reason to think that Rex is deliberately lying when he denies that he makes

⁷ For just one extremely striking recent example, see Moore (2013).

⁸ See, e.g., Blair (2005).

⁹ See, e.g., Geuss (2008); Bourke (2009); Williams (2005: ch. 1); Philp (2010). I develop the charge of cryptonormativity against political realism (both in philosophy and in mainstream political discourse) in co-authored work in preparation (Leader Maynard & Worsnip ms.).

¹⁰ For a striking example, see Mackie (1977: ch. 8), who having declared that all value-claims are false, goes on to consider which (fictional) system of morality would be most "practical". Though Mackie thinks that the moral claims within this most "practical" system are themselves false, he does not seem to worry about the normativity of the metaprinciple enjoining us to choose the most "practical" system of moral imperatives (where practicality constitutes contributing to a well-ordered and well-functioning society).

moral judgments. What reason would he have to do so? It's not like his case against his partner agent fundamentally depends on being presented in non-moral terms. Rex's confusion about what makes something a moral consideration leads to his mistaken impression that he does not make moral judgments. It's part of his being genuinely mistaken here that he is not being insincere. Now, there may be some politicians who are deliberately disingenuous in presenting their agenda as value-neutral. But it seems to be an unwarranted empirical assumption to assume that every politician who presents their agenda this way is being insincere.

Strategy 1 for resisting (i): the agent is an amoralist (or anormativist)

Remember that we said that Rex thinks that moral considerations are absolute in a way that makes them blind to consequences. It might be thought that it just follows straightaway from this that Rex does *not* make the judgment that torture is morally justified. For Rex, one might think, has a particular view of morality, one on which morality consists of absolute prohibitions and is blind to consequences. Even if this moral view is mistaken, the argument goes, it might still be Rex's view. It's just that, as he says himself, Rex does not care about morality (perhaps partly *because* of his absolutist moral views!) He is an amoralist. We can therefore say that he is not mistaken about his own mental states.

This is not a plausible diagnosis of Rex's state of mind. First, if Rex's judgment that the lives of millions of innocent people are more important than the rights of the suspect to be protected from torture is not a moral judgment, then what is it? It is certainly not a judgment about what is prudentially good for Rex. It seems to play the exact functional role of a moral judgment – justifying the prioritization of one good over another, and effectively holding Rex's own action accountable to the people whose lives are at stake. Moreover, we can easily imagine Rex being disposed to have the central reactive attitudes associated with a moral judgment that he ought to save the millions: being disposed to feel guilt if he failed to do so; being disposed to blame others (such as his partner) who do otherwise, and so on, without changing the fact that he sincerely denies making a moral judgment. By contrast, he may be disposed not to have any of these functional or reactive associations with the converse judgment. On prominent and plausible theories, the way to characterize the distinctiveness of moral judgment in particular (and the way to individuate different kinds of normative judgments more broadly) is in terms of the distinctive reactive attitudes that it is associated with.¹¹ Given that, if Rex has these reactive attitudes in association with his judgment that saving the millions is more important, then that judgment is a moral judgment, regardless of what he says about it, and is thus cryptonormative.¹²

¹¹ For a detailed work on the connection of morality and moral judgment with the reactive attitudes, see Darwall (2006). The inspiration for this kind of approach is from Strawson (1962).

¹² So being able to say whether someone makes a cryptonormative judgment *does* turn on the correct account of what a moral judgment is. Since the answer to this question is non-obvious, there will be cases where whether someone makes a cryptonormative judgment is non-obvious. This is fine, I think – we should expect classifying agent's mental states to be hard in some cases. Nevertheless, there are also clear cases, and I think Rex's judgment will count as a moral judgment on any plausible theory of what the mental state of moral judgment is.

So Rex simply misdescribes his own state of mind. It's true that he officially avows the view that morality is blind to consequences. But this is not for him a substantive view which he has about what morality requires;¹³ rather, he thinks (wrongly) that it's constitutive of something being a moral judgment that it not concern itself with consequential considerations.¹⁴ It's very important to draw this distinction between a substantive view of what (e.g.) morality requires and a constitutive view of what (e.g.) moral judgments concern themselves with. In general, an agent's answers to the two questions will differ – unless they think that only those who share their specific first-order moral views even count as making moral judgments!

Fortunately, there is a good test for whether a given view is substantive or constitutive: we can focus on how the agent classifies *others*. Suppose Rex had the substantive moral view that torturing is actually wrong. Then there would be no reason to expect him to classify others with consequentialist views as making non-moral judgments; he could simply classify them as making moral judgments, albeit ones he himself disagreed with. But if Rex had the constitutive view that moral judgments do not concern themselves with consequences, then he would classify others with consequentialist views as making non-moral judgments – since he thinks that it's part of what it is to be a moral judgment to ignore consequences. As I imagine Rex – and as I think it is most plausible to characterize real-world agents like Rex – he is in this second category.¹⁵ That is to say, he would classify others who make consequentialist judgments as making non-moral judgments. In fact, Rex's substantive moral judgment is really that one morally ought to torture the suspect, as his reactive attitudes reveal. But because he is confused about what a moral judgment is, he misdescribes his own mental states, and denies that he makes this moral judgment.

Classifying Rex as an amoralist, I have just tried to show, conflates distinct phenomena. But even if the amoralism strategy worked in the case of Rex, it does not generalize in any obvious way to cases like that of Roland, who denies that he makes any normative judgment at all. One might try to say of Roland that he is a kind of general 'anormativist', who ignores normative judgments entirely and simply acts from his desires. But this reply once again passes over the fact that Roland may have all the functional dispositions and reactive attitudes associated with normative judgment. What seems much the better diagnosis is to say that Roland simply doesn't realize that he is making a normative judgment when he characterizes women as ill-suited to the workplace.

Strategy 2 for resisting (ii): linguistic error precludes genuine disavowal

In the light of the previous reply, one may now worry that the cryptonormativity account does not make the agents in question very deeply mistaken about their own mental states after all. For it might

¹³ Relatedly, classifying Rex as making a cryptonormative judgment does not require us to have a prior fix on the correct first-order normative theory, even though it requires us to have an account of what a normative judgment is. These are theories of quite different things. A good theory of what a normative judgment is will leave lots of room for disagreement about which considerations actually carry normative weight, and as such our theory of cryptonormative judgments will too.

¹⁴ I don't think Rex is alone here: I've experienced, and heard reports of others experiencing, this view often from students and others. It's interesting to ask sociologically what brings about this false belief.

¹⁵ Of course, one *can* develop the case so that Rex is in the first category, too. Not every agent in the ballpark of our description of Rex will be making a cryptonormative judgment. This is immaterial to my argument.

be starting to look like Rex (who I'll again use as my main example) just misuses the term 'moral judgment'. Suppose, by analogy, I've been misinformed and think that 'ennui' refers to a state of intense excitement. I might now sincerely assent to the sentence 'I am in a state of ennui', even though the sentence actually expresses a falsehood. But this is a rather uninteresting phenomenon, and it is not a very deep way of my mental states being non-transparent to me! For all that has been said, my mental states are fully transparent to me. What I am mistaken about is the meaning of a word, not my own mental states. In fact, we might even say that in the relevant sense, I don't really claim that I am in a state of ennui. I utter the words, 'I am in a state of ennui', but my mistake about the meaning of the word 'ennui' means that I don't really think I'm in a state of ennui at all. And one might worry that the same is true of Rex: his error is merely a linguistic one about the meaning of 'moral judgment'.

However, I think that this too misdiagnoses what is going on with Rex. My claim has not been that Rex doesn't understand the meaning of 'moral judgment'. Rather, it is that Rex doesn't understand what a moral judgment *is*. To try to convince you that these two possibilities are distinct, here is an analogy. Suppose that I announce that I think our friend Tom is depressed, and you announce that you disagree. Suppose also that I am wrong and you are right. There are at least three possibilities here:

- (1) You and I understand equally well what depression is, and I am just mistaken about whether Tom is in this state
- (2) I lack a full understanding of what it is to be depressed, and thus mistakenly diagnose Tom as depressed
- (3) I fail to know the meaning of the word 'depressed', and thus announce something which I do not really believe

What I am insisting on here is the distinctness of (2) from both (1) and (3). (2) is a way of being genuinely mistaken about whether Tom is depressed, even though this mistake is traceable to a misunderstanding of depression in a way that the mistake in (1) is not. (2) is the sort of misunderstanding that might be a result of not really even having been seriously depressed oneself (amongst other causes); unlike the error in (3), it might not necessarily be put right by consulting a dictionary. We could even say that this misunderstanding results in a kind of impoverishment of my concept of depression. But that does not make it a trivial error, in the way that (3) is a trivial error. We would not say in (2) that you and I are just using 'depressed' in two different ways, and that both are right in our own idiolect. (I can have an impoverished concept of something relative to yours, without our talking past, or failing to contradict, each other.) I have made a mistake deeper than that. Moreover, it is *correct* to attribute the belief that Tom is depressed to me; that would not be the case in (3).

We can now cross-apply this to the case of Rex. If Rex really fully understood what it is for something to be a moral judgment, then he would probably classify himself as making a moral judgment. But, as the above analogy shows, that he lacks this full understanding does not just make it the case that he makes a mistake that is trivial or merely linguistic in classifying himself as not making a moral judgment. That ignores the possibility analogous to (2) – that Rex is mistaken about what a

moral judgment is (as opposed to the meaning of ‘moral judgment’). Just as it was correct to attribute to me the belief that Tom is depressed, it will then be correct to attribute to Rex the belief that he does not make any moral judgment (which would not be correct if he was merely misusing a phrase). Still, Rex is wrong: he does make a moral judgment.

As I have already alluded to, the question of what constitutively makes something a moral judgment, rather than some other kind of judgment, is non-trivial philosophical question, and one about which there are competing philosophical theories.¹⁶ It is one which is independent from the question of which considerations actually matter morally, since there can be moral judgments based on considerations that do not constitute genuine moral reasons. But nor is it a mere matter of linguistic stipulation. When someone offers a theory of the distinctiveness of moral judgment, for example in terms of its connection with particular reactive attitudes, she is not just stipulating a meaning for the term ‘moral judgment’. So, in parallel with the depression case, it is possible for someone like Rex to be mistaken about what moral judgment is without lacking a basic grasp of the term ‘moral judgment’.¹⁷

Again, things are even clearer with Roland. Roland might be a perfectly competent user of the term ‘normative judgment’ (or, more realistically, ‘value judgment’). He simply doesn’t realize that the judgment that women are ill-suited to the workplace is a normative judgment, because he hasn’t thought hard enough about exactly what makes something a normative judgment.¹⁸ The same goes for Christine, who hasn’t reflected enough to realize that her judgments about what works involve her in certain normative judgments. This is a failure of understanding rather than a linguistic error.

Strategy (ii) for resisting (1): the mental state is alief

In important recent work, Tamar Gendler (2008a, 2008b) has argued that there are mental states that typically move us to action in the way that beliefs do, but which lack the cognitive characteristics of belief. Gendler calls these states ‘aliefs’. One cognitive feature of aliefs that helps to explain why they are not beliefs is that we will not typically endorse them as beliefs.¹⁹ So one might think that what I have been classifying as cryptonormative judgments are in fact aliefs. This might encourage a critic to claim that the agents in our examples do not really make the normative judgments I attribute to them:

¹⁶ For example, Darwall’s reactive attitude-based theory, which I have already alluded to above.

¹⁷ Of course, some agents will just be in simple linguistic error. But we cannot chalk every purported instance of cryptonormativity up to linguistic error, since some substantive error about what moral (or normative) judgment is also possible, and (I think) widespread. I’ll come back to this point in the final section, and reinforce in a final way how the error need not be purely linguistic.

¹⁸ At the outset I said that one could make a cryptonormative judgment without having the concept *normative judgment*. But it may look like what I’ve just said traces a cryptonormative judgment to a misunderstanding about what makes something a normative judgment, and that may seem to require the agent to have the concept *normative judgment*. However, the mistake could simply manifest itself, not specifically in explicit denials that one makes a normative judgment, but in refusal to assent to the content of the judgment presented in explicitly normative terms. In that case one shows an implicit misunderstanding through failure to apply normative terms correctly. (Also, note that in many cases it’s not a misunderstanding of what makes something a normative judgment, but a misunderstanding of what makes something a more specific kind of judgment such as a moral judgment. And probably more people have the concept *moral judgment* than have the concept *normative judgment*.)

¹⁹ Though, for a different characterization, on which they are ‘in-between’ beliefs, see Schwitzgebel (2010).

rather, they alieve them. This might be particularly encouraged because Gendler has used the alief framework to analyze implicit bias (Gendler 2011), which reminds us of Roland.

Gendler is surely right that one can be acting as if one believed something without actually believing it. And I am open to the possibility that the notion of alief, or something like it, is what is needed to make sense of this phenomenon. However, I do not think that the cases of cryptonormative judgment that I have described are cases of alief. There is a big difference between Roland's brand of implicit sexism and the sort of implicit sexism that Gendler takes as paradigmatic. Gendler is interested in people who reflectively endorse egalitarian, liberal or feminist sentiments, but end up manifesting their bias unconsciously in action and decision. But Roland is not like this. Roland's bias is essentially overt, insofar as he admits that he prefers not to hire women. What Roland denies is that this constitutes any kind of normative judgment. So he is not explicit about his attitude *qua normative judgment*; still, he is explicit about the attitude under at least some description, and he reflectively endorses it: he says that women are ill-suited to the workplace.

Consequently, Roland is not in some way acting contrary to his judgments. Nor does he act in some arational or habitual manner. So he is not an aliever. He is a believer in denial. Gendler's theory points us to the ways in which certain cognitive characteristics are important for bona fide belief, but this does not go so far as to make it impossible to be mistaken about what one believes, and nothing Gendler says suggests that she intends her theory that way.

The same thing holds for Rex and for Christine. For example, Rex reflectively accepts that it is more important to save the millions – he just denies that this is a moral judgment. It's characteristic of a cryptonormative judgment that it is endorsed under some misleading description that is non-normative. This is a fundamental difference with an alief, which is typically not endorsed under any description. Relatedly, cryptonormative judgments are sensitive to factors which their bearers take to be truth-indicative, in a way that aliefs are not.

4. Cryptonormative judgments as failures of self-knowledge

I have argued that cryptonormative judgments are pervasive. Since agents who make sincere cryptonormative judgments are mistaken about whether they make normative judgments, this suggests that normative judgments are mental states which are non-transparent to their bearers.

On its own this may not come as much of a surprise. The strong doctrine that one enjoys genuine infallibility with respect to one's mental states was once popular, counting amongst its adherents at least arguably Descartes (M III: 37), Hume (T 1.4.2.7) and Wittgenstein (PI: 246-7). Few contemporary philosophers would defend this very strong claim.²⁰ However, many philosophers nevertheless hold that we enjoy some kind of qualitatively unique privileged access with respect to our own mental states.²¹ I venture that many philosophers would side with Donald Davidson (1987: 441) when he writes that there is “an overriding presumption that a person knows what he or she believes.”

²⁰ An exception may be Burge (1988).

²¹ See, e.g., Heil (1988); Davidson (1984; 1987); Shoemaker (2009).

The pervasiveness of cryptonormative judgment suggests that this ‘overriding presumption’ may be a mistake, at least with respect to the normative domain.²²

Moreover, I think that cryptonormative judgment illustrates the way in which one can be mistaken about one’s own mental states in a quite deep way. In particular, I want to argue that they show that one can be mistaken about one’s own mental states without self-deception, which some people might think of as the paradigmatic way to be mistaken about one’s own mental states. To do this, I need to argue that cryptonormative judgments are not (necessarily) instances of self-deception. So they cannot be assimilated to self-deception. The attempt to perform such an assimilation would not be an attempt to deny that there are cryptonormative judgments, or that they are not common, as the strategies considered in the previous section were. Rather, the idea is to query how deep mistakes about one’s own mental states can go.

In order for the assimilation to achieve this, one needs a notion of self-deception that can be characterized as only involving failure of self-knowledge in some limited or shallow sense. This puts some constraints on the notion of self-deception in play. On one hand, it won’t do to just define self-deception in terms of being mistaken about one’s own mental states.²³ If it is defined this way, then trivially cryptonormative judgments (and all other instances of being mistaken about one’s own mental states) would be instances of self-deception; but this would obviously not show that they do not involve especially deep failures of self-knowledge.

On the other hand, someone who wants to assimilate cryptonormative judgment to self-deception also won’t want a notion of self-deception that precludes cases of mistaken beliefs about one’s own mental states from counting as self-deception. For example, Gendler (2007) characterizes self-deception as a kind of pretense: when you are self-deceived about some proposition *p*, you do not really believe *p*, but rather pretend (in a particular way) that *p* is true (in fact, Gendler urges, you may really believe the negation of *p*). But then, if *p* is some proposition about your mental states, it follows that when you are self-deceived about your own mental states in Gendler’s sense, you do not actually have a false *belief* about your own mental states. So this notion of self-deception is also not good for the purposes of someone who wants to claim that the only way to be mistaken about your own mental states is through self-deception.²⁴

What is needed for the assimilator to steer between these two dangers is some notion of self-deception which cashes self-deception out in terms of one’s believing *p*, perhaps due to some motivational bias, despite at some level knowing or being in a position to know that *p* is false. This

²² In different ways, numerous philosophers have held that it is a mistake more generally. See, e.g., Ryle (1949: ch. 6); Williamson (2000: ch. 4); Schwitzgebel (2011). For a reply to Williamson see Berker (2008).

²³ C.f. Shoemaker (2009). As well as being too broad, this definition is also too narrow, since it limits cases of self-deception to beliefs about one’s own mental states. But clearly one can be self-deceived about matters that do not concern one’s own mental states. For example, one can be self-deceived about whether one’s long-lost relative is still alive. The over-narrowness and over-broadness here both result, I think, from the same mistake: that of confusing the notion of being deceived *about* oneself with the notion of being deceived *by* oneself.

²⁴ Perhaps there are good reasons to favor Gendler’s characterization of self-deception as pretense, or at least to think that many of the cases that are usually referred to as self-deception involve pretense rather than belief. If so, that just reinforces the problem for executing the strategy of assimilating cryptonormative judgment to self-deception. For cryptonormative judgments are genuine instances of mistakes about one’s own mental states.

vague characterization leaves a lot of room open.²⁵ Still, I think it is determinate enough for the purposes of arguing, against the assimilator, that cryptonormative judgments need not involve any self-deception in this sense.

First, when you make a cryptonormative judgment, you may not be in a position to know that you are making the judgment that you make. Consider again Rex. Rex is precluded from knowing that he makes a moral judgment by his failure to fully understand what a moral judgment is. Consequently, he is not in fact in a position to know that he is making a moral judgment. Of course, one could interpret “position to know” loosely enough that one is in a position to know *p* when one could, by careful reflection that one may not actually be fully capable of, come to know *p*. This would then count Rex as being in a position to know that he makes a moral judgment. But this is obviously too loose a gloss on “position to know” to be of use in an account of self-deception. For example, it would count one as self-deceived about complex *a priori* mathematical truths that are beyond one’s understanding.

Second, relatedly, cases of cryptonormative judgment do not require the belief that you do not make the normative judgment in question to be due to motivational bias. It can arise from entirely blameless misunderstanding. In some cases it may be that the agent at some level willfully misunderstands, but nothing about the phenomenon of cryptonormative judgment requires this.

It also won’t work to try to assimilate cryptonormative judgments to the phenomenon of fragmented belief (a phenomenon closely related to that of self-deception). In recent work, Daniel Greco (forthcoming) has appealed to fragmentation to defend the ‘iteration principle’ that if you believe something, then you believe that you believe it. The fragmentation strategy, applied to cryptonormative judgment, would say that those who make cryptonormative judgments are fragmented in their mental states. Take Rex, for example. On the fragmentation view, in one sense, or relative to one purpose (perhaps the purpose of making declarations), Rex doesn’t believe that it is morally justified to torture the suspect. But in another sense, or relative to another purpose (perhaps the purpose of making decisions), Rex does believe that it is morally justified to torture the suspect. The idea is then that, as long as we confine ourselves to one ‘fragment’ of Rex’s mental states, holding the relevant purpose constant, we can then apply the iteration principle. So, given the purposes relative to which it’s true to say that Rex believes it’s morally justified to torture the suspect, it’s also true to say that Rex believes that he believes it’s morally justified to torture the suspect. So in an important way, the extent to which Rex can be mistaken about his own judgments here is limited.²⁶

²⁵ One characterization of self-deception that fits this mould is that of Mele (1997), but it leaves substantial room for disagreeing with Mele.

²⁶ Greco himself actually thinks that his strategy preserves the thought that our mental lives are sometimes opaque to us, on the grounds that we’ll sometimes lack accurate *explicit* beliefs about what our implicit beliefs are, where explicit and implicit belief are fragmented (Greco forthcoming: 13). But it seems clear that this still severely limits how opaque our mental lives can be to us. After all, it precludes inaccurate explicit beliefs about our explicit beliefs, as well as precluding inaccurate implicit beliefs about our implicit beliefs. And ultimately, it guarantees that every time I have a belief, I also have a (true) belief that I have that belief – after all, that is just his iteration principle. On his view, it’s impossible to lack that belief. This surely is an important way in which our mental lives are not opaque to us on his view.

The point is even clearer when we move away from a dichotomy between implicit and explicit belief. Greco uses this as a working model for fragmentation to make his points, but has indicated (p.c.) that he thinks it is overly simplified. Rather, he thinks our mental states fragment into as many different compartments as we can have distinct

However, I do not think that even this more limited supposition is defensible. Suppose we implement Greco's strategy by saying that relative to the purposes of making declarations, Rex doesn't believe that torture is morally justified; but relative to the purposes of making decisions, Rex does believe that torture is morally justified. I do not see what justifies the claim that, even just relative to the purposes of making decisions, Rex *believes that he believes* that torture is morally justified. Attributing this belief is not required to make sense of Rex's decisions or dispositional profile: all that is required is the first-order belief that torture is morally justified. That fact is that Rex has no dispositions of *any* sort that indicate that he *believes that he believes* that torture is morally justified. So there does not seem to be *any* purpose relative to which there is any reason to attribute this belief to him.

Finally, one might wonder whether cryptonormative judgments really constitute failures of *self-knowledge*. In the previous part, I traced cryptonormative judgments in part to failures of understanding of what a normative judgment *is*. But it might be thought that this is not a failure of knowledge of one's own mental states. Rather, it is failure of knowledge of how to describe or classify one's own mental states.²⁷

I do not think this distinction holds up. Knowledge of how to describe and classify a subject matter is an important part of comprehensive knowledge of that subject matter. Suppose that an expert biologist and I are both shown a much-enlarged image of a tuberculosis bacterium. Asking 'what is that thing on the screen?', we give different responses. 'It is a tuberculosis bacterium,' replies the biologist. 'It is a long thin purple thing on a blue background', I reply. I take it that it would be absurd here for an observer to react in the following way: 'the biologist and you have equally good knowledge of the tuberculosis bacterium. It's merely that the biologist knows how to describe it and classify it better than you.' For my failure to describe or classify the image on the screen *constitutes* a failure to understand the thing that I'm seeing. And this failure is not somehow trivial or non-substantive; it is a paradigmatic instance of my lacking substantive knowledge than the expert biologist has.

In that case, the subject-matter of my failure of understanding is the thing on the screen, the image of a bacterium. In the case of cryptonormative judgment, the subject-matter of my failure of understanding is my own mental state. Just as the traceability of the former to my inability to describe and classify images of bacteria does not preclude it from being a failure to know about the image of the bacterium, so the traceability of the latter to my inability to describe and classify my mental states does not preclude it from being a failure to know about my own mental states. As I argued earlier, the failure is not merely linguistic; it is a failure of self-understanding.

purposes. But if that's the case, then we don't have a single notion like *explicit belief* which can plausibly be the thing that we lack when our mental lives are opaque to us. Rather, it'll just be the case that when I believe something relative to purpose A, I won't necessarily believe relative to purpose B that I believe it relative to purpose A. But why would we think that the right way to see whether my believing something relative to purpose A is transparent to me is to ask whether I believe relative to purpose B that I believe it relative to purpose A? This doesn't seem like a very fair test! Nor does it seem to reveal any significant opacity.

²⁷ Such an objection might be inspired by Burge (1988: esp. 661-63). Burge makes his point in the context of trying to reconcile a kind of infallible self-knowledge with externalism about mental content. He plausibly contends that knowledge that one is in a mental state does not require that one grasp the conditions for knowing that one is in that mental state. I do not have to deny this; what I do deny (and what does appear to put me at odds with Burge) is that knowledge that one is in a mental state might be blocked by failing to grasp the conditions for *being* in that mental state.

5. Broader metaethical lessons: the internalism/externalism debate

Given that (sincere) cryptonormative judgment is a pervasive phenomenon, we can draw an important general lesson: sincere assent to a normative judgment is not a necessary condition for making such a normative judgment. This conclusion has significant ramifications for some existing metaethical debates – in particular, debates about the relationship of normative judgment to motivation, including debates between so-called motivational ‘internalists’ and ‘externalists’.²⁸

According to a widespread kind of motivational externalism, whether an agent makes a normative judgment is in no way constituted (not even partially) by whether that agent is motivated to act on the judgment in question.²⁹ On this view, the agent’s being motivated could at most be construed as *evidentially* bearing on whether the agent makes the normative judgment in question.

What could be used to motivate this kind of externalist view? One classic way has been to appeal to cases of *amoralist* agents.³⁰ Amoralists are supposed to be agents who make moral judgments, yet are not at all motivated (even in the slightest) to act on them, since they don’t care about morality. Yet the externalist cannot just describe the agent this way in the presenting the cases, since the description begs the question, describing the agent in such a way as to guarantee the truth of externalism. As such, the internalist will disagree on the description of the case as one in which the agent genuinely makes a moral judgment. What the externalist needs, then, is a way of describing the agent’s characteristics that is itself neutral between internalism and externalism, but which is such that it is *then* intuitively plausible to describe the agent as an amoralist, and thus as a counterexample to internalism.

This is exactly what externalists try to do. So what agreed, non-question begging characteristics can the externalist use to describe such an agent? Typically, one such characteristic – perhaps the central such characteristic – that externalists appeal to is that of an agent who *sincerely assents* to a moral claim. So, stories of purported amoralists involve an agent who sincerely assents to a moral claim. Externalists claim, plausibly, that we can easily imagine someone who sincerely assents to moral claims without being motivated. And then argue that given this sincere assent, the internalist diagnosis that no moral judgment is in fact made would be *ad hoc*.³¹

However, if what I have argued here is right, the phenomenon of cryptonormative judgment shows that normative judgment is a state that can be far from transparent to its bearer, and consequently can become divorced from sincere assent. Our cases shows that sincere assent is not a

²⁸ ‘Internalism’ and ‘externalism’ here are being used to mean something entirely different from what they often refer to in epistemology and the philosophy of mind, where internalism is often a claim about access to one’s own mental states and epistemic states. This paper has been putting pressure on such claims of privileged access; the ‘motivational internalism’ I am going on to offer some support for here, by contrast, concerns the relationship between normative judgment and motivation, and is not only distinct from but essentially unrelated to externalism about mental or epistemic content.

²⁹ For examples of this view, see, e.g., Brink (1986), Svavarsdóttir (1999) and Copp (2007: ch. 8).

³⁰ All those authors mentioned in the previous footnote appeal to such cases.

³¹ I offer a similar diagnosis of the debate between internalists and externalists, and a different way of resisting the externalist’s main strategy, in Phillips & Worsnip (ms.).

metaphysically necessary condition for normative judgment – nor is it necessary for us to have adequate epistemic reason to attribute normative judgment. Now, admittedly, what the externalist needs is the claim that sincere assent is a *sufficient*, rather than a necessary, condition for normative judgment.³² But if one can mistakenly think that one does not make a normative judgment – and thus falsely but sincerely deny that one makes a normative judgment – then it seems very plausible that one can also mistakenly think that one *does* make a normative judgment – and thus falsely but sincerely deny that one makes a normative judgment. Just as the natural name for the first phenomenon is *cryptonormative judgment*, the natural name for the second phenomenon is *pseudonormative judgment*. Effectively, what the internalist claims is that supposed amoralists in fact make pseudonormative judgments.³³

Externalists typically claim that it is *ad hoc* to say that so-called amoralists do not really make moral judgments. But if they are willing to countenance the phenomenon of (sincere) cryptonormativity – as I have argued we all should be – then they should also be willing to countenance the phenomenon of (sincere) pseudonormativity without its seeming bizarre, or an *ad hoc* notion to invoke. In fact, the explanations of the two possibilities run in a strongly analogous way. In the case of cryptonormativity, an agent’s understanding of what normative judgment is is overly narrow, thus leading to a misclassification of the agent’s mental states. In the case of pseudonormativity, the agent’s understanding of what normative judgment is overly broad. So, for example, on the internalist story, the putatively amoralist agent might err in self-classification because he thinks that to judge that some act is morally wrong is merely to judge that it falls into the category of behavior commonly described as wrong by ordinary people (and he does make *that* judgment). But if the internalist is right, this does not suffice for making a moral judgment: really judging that some act is morally wrong involves seeing it in such a way that one is at least somewhat motivated to refrain from performing it. No wonder, for the internalist, that sincere assent can come apart from normative judgment, since the agent’s sincere assent rests on what, by internalist lights, constitutes a misunderstanding of what normative (or moral) judgment is.

Indeed, in many purported amoralist cases, I think it is quite natural to claim that the agent may make both a cryptonormative *and* a pseudonormative judgment, and the attribution of the former

³² To be clear, it’s not that this claim actually has to be true for externalism to be true, but rather that it has to be true for amoralism cases to give the cases for externalism any purchase. The externalist could admit that sincere assent is a poor guide to the presence of normative judgment, and maintain that normative judgment is just something you either have or don’t have, deep down, and that identifiable surface characteristics such as assent and motivation are always insufficient to determine its presence or absence. But then by the externalist’s own lights, no case of purported amoralism, described merely in terms of the presence of sincere assent and absence of motivation, will be enough to constitute a counterexample to the internalist view. So for amoralism to constitute an argument for externalism (and it has historically been the main argument), the externalist needs a connection between sincere assent and judgment.

Does the externalist need sincere assent to be a metaphysically sufficient condition for normative judgment, or for it to give us sufficient epistemic reason to attribute normative judgment? That is not totally clear. Things would certainly be cleaner for the externalist if sincere assent were a metaphysically sufficient condition. For otherwise the internalist can simply claim that although generally sincere assent gives us sufficient epistemic reason to attribute moral judgment, the absence of motivation is an adequate defeater for this epistemic reason. In any case, I suggest that cryptonormative judgments undermine both the metaphysical and the epistemic claim.

³³ What R.M. Hare, an arch-internalist, calls an ‘inverted commas moral judgment’ seems to be a particular instance of a pseudonormative judgment. See Hare (1952: ch. 11).

strengthens the case for the latter, thus lending support to internalism. Go back to Rex, for example. As well as sincerely denying that he make the moral judgment that torturing the suspect is right, Rex might well also sincerely assent to the proposition that torturing the suspect is wrong. Earlier I argued that Rex's real moral judgment here is that torturing the suspect is the right thing to do: that is a cryptonormative judgment. We could say that Rex contradicts himself by also judging that torturing the suspect is the right thing to do. But in general it seems that the same considerations that favor counting him as judging that torture is right also favor not counting him as judging that torture is wrong. He thinks he makes the latter judgment, but in fact makes the former. So he makes both a pseudonormative and a cryptonormative judgment.³⁴

The phenomenon of (sincere) cryptonormative judgment, then, weakens the externalist and strengthens the internalist in the dialectic familiar from the literature about normative judgment and motivation. It does not, of course, constitute a knock-down proof of internalism. But I do think that the phenomenon of cryptonormative judgment lends the internalist position some positive support. If it turns out that sincere assent is less important for normative judgment than we thought, we are left wondering what else, if anything, determines whether an agent counts constitutively as making a normative judgment. And the obvious candidate for this role is whether the agent is motivated to comply with the judgment in question. After all, part of what makes it plausible to classify the agents we began with as making normative judgments despite their denials that they do is their dispositional and motivational profile. To crudely oversimplify, with these agents, it is not what they say but what they do that matters in attributing normative judgments to them. And clearly this fits in nicely with the internalist picture.

6. Conclusion: the relevance of cryptonormative judgment to socio-political critique

I have argued for three main claims in this paper. First, cryptonormative judgments are common: they are the natural diagnosis of familiar cases from everyday life. Second, they reveal that normative judgment is a state which can be quite deeply non-transparent to its bearer, in a way that is not, for example, assimilable to the phenomenon of self-deception. Third, they shed light on debates over amoralism and lend some support to a picture of normative psychology that links normative judgment constitutively to motivation.

I hope to have shown that cryptonormative judgments are a philosophically interesting phenomenon, and this is one reason why they are worth having a theoretical account of and a vocabulary to describe. However, I also think that cryptonormative judgments (and cryptonormativity more broadly) are of socio-political interest. I can only touch on the issues they raise here, but want to close by noting some features of cryptonormative judgments that make them potentially insidious.³⁵ Here I drop the restriction to sincere cryptonormative judgments, although I think that sincere

³⁴ Strictly speaking, the pseudonormative judgment that Rex makes is not that torturing the suspect is morally wrong, since he does not make that judgment at all. Rather, it is the actually non-normative judgment that Rex actually makes but presents as the normative judgment that torturing the suspect is morally wrong. This might be, for example, that torturing the suspect is the sort of thing that most moralistic people call 'wrong'.

³⁵ Much more on this in Leader Maynard & Worsnip (ms.).

cryptonormative judgments too can, if anything, be even more insidious in their effects than insincere cryptonormative judgments.

When one presents a normative judgment as non-normative, one effectively takes oneself out of the space of normative justification, or out of the space of a particular kind of normative justification. Such a technique can be used to avoid having to defend a normative view *as a normative view*, as well as to make it seem like an incontestable matter of descriptive fact. This reinforces the point made earlier about how cryptonormative judgments are not based on mere linguistic error. When Rex presents his judgment as non-moral, he does not harmlessly apply an idiosyncratically narrow use of the term ‘moral’ – rather, he avoids having to defend his view on own moral grounds and, just as importantly, having to engage with the rival view that it competes against, portraying them instead – wrongly – as simply orthogonal to one another.

Much existing work calls our attention to the ways in which *normative* concepts and vocabulary can have ideological, power-laden effects.³⁶ But, consistently with this doubtlessly true observation, it is important to see that *non-normative* concepts and vocabulary – especially when offered in the stead of normative claims, or in a purported attempt to eliminate the need for them, or in a way that encodes unarticulated and unrecognized normative assumptions – can have equally or perhaps even more dangerous effects of the same general kind. In this way, the vocabulary of cryptonormativity and of cryptonormative judgments gives us a theoretical framework for talking about, diagnosing, and calling out the conscious and unconscious exercises of this ideological device.

References

- Anderson, A. (1992). ‘Cryptonormativism and Double Gestures: The Politics of Post-Structuralism,’ *Cultural Critique*, 21: 63-95.
- Berker, S. (2008). ‘Luminosity Regained,’ *Philosophers’ Imprint*, 8/2.
- Blair, T. (2005). Speech to the Labour Party’s 2005 conference in Brighton. Published in full at BBC News online, dated 27 September 2005.
- Bourke, R. (2009). ‘Theory and practice: the revolution in political judgement,’ in Geuss & Bourke (eds.), *Political Judgement: Essays for John Dunn*. Cambridge, UK: Cambridge University Press.
- Brink, D.O. (1986). ‘Externalist Moral Realism,’ *Southern Journal of Philosophy*, 24: 23-41.
- Burge, T. (1988). ‘Individualism and Self-Knowledge,’ *Journal of Philosophy*, 85/11: 649-63.
- Copp, D. (2007). *Morality in a Natural World*. Cambridge, UK: Cambridge University Press.
- Darwall, S. (2006). *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge, MA: Harvard University Press.
- Davidson, D. (1984). ‘First Person Authority,’ *dialectica*, 38/2-3: 101-111.
- (1987). ‘Knowing One’s Own Mind,’ *Proceedings and Addresses of the American Philosophical Association*, 60/3: 441-58.
- Descartes, R. (1641/1996). *Meditations on First Philosophy* (ed. & trans. Cottingham). Cambridge, UK: Cambridge University Press.

³⁶ Of course, this was a central theme of some of Marx’s work; others to explore it more recently include (amongst many others) Williams (e.g. 1985: ch. 10), Leiter (2004), and Geuss (e.g. 2010: essay III).

- Gendler, T.S. (2007). 'Self-Deception as Pretense,' *Philosophical Perspectives*, 21: 231-58.
- (2008a). 'Alief and Belief,' *Journal of Philosophy*, 105/10: 634-63.
- (2008b). 'Alief in Action (and Reaction),' *Mind and Language*, 23/5: 552-85.
- (2011). 'On The Epistemic Costs of Implicit Bias,' *Philosophical Studies*, 156: 33-63.
- Geuss, R. (2008). *Philosophy and Real Politics*. Princeton: Princeton University Press.
- (2010). *Politics and the Imagination*. Princeton: Princeton University Press.
- Greco, D. (forthcoming). 'Iteration and Fragmentation,' *Philosophy and Phenomenological Research*. Page references are to the draft manuscript.
- Habermas, J. (1987). *The Philosophical Discourse of Modernity: Twelve Lectures* (trans. Lawrence). Cambridge, MA: MIT Press.
- Heil, J. (1988). 'Privileged Access,' *Mind*, 97/386: 238-51.
- Hume, D. (1739/2000). *A Treatise of Human Nature* (ed. Norton & Norton). Oxford: Oxford University Press.
- Kamolnick, P. (1998). 'Visions of Social Justice in Marx: An Assessment of Recent Debates in Normative Philosophy,' in Panasiuk & Nowak (eds.), *Marx's Theories Today*. Amsterdam: Rodopi.
- King, M. (2009). 'Clarifying the Foucault-Habermas debate: Morality, ethics, and 'normative foundations',' *Philosophy & Social Criticism*, 35/3: 287-314.
- Kolodny, N. (1996). 'The ethics of cryptonormativism: A defense of Foucault's evasions,' *Philosophy & Social Criticism*, 22/5: 63-84.
- Leader Maynard, J. & Worsnip, A. (ms.). 'Normativity and cryptonormativity in political realism,' draft manuscript, University of Oxford/New York University.
- Leiter, B. (2004). 'The Hermeneutics of Suspicion: Recovering Marx, Nietzsche, and Freud,' in his (ed.) *The Future for Philosophy*. Oxford: Oxford University Press.
- Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- Maes, J., Tarnai, C. & Schuster, J. (2012). 'About Is and Ought in Research on Belief in a Just World: The Janus-Faced Just World Motivation,' in Kals & Maes (eds.), *Justice and Conflicts: Theoretical and Empirical Contributions*. Heidelberg: Springer.
- Mele, A. (1997). 'Real Self-Deception,' *Behavioral and Brain Sciences*, 20/1: 91-102.
- Moore, C. (2013). 'This obsession with Ethics is one of the great curses of our time,' *The Telegraph*, 22nd November.
- Phillips, J. & Worsnip, A. (ms.). 'Motivating internalism,' draft manuscript, Harvard University/New York University.
- Philp, M. (2010). 'What is to be done? Political theory and political realism,' *European Journal of Political Theory*, 9/4: 466-484.
- Ryle, G. (1949). *The Concept of Mind*. London: Hutchinson.
- Schwitzgebel, E. (2010). 'Acting Contrary to our Professed Beliefs, or, The Gulf Between Occurrent Judgment and Dispositional Belief,' *Pacific Philosophical Quarterly*, 91: 531-553.
- (2011). *Perplexities of Consciousness*. Cambridge, MA: MIT Press.
- Shoemaker, S. (2009). 'Self-Intimation and Second-Order Belief,' *Erkenntnis*, 71: 35-51.
- Strawson, P. (1962). 'Freedom and Resentment,' *Proceedings of the British Academy*, 48: 1-25.
- Sturgeon, S. (2007). 'Normative Judgement,' *Philosophical Perspectives*, 21: 569-87.

Svavarsdóttir, S. (1999). 'Moral Cognitivism and Motivation,' *Philosophical Review*, 108/2: 161-219.

Williams, B. (1985). *Ethics and the Limits of Philosophy*. London: Fontana.

----- (2005). *In the Beginning Was the Deed: Realism and Moralism in Political Argument* (ed. Hawthorn). Princeton: Princeton University Press.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

Wittgenstein, L. (1953). *Philosophical Investigations* (trans. Anscombe). Oxford: Blackwell.