

Replies

Mark Eli Kalderon

November 14, 2007

§1 THE TROUBLE WITH TERMINOLOGY

Language has a life of its own that transcends authorial intention. This can make communication fraught. Geoffrey Pullum describes some of the obstacles:

Producing language that other people will be able to understand involves not just having a picture in your mind of the scenario... You have to deploy a shared linguistic system, according to established rules, using lexemes of known meaning, to present that picture to others in a way that will work for them. You have to consider whether there are other ways of viewing the situation at hand. You have to examine the wording you have chosen to see if it has ambiguities or unclarities. You have to put yourself in the place of a person who did not [develop the picture], and you have to ask yourself whether they would understand [it]. (*Language Log*, <http://itre.cis.upenn.edu/~myl/language-log/archives/003338.html>)

The development of a specialized vocabulary, explicitly explained, is one way of overcoming these obstacles. But there is no guarantee of success, and terminology can yet mislead.

Philosophical terminology, self-consciously deployed, aspires to the status of stipulative labels. A piece of terminology should be understood in terms of the explicit explanation of it. But the autonomy of language is not so easily quelled. If the terminology is derived from more familiar usage, it can still carry with it unwelcome associations.

There is, perhaps, too much philosophical terminology as it stands. Mindful of this, I tried to adapt prevalent terminology to my own purposes, though sometimes

refining and extending its sense. One prominent example of such refinement is the distinction I drew between noncognitivism, nonfactualism, and expressivism. These terms are often used interchangeably. In using them to draw this distinction, I was thus offering a precisification of established usage. These are not the labels I would have chosen *ab initio*, in the absence of the linguistic history of philosophy. ‘Expressivism’ can mislead (and has misled) because of the semantic and pragmatic interpretations of ‘express’. And ‘nonfactualism’ can mislead because its contrary ‘factualism’ is clearly inapt (a fictionalist discourse lacks a nonfactualist interpretation but it is not fact stating as ‘factualism’ can suggest). ‘Nonrepresentationalism’ would have been a better label. But faced with a choice between using an established, if inapt, piece of terminology and introducing yet more terminology, I decided the conservative route was best.

‘Fictionalism’ itself carries unwelcome associations since it is natural to wonder, as Matthew Chrisman does, in what respects is a fictionalist discourse like genuine fiction. To be clear, I only ever intended a limited analogy: like fictional utterances, the utterance of a sentence in a fictionalist discourse is not assertion and the acceptance of the uttered sentence is not belief in the proposition expressed. In the specific version of fictionalism that I developed, however, I drew upon a further, deeper analogy likening the expressive power of moral language to the framing effects of metaphor.

One, not terribly central, piece of terminology left unexplained was my use of ‘pretense’. Here we can distinguish thin and thick appeals to pretense in the development of a fictionalist alternative to realism. In the thin sense, ‘pretense’ merely registers the mismatch between cognitive appearance and noncognitive reality—in accepting a claim we merely make as if to believe. In the thin sense, pretense is not an explanatory notion. In the thick sense it is. We will return to this distinction in subsequent discussion.

Were I to write *Moral Fictionalism* now, I would eschew much of the specialized vocabulary insofar as possible. Alas, that was not the case, and I welcome the opportunity to explain more clearly some of what I meant to convey.

§ 2 FRIEND

Fictionalism is the view that the sentences of the target discourse express truth-evaluable propositions and yet the acceptance and sincere utterance of these sentences do not require belief in the proposition expressed. The case for moral fictionalism thus takes the form of an argument for moral acceptance being other than belief and an argument against a nonrepresentational semantics. There are many arguments for moral acceptance being some attitude other than belief. In chapter one I offered my own—in part because it was novel and worthy of discussion, and in part because it animates a normative worry I have about the conduct

of public moral discussion, a normative worry discussed in chapter four. One thing I should have been clearer about, and which Stacie Friend is thankfully not taken in by, is that moral fictionalism does not stand or fall with the specific argument for noncognitivism that I offered. If it fails, moral fictionalism may yet be true—it is just that the noncognitive component must be justified on other grounds.

And indeed Friend believes it fails:

I do not think the argument for intransigence is convincing. It is compatible with Putnam's claim—and his comment is cited by Kalderon only as an illustration—that he and Nozick each believe they are right about the moral facts. Where we have given up trying to convince our opponents, this need not be because we take further reasons to be without force; it could be because we take our opponents to be unable to see their force.

I agree that we can legitimately give up trying to convince our opponents because we take them to be unable to see the force of our arguments. Indeed are a number of legitimate reasons we can give up trying to convince our opponents. Mark Sainsbury has suggested to me that one may do so because continuing the discussion would be inimical to one's affection for the interlocutor. But intransigence is more than an unwillingness to convince a specific person of the error of their ways. It is, rather, an unwillingness to keep an open mind, a disposition not to reconsider in light of new evidence. But we can be open minded and yet not persist in fruitless disagreements or disagreements that would endanger valued friendships.

Despite, being unconvinced by the argument from intransigence, Friend recognizes that the noncognitive component of moral fictionalism can be motivated on other grounds. However she worries whether noncognitivism without nonfactualism is a stable position:

On the one hand, HMF requires that our moral practices be suited to the expression of non-cognitive attitudes: else why think that they are non-cognitivist? But on the other hand, HMF requires that they must be unsuited to this expression: else why think there is a cognitivist fiction?

The first horn of the dilemma presupposes that the best case for noncognitivism concerns the way moral language mediates and makes public out noncognitive attitudes. Here we disagree. Though standard noncognitivists often reflect on the expressive capacity of moral language, the central claim of noncognitivism is psychological not linguistic. It is a claim about the nature of moral judgment. Think, for example, of motivational arguments for noncognitivism—that moral judgements

motivate in a way that beliefs could not. These arguments make no mention whatsoever of public language and how it can serve to express noncognitive attitudes.

Concerning the second horn, we also disagree. I would not claim that moral language is unsuited to the expression of noncognitive attitudes despite its cognitive appearance. If, in a fit of bad temper, Bernice calls Edgar a pig, her utterance is well suited to the expression of her contempt for Edgar despite its cognitive appearance, so what's the problem?

The problem, according to Friend, is that we ordinarily regard our moral practice as cognitive and our moral language as assertoric, and she regards it as implausible that we could be so mistaken about this. The pretense involved in genuine fiction is not unwitting, why then should the pretense involved in the moral fiction be? Friend doubts as well, Yablo's corresponding view about mathematics—the vulgar may treat worries about whether the numbers 'really' exist lightly, but no so professional mathematicians. Here, I believe, Friend is simply mistaken about the sociology of mathematics. While philosophers may regard a platonism as a natural, if philosophically problematic, interpretation of mathematics, mathematicians do not. Indeed they tend to deploy constructivist metaphors more amenable to a fictionalist interpretation in describing what they are doing.

One worry that Friend has about fictionalists simply helping themselves to a representational semantics favored by realists concerns the denotations of moral predicates. Like Jimmy Lenman, and unlike Chrisman, she correctly recognizes that fictionalism is a form of ontological noncommitment more akin to agnosticism than atheism. But that's a matter of our attitudes and not metaphysics. Suppose there are no moral properties. What then could serve as the denotations of moral predicates? That's a good question, let me canvass a few possible replies.

First, notice, this is not just a problem for the fictionalist. It is equally a problem for an error theorist. Perhaps, like some error theorists, one could claim that while no moral properties are instantiated uninstantiated moral properties exist and can serve as the denotations of moral predicates. Not all will be happy with this suggestion, and I am not myself blind to the attractiveness of immanent universals. So let's consider something else.

One option, explored by some fictionalists, is to invoke a *thick* pretense in providing the relevant semantics. In the absence of a suitable property to serve as the semantic value of our predicates, we make believe that such properties exist, and the content of this pretense serves as the content of our predicates. However, a fictionalist semantics is controversial and would lack a clear advantage over non-factualist semantics.

One final option, suggested by Mark Sainsbury, is that a Russellian semantics for predicates be abandoned. Instead of thinking of the content of a moral predicate as the moral property denoted, think instead of its content as determined by

its satisfaction conditions. When combined with a negative free logic, this could perhaps avoid the present difficulty.

§3 CHRISMAN

Chrisman, like Friend, doubts whether moral fictionalism can be sustained in the absence of a nonfactualist semantics. Chrisman distinguishes two ways of understanding the noncognitive attitudes conveyed by moral utterances and argues that this is the basis of a dilemma for moral fictionalism:

1. Moral utterances convey the kind of noncognitive attitudes that the cognitivist claims to be constitutive of the acceptance of moral sentences.
2. Moral utterances convey the pretense that something instantiates a moral property.

On the first understanding, moral fictionalism collapses into standard noncognitivism, and on the second understanding, moral fictionalism is phenomenologically implausible. I only ever meant to invoke the thin notion of pretense in the book, and the second option pretense is doing explanatory work. And I agree with Chrisman that it would be phenomenologically implausible. So let's focus on the first horn of the dilemma.

Chrisman spends some time trying to locate my worry about expressivism. However, the target of chapter two was not expressivism *per se*, but an expressivist nonfactualism. Expressivism, as I understood the notion, was a positive semantic claim about the determinants of content—that the content of the relevant vocabulary was partly determined by its use to convey the relevant attitudes. One could be an expressivist in this sense without being a nonfactualist. In the book I gave the following toy example. One might hold that 'good' expresses commendation and that its use to convey this attitude might determine the meaning of 'good'. But that's consistent with a realist interpretation since 'good' might thereby come to refer to that property that warrants commendation, namely, goodness. (Interestingly in recent work, Gibbard retains his commitment to expressivism but is now at least agnostic whether an expressivist discourse would be nonrepresentational.) Not only is expressivism, so understood, consistent with a representational semantics, it is consistent as well with the discourse being genuinely cognitive and assertoric: by commending something in this way one thereby asserts that it is good—an assertion, if sincere, that one also believes.

I agree with expressivism at least to this extent—moral utterances convey noncognitive attitudes, but this falls short of the expressivist's distinctive semantic claim that so doing partly determines the meaning of moral vocabulary. Perhaps it does, perhaps it doesn't. I remained silent on the determinants of moral semantics.

What I am committed to, however, is the denial of an expressivist nonfactualism—hopeless on the grounds that it can never escape the lingering grip of the Frege-Geach problem.

Chrisman begins his essay by describing the attractions of noncognitivism thus: ‘by construing ethical claims as expressing noncognitive attitudes rather than beliefs, one can reject a realist ontology of morality without rejecting moral discourse wholesale.’ I think he is wrong about this. While some may be moved by the rejection of a realist metaphysics (too often by assimilating it too closely to a Cambridge platonism), as I mentioned above, there are grounds pertaining purely to the nature of moral judgment that warrant noncognitivism. And, I would dogmatically add, these are the important grounds from the perspective of moral philosophy. This initial claim ramifies into his discussion of the dilemma. For the fictionalist is allegedly committed to the systematic falsity of moral sentences but inexplicably offers no metaphysical grounds for this.

Unfortunately, hermeneutic fictionalism is not so committed, only a revolutionary fictionalism would be (indeed revolutionary fictionalism is really just a noneliminativist form of error theory). Hermeneutic fictionalism is a kind of ontological noncommitment more akin to agnosticism than atheism as Lenman observes. So the hermeneutic moral fictionalist is not committed to the truth of the moral claims he utters, but neither to their falsity. Since it is not a kind of error theory, there need be no metaphysical grounds for attributing such error. Rather, reflection on what moral judgment must be like given its epistemic role was the grounds for noncognitivism offered in *Moral Fictionalism*.

§ 4 LENMAN

Lenman begins with an appeal to authority that rests on a misconstrual:

The big advantage of fictionalism, thinks Kalderon, the thing that entitles it to be the “default hypothesis” for noncognitivists (p. 118), is that it is “semantically uncontroversial” (p. 119). This is an odd claim at face value. Fictionalism after all is a theory about moral semantics that respected philosophers are on record as rejecting.

Fictionalism is not primarily a semantic thesis. It is a claim about the nature of acceptance. Perhaps, Lenman was misled because some fictionalists also invoke thick pretense to do semantic work. But that is to overlook or ignore the official explanation in the book which was that fictionalism is the view that the sentences of the target discourse express truth-evaluable propositions and yet the acceptance and sincere utterance of these sentences do not require belief in the proposition expressed. The misconstrual is not harmless, since much of Lenman’s criticism depends on it.

Lenman worries how the moral fictionalist can account for the validity of moral argument. This worry would be serious if genuine, for fictionalism would itself be subject to something akin to the Frege-Geach problem. And if it is, the central argument of the book simply collapses.

The Frege-Geach problem is a dilemma for expressivist forms of nonfactualism, the second horn of which concerns, in part, entailment. Suppose that freestanding and embedded occurrences of ‘wrong’ mean the same. The problem is that words can occur in embedded contexts and fail to express the attitudes they do in freestanding contexts. This difficulty would be avoided if the account applied only to the meaning of freestanding occurrences. However, if the account applies only to freestanding occurrences, then it is incomplete, for the expressivist would lack an account of the meaning of moral words in embedded contexts. Furthermore, some guarantee must be given that freestanding and embedded occurrences mean the same despite their meanings being differently determined. For, if freestanding and embedded occurrences differ in meaning, then the expressivist is apparently committed to the invalidity of recognized forms of valid argument (due to the fallacy of equivocation). The validity of an argument is a matter of its premises entailing its conclusion. Entailment supervenes on the semantic properties of the sentences that it relates. That is why the meaning of freestanding and embedded occurrences being differently determined raises a question about entailment.

Notice, however, that *there could be no problem about entailment* on a fictionalist construal of moral discourse—the fictionalist maintains that moral sentences have just the truth-conditional contents that the realist assigns to them, truth-conditional contents that unproblematically determine the relevant entailment relations. The ‘fictionalist semantics’ Lenman claims are canvassed on pages 119–129 are not semantics but rather accounts of what might be pragmatically conveyed by utterances in a fictionalist discourse and are explicitly described as such.

Lenman really misses a beat here, since in the vicinity there is a legitimate query not about the validity of moral argument, but the reasonableness of moral inference. How could a moral argument, as the fictionalist understands it, be a good argument, where the goodness of an argument, here, consists in providing a reason to accept the conclusion? Whereas entailment is an abstract relation on sentences or the propositions they express, inference is a change of view, a transition in the epistemic states of the subject. So whereas (a horn of) the Frege-Geach problem concerns entailment, at least in part, the present challenge concerns inference. This is less of an objection and more of a reasonable challenge. It’s a pity that Lenman’s misconstrual, perhaps abetted by an expressivist tendency to conflate inference and entailment, prevents him from exploring such a challenge.

Let’s consider the argument that is traditionally used to display the failings of Ayer’s emotivism:

1. If lying is wrong, then getting one's little brother to lie is wrong.
2. Lying is wrong.
3. Getting one's little brother to lie is wrong.

First, let's consider the attitude involved in accepting the sentence 'Lying is wrong'. If a competent speaker accepts this sentence, then there will seem to be a reason not to lie, and there will be a tendency for countervailing considerations to cease to be salient, and even where they are, there is a tendency for these to seem to be outweighed or even ruled out as reasons for doing otherwise. Moreover, the reason not to lie does not appear to be a reason that is contingent upon our acceptance of it the way that reasons of taste are. Moreover the reason not to lie seems to be a reason no just for the speaker but for everyone else as well (including the speaker's younger sibling). Indeed, the reason appears to be sufficient to accept the claim on behalf of others. The speaker can seem justified in demanding that others accept this claim about the wrongness of lying and so appropriately reconfigure their affective sensibility.

Similar remarks apply to the attitude involved in accepting the conclusion 'Getting one's little brother to lie is wrong'. The only difference is a difference in the object of the affect.

What's the attitude involved in accepting the initial, conditional premise? The attitudes involved in accepting the antecedent and in accepting the consequent are complex functional states involving the tendency for lying and getting one's little brother to lie to become salient and to present a complex normative appearance. It is natural to think of the attitude involved in accepting the conditional claim as a higher-order functional state that structures the speaker's affective sensibility—specifically, as the tendency to have the affect involved in accepting the consequent when having the affect involved in accepting the antecedent. It is a higher-order functional state since the affects in question are themselves complex functional states. However, it strikes me that the attitude involved in accepting the conditional premise is more than having one's affective sensibility configured in this way. It involves, as well, the endorsement of this affective sensibility. If a competent speaker accepts the conditional sentence, 'If lying is wrong, then getting one's little brother to lie is wrong,' then there appears to be a reason to have an affective sensibility configured in this way, a reason that seems to outweigh or even rule out reasons for having a differently configured sensibility. Moreover, the reason to have this kind of sensibility does not appear to be a reason that is contingent upon our acceptance of it the way that reasons of taste are. Moreover the reason seems to be a reason no just for the speaker but for everyone else as well. Indeed, the reason appears to be sufficient to accept the claim on behalf of others.

The speaker can seem justified in demanding that others accept this claim and so appropriately reconfigure their own affective sensibility.

The reasonableness of this inference, as the fictionalist understands it, should now be plausible, at least to a first approximation. The speaker accepts that lying is wrong and so has the appropriate affective attitude towards lying. Not only is the speaker conditionally disposed to have the affective attitude towards getting his little brother to lie given the affective attitude involved in accepting the wrongness of lying, but he also authoritatively endorses that his affective sensibility be configured in this way. It is reasonable, then, that, given these commitments, he should come to have the affective attitude towards getting his little brother to lie that is involved in accepting that it is wrong to get one's little brother to lie. Given the nature and content of these attitudes and how they bear on the configuration of a person's moral sensibility, the reasonableness of moral inference is *prima facie* plausible.

Lenman is misled by misconstruing fictionalism as a distinctive kind of semantics. He is also, perhaps, misled by a thick understanding of pretense. When raising the worry about entailment he complains that '[h]ard questions about the logic of moral make-believe abound and Kalderon tells us very little about how to address them.' If pretense were doing explanatory work in the account on offer, as it would be if posits of a pretense were the semantic values of moral predicates, the complaint would be justified. But the pretense invoked is thin. It is just the mismatch between cognitive appearance and noncognitive reality—a mismatch that Lenman himself is committed to.

Lenman considers the following principle:

Suppose that moral predicate 'F' denotes moral property *p*. It is fictionally true that *x* is F iff *x* instantiates nonmoral property *p** that would elicit the relevant affective response from a person with virtuous moral sensibility.

and raises a dilemma on the basis of it: 'What's not clear to me is whether the right hand side is supposed to be a move internal to the fictional practice or a move external to it intended to connect pretenceworthiness with virtue.' If external, then there moral fictionalism collapses into a virtue theory, and if internal, then a fictionalist moral practice is a sham.

I am unsure what the internal/external distinction means in this context and Lenman does not explain, but let's consider the dilemma nonetheless.

Why not just drop the 'fictionally' and be a virtue theorist? Because, as I argued in chapter three, the relevant affects cannot be understood apart from the fictional moral properties. So I must embrace the second horn, and face the possibility that morality is a sham.

Now I do entertain the possibility that moral practice might require revision if moral acceptance were intransigent. But that flows from a particular grounds for regarding moral acceptance to be noncognitive. It does not follow that all forms of moral fictionalism are unstable in this way, nor does it follow that the principle adumbrated above results in such instability. Are there general grounds for regarding a fictionalist moral practice to be a sham? Lenman does not say, and to that extent at least, the alleged dilemma is no dilemma at all.

§5 CONCLUSION

I would like to thank the contributors to this symposium for their comments which were both thoughtful and thought-provoking. Moral fictionalism faces many challenges. But again, let me register my hope that if anything is taken away from *Moral Fictionalism*, it is the need to carefully distinguish the semantic, pragmatic, and psychological claims of standard noncognitivism.