

Moral Fictionalism, the Frege-Geach Problem, and Reasonable Inference

Mark Eli Kalderon

5 May 2007—St Andrews Metaethics Workshop

§ I THREE CLAIMS

CHANGE SLIDE Go through outline of talk **CHANGE SLIDE**

It is my sincerest hope that if there is one thing that people take away from *Moral Fictionalism*, it is the recognition that standard noncognitivism involves a syndrome of three, logically distinct claims.

Standard noncognitivists claim that moral judgment is not belief or any other cognitive attitude but is, rather, a noncognitive attitude more akin to desire; that this noncognitive attitude is expressed by our public moral utterances; and, hence, that our public moral utterances lack a distinctively moral subject matter and so are not answerable to the moral facts. Notice, however, that these are logically distinct claims—the first is a *psychological* claim, the second and third, positive and negative *semantic* claims, respectively. We can regiment the familiar technical vocabulary as follow: **CHANGE SLIDE**

- Cognitivists claim that moral acceptance is best explained by moral belief while noncognitivists claim that moral acceptance is best explained by attitudes other than moral belief. This difference is partly a difference in the nature of the attitudes involved in moral commitment (whether they are cognitive or noncognitive) and partly a difference in the content of these attitudes (whether the cognitive attitudes involved, if any, have moral propositions as their objects). Let *noncognitivism*, then, be the claim that the acceptance of a moral sentence is a some attitude other than belief in a moral proposition. This is a *psychological* claim—it is a claim about the nature and content of the attitudes involved in

moral acceptance. Standard noncognitivists typically claim more than this. They make, in addition, positive and negative claims about moral semantics. **CHANGE SLIDE**

- Let *expressivism* be the claim that the meaning of a moral sentence is determined by the attitude its utterance expresses. This is a positive claim about the *semantics* of moral discourse—it provides positive information about the determinants of moral meaning. Standard noncognitivists also typically make a negative semantic claim. **CHANGE SLIDE**
- Let *nonfactualism* be the claim that the meaning of a moral sentence fails to represent a distinctive domain of moral fact. Specifically, moral sentences do not express moral propositions—propositions that represent moral objects, qualities, and relations. The meaning of a moral sentence may be representational, at least in part; but if it is, it represents nothing distinctively moral.

These claims are logically distinct. The meaning of moral sentence may be determined by the noncognitive attitude its utterance expresses, but it may thereby come to represent the moral facts—it may thereby come to represent the existence and distribution of moral objects, qualities, and relations. So, for example, freestanding attributions of goodness may commend their objects, but in so commending they may represent their objects as good. This would be a kind of moral conceptual role semantics and one that determines a representational semantics for attributions of goodness. Horwich (2005) and Wedgwood (2001) have similar approaches to moral semantics and Gibbard (2003) is now at least agnostic about this possibility. Thus, expressivism does not entail nonfactualism. Nor does the converse entailment hold. Consider redundancy theories of truth that maintain that a sentence and the corresponding truth-ascription mean the same. Truth-ascriptions, so understood, do not represent sentences or the propositions they express as instantiating the property, truth (for, if they did, sentence and corresponding truth-ascription would differ in meaning). But the meaning of truth-ascriptions needn't be determined by the noncognitive attitudes expressed. Crucially, and more obviously once stated, the *psychological* claim of noncognitivism entails neither *semantic* claim. **CHANGE SLIDE**

This last observation has an important taxonomic consequence since it allows for the possibility of a fictionalist form of noncognitivism distinct from the more familiar standard noncognitivism. Specifically, if noncognitivism just

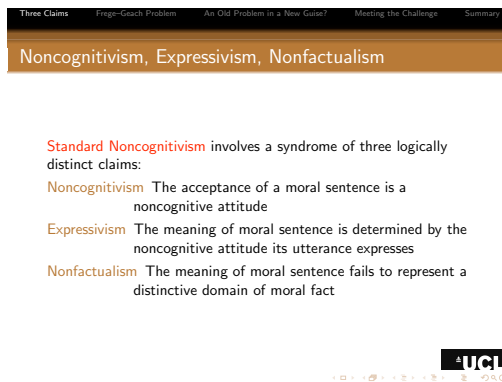


Figure 1: Noncognitivism, Expressivism, Nonfactualism

is the claim that moral acceptance is a noncognitive attitude, then what's to prevent this attitude being expressed by utterances that purport to represent the moral facts?

Moral realists maintain that morality has a distinctive subject matter. Specifically, realists maintain that moral discourse is representational, that moral sentences express moral propositions—propositions that represent the existence and distribution of moral objects, qualities, and relations. Standard noncognitivists, in contrast, typically maintain that the realist imagery associated with morality is a fiction, a reification of our noncognitive attitudes. The thought that there is a distinctively moral subject matter is regarded as something to be debunked by philosophical reflection on the way moral discourse mediates and makes public our noncognitive attitudes (hence the familiar deployment of the Humean metaphor of projection). The realist fiction is understood as a philosophical misconception of a discourse that is not fundamentally representational but whose intent is rather practical.

There is, however, another way to understand the realist fiction. Perhaps the subject matter of morality is a fiction that stands in no need of debunking but is rather the means by which our attitudes are conveyed. Perhaps moral sentences express moral propositions, just as the realist maintains, but in uttering a moral sentence competent speakers do not assert the moral proposition expressed but rather convey by means of it the noncognitive attitude centrally involved in moral acceptance. Moral fictionalism is this alternative to standard noncognitivism. Moral fictionalism is noncognitivism without nonfactualism.

CHANGE SLIDE

Not only does the distinction between noncognitivism, expressivism, and nonfactualism raise the logical possibility of moral fictionalism, but it also makes clear the general form of argument required to establish its truth. To establish the truth of moral fictionalism, it suffices to: **CHANGE SLIDE**

1. argue for the psychological claim of noncognitivism **CHANGE SLIDE**
2. argue against expressivist nonfactualism

Indeed, this was the argumentative strategy that I pursued in *Moral Fictionalism*. I am not going to talk about my case for noncognitivism, but I will talk about the problem that formed the basis of my case against expressivist nonfactualism. **CHANGE SLIDE**

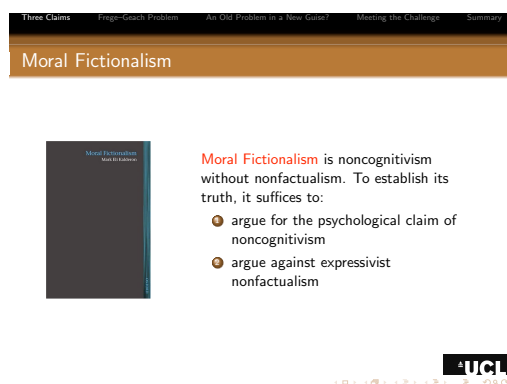


Figure 2: Moral Fictionalism

§ 2 FREGE-GEACH PROBLEM

The Frege-Geach problem was first raised by Ross (1939, 33–34) and independently by Geach (1958, 1960, 1965) and Searle (1962, 1969) and was originally directed at expressivist proposals such as Ayer’s emotivism:

It is worth mentioning that ethical terms do not serve only to express feeling. They are calculated also to arouse feeling, and so to stimulate action. . . . In fact we may define the meaning of the various ethical words in terms both of the different feelings they

are ordinarily taken to express, and also the different responses which they are calculated to provoke. (Ayer, 1946, 108)

Ayer's talk of 'defining' the meanings of moral words suggests that he identified the meaning of a moral word with the attitudes it expresses and provokes.

Ayer's emotivism is subject to the Frege-Geach problem because it is a form of atomistic reduction. It is *atomistic* in that it assumes a one-one correlation between the meanings of moral words and the linguistic actions they perform (such as the expression and provocation of the relevant attitudes). It is a *reduction* in that it identifies the meaning of a moral word with attitudes it expresses and provokes. Atomistic reductions face the following dilemma: **CHANGE SLIDE** 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. **CHANGE SLIDE** This difficulty would be avoided if the atomistic reduction 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). There are other potential problems as well. That freestanding and embedded occurrences differ in meaning is apparently incompatible with the possibility of indirect discourse, the possibility of answering a moral question, the possibility of denying a moral claim, and so on.

The proper target of the Frege-Geach problem is not expressivism per se, but those expressivist semantics, like Ayer's, that take the form of an atomistic reduction. A natural Fregean diagnosis for the failure of atomistic reduction is that it runs afoul of the attitude/content distinction. The meaning common to freestanding and embedded occurrences of a word despite the difference in attitude expressed is its *content*. Ayer's emotivism fails because it mistakes the attitude expressed by a moral word for its content. This is what I called, in *Moral Fictionalism*, the *pragmatic fallacy*. Dewey, in a review of Stevenson, puts the point as follows:

One can agree fully that ethical sentences (as far as their end and

use is concerned) ‘plead and advise’ and speak ‘to the connative-affective natures of men.’ Their use and intent is practical. But the point at issue concerns the means by which this result is accomplished. It is . . . a radical fallacy to convert the end-in-view into an inherent constituent of the means by which, in genuinely moral sentences, the end is accomplished. To take the case in which ‘emotional’ factors accompany the giving of reasons as if this accompaniment factor were an inherent part of the judgment is, I submit, . . . a theoretical error . . . (Dewey, 1945, 702–3)

The challenge posed by the Frege-Geach problem is to explain the meanings of words in embedded contexts without assuming that they express the attitudes they do in freestanding contexts and do so in a way that will sustain a nonfactualist interpretation. In *Moral Fictionalism*, I argued that this could not be done—that an echo of the pragmatic fallacy besets even sophisticated forms of expressivism, of the kind developed by Blackburn and Gibbard, that do not take the form of an atomistic reduction. Whether or not that argument is ultimately successful, moral fictionalism at least has the advantage of not facing the familiar semantical difficulties that standard noncognitivism faces.

Or does it?

One might worry that this advantage is illusory—that the nature of moral argument is a challenge for moral fictionalism as well. Or so some have recently alleged. **CHANGE SLIDE**

The slide is titled "The Frege-Geach Problem" and is part of a presentation with navigation tabs at the top: "Three Claims", "Frege-Geach Problem", "An Old Problem in a New Guise?", "Meeting the Challenge", and "Summary". The slide content includes two portraits of philosophers: the top one is Gottlob Frege and the bottom one is Gilbert Harman. To the right of the portraits is the text "The Dilemma:" followed by two bullet points: "● Freestanding and embedded occurrences of a moral word have the same meaning" and "● Freestanding and embedded occurrences of a moral word differ in meaning". At the bottom right of the slide is a small logo for "VCL" and some navigation icons.

Figure 3: The Frege-Geach Problem

§ 3 AN OLD PROBLEM IN A NEW GUISE?

Matti Eklund is prominent among those who worry whether moral fictionalism really enjoys this advantage over standard noncognitivism. Eklund writes:

However, one may reasonably worry that the envisaged fictionalist move doesn't in fact help at all with the Frege-Geach problem. For someone actually presenting an argument like this would be presenting a good argument. But for this to be so, what the arguer actually expresses by the premise sentence must provide good reason to accept what she actually expresses by the conclusion sentence. Kalderon's moral fictionalist faces the problem of giving an account of what is actually expressed that respects this. But this seems not to be essentially different from the traditional non-cognitivist's problem of saying what the meanings of the relevant sentences are such that the impression that argument is valid can be respected. (Eklund, Summer 2007)

CHANGE SLIDE

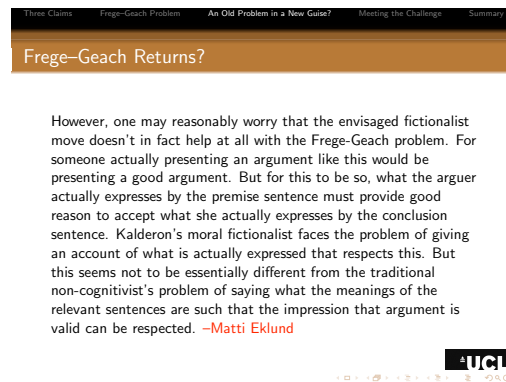


Figure 4: Frege-Geach Returns?

I believe that Eklund is misrepresenting a reasonable *challenge* for moral fictionalism as a determinate *problem*. To see this, let's detail the way this challenge differs from, and so is nonidentical to and disanalogous with, the Frege-Geach problem. **CHANGE SLIDE**

First, the Frege-Geach problem is a dilemma for expressivism, the second horn of which concerns, in part, *entailment*. Recall, atomistic reductions face

the following dilemma: 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 atomistic reduction 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 Eklund doesn’t raise a problem about entailment. Indeed, *there could be no problem about entailment*—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. Eklund’s challenge, instead, concerns the reasonableness of moral inference, as the fictionalist understands it. Eklund asks how a moral argument, as the fictionalist understands it, could 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, Eklund’s challenge concerns *inference*. **CHANGE SLIDE**

Second, the issue about entailment, as it arises in the context of the Frege-Geach problem, is motivated by a genuine, and potentially persisting, problem. Expressivism, in its atomistic form, lack an account of the meaning of moral words in embedded contexts and this raises a question about the validity of moral argument. The problem potentially persists since the original rationale for a nonfactualist interpretation is linked to the problematic, atomistic reduction. If there is nothing more to the content of a moral sentence than the noncognitive attitudes that its utterance conveys, and if these attitudes can be individuated independently of any moral object, quality, or relation, then

the content of a moral sentence could not consist in any moral proposition expressed. So, if expressivism abandons atomistic reduction, it must provide some other explanation for why it determines a nonfactualist interpretation for moral discourse. Expressivists must explain the meanings of moral words in embedded contexts and do so in a way that secures a nonfactualist interpretation for them. The problem, then, potentially persists since it is far from obvious how this could be achieved. Indeed, in *Moral Fictionalism*, I argued that it could not.

Notice, however, that Eklund’s challenge is not so motivated. There is no determinate problem for a natural, if primitive, form of fictionalism that potentially persists for sophisticated variants. What there is is nothing more than a reasonable query—to explain the attitudes conveyed by sincerely uttering a moral sentence, attitudes that constitute the acceptance of that sentence, in such a way that inferences involving that sentence count as reasonable.

I think this is a good question. But it is not the Frege–Geach problem. Nor is it even a problem. At least not yet. **CHANGE SLIDE**

Three Claims Frege–Geach Problem An Old Problem in a New Guise? Meeting the Challenge Summary

Differences

Why Eklund’s challenge is not the Frege–Geach problem:

- entailment versus inference
- dialectical difference

UCI

Figure 5: Differences

§4 MEETING THE CHALLENGE

The remainder of this talk will explore how Eklund’s challenge might be met given the resources of *Moral Fictionalism*.

According to noncognitivism, moral acceptance is not belief in a moral proposition. The difference between cognitivism and noncognitivism is partly a difference in the nature of the attitudes involved in moral acceptance

(whether they are cognitive or noncognitive) and partly a difference in the content of these attitudes (whether the cognitive attitudes involved, if any, have moral propositions as their objects).

Eklund's challenge is to specify the nature and content of the attitudes involved in moral acceptance such that the inferences represented by good moral arguments count as reasonable.

A variety of accounts of moral acceptance are consistent with the abstract possibility of moral fictionalism. In *Moral Fictionalism*, I argued that moral acceptance is a kind of affect akin to Scanlon's (1998, chapter 1) desire in the directed attention sense. A desire in the directed attention sense involves a tendency to focus on the object of desire as well as a tendency for the object of desire to appear in a favorable light.

The relevant affect is a generalization of this notion. First, it is unclear what a tendency to focus comprises. As I understood this notion, the relevant affect involved a tendency for certain features of the circumstance to become salient in perception, thought, and imagination. Second, there is a sense in which the label 'desire in the directed attention' sense is inapt—it suggests a too narrow construal of the relevant kind of affect. Specifically, it suggests that the constituent normative appearance is invariably positive. However, whatever Scanlon's intention, I am not assuming that the object of the affect invariably appears in a favorable light, only that there is a tendency for the object of the affect to have a certain normative appearance, whether or not that appearance is positive.

So far, the relevant kind of affect has been characterized in terms of its functional role—in terms of the tendency of certain features of the circumstance to become salient in perception, thought, and imagination and the tendency for these features to present a certain normative appearance. Scanlon commits himself to nothing further than this functional characterization.

It is natural to ask 'Why do these effects hang together?' What is it about the nature of this attitude that explains and renders intelligible that it should have this functional role? A question forcefully pressed by Johnston (2001). McDowell (1998, essays 3, 10) suggests that it is 'natural' for the noncognitivist to conceive of this attitude as a mixed state, a noncognitive refinement of sensing, where the noncognitive component is the source of the normative appearance. This is a substantive metaphysical proposal about the nature of the attitude that goes beyond anything that Scanlon commits himself to. McDowell has done much to discredit the claim that the relevant kind of affect can be understood as a mixed state involving perceptual and

noncognitive components that can be independently specified.

It is for this reason that in *Moral Fictionalism* I eschewed such an account. If the affect were conceived to be a particular event in a person's consciousness, then it would be reasonable to assume that its nature would be manifest in the way it structures a person's consciousness, and so reasonable to assume that the nature of the affect would explain and render intelligible the tendency of the object of the affect to become salient and the tendency for it to present a certain normative appearance. However, to assume at the outset that the affect is a particular conscious event is to overlook a metaphysical option available to the noncognitivist. Perhaps the affect is not some particular event in a person's consciousness, but the way in which events in the person's consciousness are structured. The suggestion is that there is nothing more to being an affect of the relevant kind than the tendency for certain features of the circumstance to become salient in perception, thought, and imagination, and the tendency for these features to present a certain normative appearance. This is itself a substantive metaphysical claim that goes beyond anything that Scanlon commits himself to (just as Blackburn's projectivism is a substantive metaphysical claim that goes beyond anything that Hume commits himself to).


This conception evades McDowell's criticism. The relevant affect is, to be sure, a mixed state—it is a noncognitive attitude that involves thoughts and perceptions about the morally salient features of the circumstances. However these attitudes are not distinct and so could not be independently specified; nor could they explain the way they structure a person's consciousness. The thoughts and perceptions involved in moral acceptance are events in a person's consciousness whose structure *constitutes* the relevant affect. **CHANGE SLIDE**

So far I have characterized a class of affects. What distinguishes the affects involved in moral acceptance from other attitudes of this class? What distinguishes them from being in a bad mood, for example? What is distinctive about the affects is the particular way that they structure a person's moral consciousness. The affects involved in moral acceptance structure a person's moral consciousness in a way that reflects the marks of moral authority.

Morality is authoritative. After all, while morality in some sense answers to our concerns, it is also in some sense independent of them. The authority of morality is manifest in the role it plays in moral discourse and in the cognitive psychology of competent speakers. A full account of that authority would involve specifying its source in a way that made it intelligible that it

Three Claims Frege–Gödel Problem An Old Problem in a New Guise? Meeting the Challenge Summary

Desire in the Directed Attention Sense



The affect involved in moral acceptance is nothing over and above:

- A tendency for the object of the affect to become salient in perception, thought, and imagination
- A tendency for the object of the affect to present a certain complex normative appearance

UCL

Figure 6: Desire in the Directed Attention Sense

should exhibit that role. The marks of moral authority are the distinguishable aspects of the role moral acceptance plays in moral discourse and the cognitive psychology of competent speakers.

- The first mark of moral authority is *precedence*, a variant of what Rawls (1971) and Scanlon (1998) call the ‘priority of right’. Specifically, in uttering a moral sentence that he understands, a competent speaker conveys a reason to act in a given circumstance that potentially overrides or cancels any conflicting nonmoral reasons available in that circumstance. **CHANGE SLIDE**
- The second mark of moral authority is *noncontingency*, a mark that distinguishes moral reasons from reasons of taste. Specifically, in uttering a moral sentence that he understands, the existence of the reason conveyed is not contingent upon the speaker’s or anyone else’s accepting it. **CHANGE SLIDE**
- The third mark of moral authority is *well-groundedness*. In uttering a sentence that he understands, the grounds the competent speaker has, if sincere, for accepting the uttered moral sentence applies not only to the speaker but to everyone else as well. So in uttering a moral sentence that he understands, a competent speaker accepts that sentence on behalf of others. Hume puts the point this way:

When a man denominates another his enemy, his rival, his an-

tagonist, his adversary, he is understood to speak the language of self-love, and to express sentiments, peculiar to himself, and arising from his particular circumstances and situation. But when he bestows on any man the epithets of vicious or odious or depraved, he then speaks another language, and expresses sentiments, in which, he expects, all his audience are to concur with him. He must here, therefore, depart from his private and particular situation, and must chuse a point of view, common to him and others. (Hume, 1740/2003, section 9)

CHANGE SLIDE

- The fourth mark of moral authority is the *demand* it makes on others, and approximately corresponds to the 'do so as well' component of Stevenson's (1944) analysis. In uttering a moral sentence that he understands, a competent speaker demands that his audience accept the uttered moral sentence.

Putting these marks of moral authority together with the account of affect derived from Scanlon, we get the following picture: In accepting a moral sentence that he understands, a competent speaker reconfigures his affective sensibility so as to render salient, in a phenomenologically vivid manner, the moral reasons apparently available in the circumstance, as he understands it. It is the structure of a person's moral consciousness, and not some further fact, that constitutes the relevant kind of affect. The relevant affect is nothing over and above the tendency for certain features of the circumstance to become salient in perception, thought, and imagination, and for these to present a certain complex normative appearance. Specifically, certain features of the circumstance become salient and appear to be reasons for acting, while other features potentially cease to be salient and can appear to be outweighed or even ruled out as reasons for doing otherwise, even if, in normal circumstances, they would count as such reasons. The salient features appear to be reasons that are not contingent upon our acceptance of them. There appears to be reasons for accepting the moral sentence, and these reasons directly or indirectly involve reasons that ground the deontic status of the relevant practical alternatives. These grounding reasons appear to be reasons not only for the speaker, but for everyone else as well. They appear to be sufficient reason for accepting that sentence on behalf of others. From this perspective, the competent speaker

can seem justified in demanding that others accept the moral sentence and so come to respond affectively in the relevant manner. The affects centrally involved in moral acceptance are in this way essentially other regarding. What distinguishes the affects involved in moral acceptance, then, is that they display a structure appropriate to the authority of our moral claims. **CHANGE SLIDE**



Figure 7: Marks of Moral Authority

Now that we have sketched the nature and content of the attitudes involved in moral acceptance, let's consider a 'good' moral argument. Indeed, 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.

Given well-groundedness—that the reason provided by a moral claim applies not only to the speaker but to everyone else as well—accepting the second premise of the argument has as a consequence that it is wrong not just for the speaker to lie but for the speaker's younger brother to lie. But that's not the conclusion of the argument. The argument is supposed to establish that *getting one's little brother to lie* is wrong. That seems right to me. The wrongness of getting one's little brother to lie is a further moral commitment, thus requiring the initial premise.

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. **CHANGE SLIDE**

Three Claims Frege–Geach Problem An Old Problem in a New Guise? Meeting the Challenge Summary

A 'Good' Moral Argument

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

UCL

Figure 8: A 'Good' Moral Argument

SUMMARY

We began with the distinction between *noncognitivism*, *expressivism*, and *nonfactualism*. Whereas the former is a *psychological* claim, the latter are positive and negative *semantic* claims. Again, if there is one thing to take away from *Moral Fictionalism*, it is the recognition that these are logically distinct claims. The logical distinctness of these doctrines raises the possibility of moral fictionalism, of noncognitivism without a nonfactualist expressivism. **CHANGE SLIDE**

One advantage that moral fictionalism has over standard noncognitivism is that it is not subject to semantical difficulties that arise in the wake of

the Frege-Geach problem. Specifically, a nonfactualist expressivism faces a potentially persistent difficulty in explaining moral *entailment*, a difficulty not similarly faced by moral fictionalism. **CHANGE SLIDE**

Nevertheless, as Eklund's worry shows, there is a question, not about moral entailment, but about the reasonableness of moral inference. How is that moral inferences that we ordinarily count as reasonable count as reasonable when understood as the fictionalist understands them? **CHANGE SLIDE**

Finally, I suggested that it is plausible, at least to a first approximation, that such inferences would in fact be reasonable, as the fictionalist understands them. In *Moral Fictionalism*, I argued that the attitudes involved in moral acceptance were a generalization of Scanlon's notion of desire in the directed attention sense. 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. Nothing that Eklund says suggests otherwise.

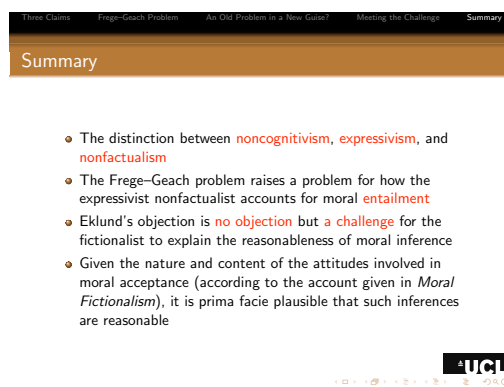


Figure 9: Summary

REFERENCES

Alfred Jules Ayer. *Language, Truth and Logic*. Gollancz, London, 1946. 5

John Dewey. Ethical subject-matter and language. *Journal of Philosophy*, 42: 701-12, 1945. 6

- Matti Eklund. Fictionalism. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Summer 2007. URL <http://plato.stanford.edu/archives/sum2007/entries/fictionalism/>. 7
- Peter Geach. Imperative and deontic logic. *Analysis*, 18:49–56, 1958. 4
- Peter Geach. Ascriptivism. *Philosophical Review*, 69:221–5, 1960. 4
- Peter Geach. Assertion. *Philosophical Review*, 74:449–65, 1965. 4
- Allan Gibbard. *Thinking How to Live*. Harvard University Press, Cambridge, Massachusetts, 2003. 2
- Paul Horwich. The Frege-Geach point. *Philosophical Issues*, 15:78–93, 2005. 2
- David Hume. *A Treatise of Human Nature*. Oxford University Press, Oxford, 1740/2003. 13
- Mark Johnston. The authority of affect. *Philosophy and Phenomenological Research*, 63(1):181–214, 2001. 10
- John McDowell. *Mind, Value, and Reality*. Oxford University Press, 1998. 10
- John Rawls. *Theory of Justice*. Belknap Press, Cambridge Massachusetts, 1971. 12
- David Ross. *Foundations of Ethics*. Oxford University Press, Oxford, 1939. 4
- Thomas Scanlon. *What We Owe to Each Other*. Belknap Press, Cambridge Massachusetts, 1998. 10, 12
- John Searle. Meaning and speech acts. *Philosophical Review*, 71:423–32, 1962. 4
- John Searle. *Speech Acts, An Essay in the Philosophy of Language*. Cambridge University Press, 1969. 4
- Charles Stevenson. *Ethics and Language*. Yale University Press, New Haven, 1944. 13
- Ralph Wedgwood. Conceptual role semantics for moral terms. *Philosophical Review*, 110:205–220, 2001. 2