

# 5

## Expressivism, Inferentialism and the Theory of Meaning

*Matthew Chrisman*

### 1. Introduction

One's account of the meaning of ethical sentences should fit – roughly, as part to whole – with one's account of the meaning of sentences in general. When we ask, though, where one widely discussed account of the meaning of ethical sentences fits with more general accounts of meaning, the answer is frustratingly unclear. The account I have in mind is the sort of metaethical expressivism inspired by Ayer, Stevenson, and Hare, and defended and worked out in more detail recently by Blackburn, Gibbard, and others. So, my first aim (Section 1) in this paper is to pose this question about expressivism's commitments in the theory of meaning and to characterize the answer I think is most natural, given the place expressivist accounts attempt to occupy within metaethics. This involves appeal to an ideationalist account of meaning. Unfortunately for the expressivist, however, this answer generates a problem; it's my second aim (Section 2) to articulate this problem. Then, my third aim (Section 3) is to argue that this problem doesn't extend to the sort of account of the meaning of ethical claims that I favor, which is like expressivism in rejecting a representationalist order of semantic explanation but unlike expressivism in basing an alternative order of semantic explanation on inferential role rather than expressive function.

### 2. Expressivism and the theory of meaning

Metaethics is often taught as beginning – in a way that has any clear distinction from normative ethics – with Moore's (1903, chapter 1) discussion of the 'naturalistic fallacy' and presentation of the 'open-question argument' against the reduction of moral terms like 'good' to non-moral terms like 'what's desired.' To be sure, almost no contemporary metaethicist thinks that the 'naturalistic fallacy' really is a fallacy or that the 'open-question argument' shows everything that Moore thought that it showed. However, it is widely assumed that one's metaethical view must take a stand on

Moore's contention that moral terms cannot be analyzed in natural terms. In response, I think we see four basic kinds of theories favored by most contemporary metaethicists:

- *Non-naturalists* agree with Moore's initial proposal, at least in its spirit even if not its details. That is, they argue that the only plausible way to accept Moore's negative argument and to continue to take ethics seriously is to posit the existence of *sui generis* moral properties, which can then be seen as the referents of moral terms.<sup>1</sup>
- *Naturalists* balk at the ontological commitments involved in non-naturalism. So, they argue that Moorean arguments for unanalyzability either can be met or are beside the point. If they think they can be met, they propose an analysis of moral terms in non-moral terms. If they think it's beside the point, they argue that, even if we cannot analyze moral terms in non-moral terms, that doesn't show that moral properties aren't natural properties any more than the fact that we cannot analyze the term 'water' with the coextensive term 'H<sub>2</sub>O'.<sup>2</sup>
- *Fictionalists* see an easier way to ontological parsimony. They argue that we can accept Moore's negative argument and agree with the non-naturalist about the ostensible referents of moral terms, but we can think of our language referring to these *sui generis* properties as a convenient fiction rather than as manifesting actual ontological commitments.<sup>3</sup>
- *Expressivists* propose something quite different. They argue that we can accept Moore's negative argument while avoiding commitment to *sui generis* moral properties by arguing that ethical claims have a distinctive expressive role in our discourse, which obviates any need for a theory of the referents of moral terms.<sup>4,5</sup>

One traditional way to think about these divisions is in terms of each view's commitments in the theory of reality. In this regard, non-naturalism and naturalism are usually seen as forms of moral realism because they are committed to the reality of moral facts alongside other sorts of facts (physical, biological, economic, etc.). By contrast, fictionalism and expressivism are usually seen as forms of moral anti-realism because they seek to avoid commitment to the reality of moral properties and correlative facts. There may be ways of talking about ethics that seem to commit one ontologically, but, if expressivists and fictionalists engage in these, they'll insist on an account of them that evades commitment in the final ontological reckoning.

Another way to think about these divisions is in terms of each view's commitments in the theory of meaning. Here, non-naturalism, naturalism, and fictionalism are allied in a 'descriptivist' (or 'factualist') order of explanation for the meaning of ethical claims. The idea is basically that ethical claims mean what they do, ultimately, in virtue of how they describe the world as being. (Of course, non-naturalists and naturalists hold that some of

these descriptions are literally correct, while fictionalists deny this, but the basic style of semantic explanation is the same.) By contrast, expressivists reject the descriptivism<sup>6</sup> implicit in the other theories, asserting instead that ethical claims mean what they do in virtue of their distinctive expressive role in our discourse.

This rejection of descriptivism purports to have two dialectical advantages. First, it lays the foundation for a novel account of the way ethical discourse seems distinctively connected to action. If ethical claims get their meaning in virtue of their distinctive expressive role, then, as long as what they express is connected to action in a special way, it's no surprise that ethical claims are distinctively connected to action. Second, it vindicates the expressivist's stance on the theory of reality without needing to treat ethical discourse as in any way fictional or second-rate. The idea is that, because ethical claims aren't in the business of describing the world, we have no theoretical need to account for the facts other theorists think they describe.

But what do expressivists mean by 'distinctive expressive role'? Or, more pertinently, what *could* they mean by this phrase, given that their view is intended as a view about the meaning of ethical claims that undercuts a core semantic assumption of the main competitors to expressivism, while fitting with a broader view of meaning in the philosophy of language?

Let's approach the second part of this question first by considering what seems to me to be a deep theoretical choice point in the philosophy of language. In a recent survey of general developments in the theory of meaning, Loar writes, 'Twentieth-century theory of meaning is divided into two: *truth* theories, and *use* theories'<sup>7</sup> (2006, p. 85). I take it his idea is that, while all philosophers of language will want to have an explanation of the relationships between meaning, truth-conditions, and rules of use, some theories of meaning afford explanatory priority to truth-conditions, while others afford explanatory priority to rules of use. In order to afford explanatory priority to truth-conditions, however, one must think that truth is more than merely a device of semantic ascent and generalization.<sup>8</sup> Thus, a 'truth theory' of meaning will assume a non-deflationary, indeed typically a realist, notion of truth.<sup>9</sup> And this, I think, makes it into what Fodor and Lepore refer to as the "'Old Testament" story, according to which the meaning of an expression supervenes on the expression's relation to things in the world' (1991, p. 329). The idea, at its core, is that meaning is a matter of word-world representational purport, which is why I will refer to it as a *representationalist* account of (what constitutes) meaning.

By contrast, use theories don't depend on any particular conception of truth, but they do need an account of how use and the correlated rules of use constitute meaning. And, typically, the idea is to start with the observation that meaning (linguistic meaning, anyway) is a conventional and thus rule-governed affair. Given rules of correct use, the use theorist claims

that some of these rules – the *semantic* rules – are constitutive of an expression’s meaning. So far this is purposefully vague, and I’ll only call it *anti-representationalism* at this stage, because I want to distinguish between two versions of the idea below. But, if we are so inclusive, then I think anti-representationalism about meaning is what Fodor and Lepore refer to as the ‘“New Testament” story, according to which the meaning of an expression supervenes on the expression’s role in a language’ (ibid.).

So, the choice I think we should press expressivists to make at this point is between representationalism and anti-representationalism as a general theory of meaning. For, if their account of the meaning of ethical claims is to fit with a broader theory of meaning in the philosophy of language, then they’ll have to find a place on one or the other side of this divide. We’ve already seen that the expressivist’s main competitors are allied in what I referred to as a descriptivist view of the meaning of ethical claims; and I think this fits pretty clearly with representationalism, as an application of a general theory to a specific case. The anti-expressivist idea (slightly recast) is that, just like non-ethical claims, ethical claims mean what they do in virtue of how they represent the world as being, that is, in virtue of their word-world representational purport.<sup>10</sup> And different anti-expressivist theories will differ based on what kinds of facts they think ethical claims purport to represent. Since expressivism begins with an attempt to undercut this idea, it may seem as if the expressivist must take the other side of the divide, that is, endorse an anti-representationalist account of meaning, in general.

Although I think all contemporary expressivists do endorse a form of anti-representationalism, it’s not initially as straightforward as I’ve just made it sound. For there is a way to be a representationalist in the theory of meaning while nonetheless endorsing what is widely thought to be a form of expressivism. This is to say that meaning can, in general, be explained representationally, but to go on to insist that, strictly speaking, ethical expressions have no meaning. In fact, this seems to be very similar to Ayer’s idea that the mixed ethical–non-ethical sentence ‘Your stealing that money was wrong’ has as its ‘factual meaning’ that you stole that money, and that the correlative ethical generalization ‘Stealing money is wrong,’ as he puts it, ‘has no factual meaning – that is, expresses no proposition which can be either true or false’ (1936/1946, p. 107). This view is consistent, of course, with thinking that ethical expressions have something broader than factual meaning: call it ‘linguistic significance.’ It’s just that this sort of expressivist will stress the fact that linguistic significance outstrips what we’re accounting for in a theory of meaning as word-world representational purport.

I think this version of expressivism is appropriately dubbed the ‘boo-hooray theory’ of the meaning of ethical expressions. For terms like ‘boo’ and ‘hooray’ arguably do not have meaning, when that is construed as word-world representational purport, though they clearly do have linguistic significance. So, if the expressivist thinks that ethical terms function

roughly like those terms, he can continue to endorse a representationalist account of meaning, while disagreeing with non-naturalists, naturalists, and fictionalists that this general theory of meaning should be applied to ethical expressions. We might thus view the expressivist's position in the theory of meaning as a version of semantic representationalism (about the non-ethical) plus semantic nihilism (about the ethical).

The problems with this way of developing the expressivist idea have become glaringly obvious over the past fifty years of metaethical debate. In the following two paragraphs, I sketch what I take to be the core of the problem.

Because of the productivity and learnability of language, it's highly plausible that meaning is compositionally systematic, in the sense that, if two sentences have the same deep grammatical form, then the meaning of each sentence can be represented as a single function on the meaning of the correlative parts. We can make the point at the level of subsentential parts, such as predicates, which compose with things like terms, quantifiers, and other predicates to form whole sentences. But the same point applies at the level of logically simpler sentences, which compose with sentential connectives and operators to form logically more complex sentences. So, for example, the meaning of a sentence such as 'Grass is green' can be represented as a function SUBJPRED from the meanings of 'Grass' and 'is green.' Likewise, the meaning of a logically complex sentence 'If grass is green, then chlorophyll is green' can be represented as a function COND from the meaning of 'Grass is green' and 'Chlorophyll is green.'

The problem this generates is that, by the normal standards of syntacticians, each ethical claim seems to have the same deep grammatical form as some non-ethical claim. If that's right, it means that we should be able to represent the meaning of a sentence such as 'Tormenting is wrong' as the (exact same) function SUBJPRED from the meanings of 'Tormenting' and 'is wrong' as before. Likewise, we should be able to represent the meaning of a logically complex sentence such as 'If tormenting is wrong, then torturing is wrong' as the (exact same) function COND from the meanings of 'Tormenting is wrong' and 'Torturing is wrong'. However, the boo-hooray version of expressivism is incompatible with this idea, since it denies that predicates such as 'is wrong' and sentences such as 'Tormenting is wrong' have meaning. Because of this, the boo-hooray expressivist is committed to defending a highly implausible presupposition of his theory, that is, that the deep grammatical form of ethical sentences such as 'Tormenting is wrong' and 'If tormenting is wrong, then torturing is wrong' is radically different from the deep grammatical form of non-ethical sentences such as 'Grass is green' and 'If grass is green, then chlorophyll is green.'<sup>11</sup>

This is one way of viewing Geach's (1965) and Searle's (1962) point in arguing that (early) expressivists cannot make sense of the semantic similarity between a logically unembedded occurrence of an ethical sentence, such

as ‘Tormenting the cat is wrong,’ and a logically embedded occurrence, such as ‘If tormenting the cat is wrong, then getting your little brother to torment the cat is wrong.’ The problem is not – as it has sometimes been supposed – that expressivists have nothing they can say about why endorsing these two sentences licenses endorsing the sentence ‘Getting your little brother to torment is wrong.’ Of course, they are theoretically prevented from saying that the meaning of the unembedded sentence and the antecedent of the conditional are both a function of their representational purport, but they can tell some other story linking the linguistic significance of the unembedded sentence to the linguistic significance of the conditionalized sentence. The problem, rather, is that a commitment to the compositional systematicity of meaning puts a very heavy explanatory burden on anyone who would explain the semantic relationship between these three sentences in any way different from the explanation of the semantic relationship between the non-ethical sentences ‘Tormenting the cat is loud’ and ‘If tormenting the cat is loud, then getting your little brother to torment the cat is loud’ and ‘Getting your little brother to torment the cat is loud.’ In short, the meaning of conditionalized sentences – whether ethical or non-ethical – needs to be represented as a *uniform* function of the meanings of their antecedents and consequents. That is encouraged by viewing meaning as compositionally systematic across deep grammatical similarity.<sup>12</sup>

Although I think this problem undermines expressivism, understood as committed to semantic representationalism plus nihilism, it’s not clear whether it undermines expressivism *tout court*. For, given the broader distinction between representationalist and anti-representationalist accounts of meaning in the philosophy of language, it remains open for expressivists to reject representationalism, quite generally, and go in for some form of anti-representationalism. Recall that these theories are ones that start their explanation of meaning not from a notion of representational purport but from a notion of rule-governed linguistic role. So, the expressivist who wants to avoid the problems with the boo-hooray theory will need to appeal to an account of how rule-governed linguistic role constitutes meaning, which comports with the compositional systematicity of meaning.

Here, the strategy which motivates the name ‘expressivism’ is, I believe, to cash out the notion of linguistic role in terms of *expressing a mental state*. For example, Gibbard writes, ‘[W]hat “rational” means is explained by saying what it is to call something “rational”, and to call something “rational”, the analysis says, is to express a state of mind’ (1990, p. 84). This approach to linguistic meaning is not a new one created for the purposes of salvaging metaethical expressivism. It draws inspiration from a general theory of meaning tracing back to Locke, who wrote ‘Words in their primary or immediate Signification, stand for nothing, but the Ideas in the Mind of him that uses them’ (1690/2008, chapter II). Grice (1957), Schiffer (1972) and Davis

(2003) have made this idea much more plausible by distinguishing between speaker-meaning and sentence-meaning, and then arguing that a sentence's meaning could be explained in terms of mentalistic notions such as belief and intention, linguistic regularities, and conventions created and eventually ossified by the speaker-meanings in a community of language users.

The general idea is to start with the observation that people outwardly express what's going on internally in all sorts of ways (e.g. wincing, smiling, winking, giving a thumbs-up, etc.), and this can be either intentional or unintentional. One more way people express what's going on internally is by using language with the intention of indicating what's going on internally. However, for this to work, there must be tractable regularities between the use of certain sign/sound-designs and certain mental states. Thus, over time, conventional rules develop, which govern the use of these sign/sound-designs. And in light of these conventional rules we can speak not just of the mental state token expressed by some person in doing something but also of the mental state type expressed by a particular sign/sound-design. If we allow the term 'idea' to refer to mentally instantiated concepts which can be expressed in language and the term 'thought' to refer to combinations of such concepts into judgments (the mental analogues of making a claim), the core thought here is to understand meaning as constituted by what we might call word-idea/thought (conventional) expressive function. And, as a version of anti-representationalism, this general program in theory of meaning is sometimes called *ideationalism*.

How does ideationalism fit with a commitment to the compositional systematicity of meaning? It was violating this plausible commitment in the theory of meaning that undermined the boo-hooray version of expressivism. But it looks as though ideationalists can do better, since, like the representationalists, they have a generic formula for specifying the meaning of any predicate and sentence. So, for example, while the representationalist can say that the meaning of the predicate 'is green' is its purporting to represent some property, in this case greenness, the ideationalist can say that this predicate's meaning is what concept conventionally expresses, that is, the concept of greenness. And, while the representationalist says that the meaning of the sentence 'Grass is green' is its purporting – via the systematic contribution (represented above by the function SUBJPRED) of the representational purport of its parts – to represent the fact that grass is green, the ideationalist can say that the meaning of the whole sentence is its conventionally expressing – via the systematic contribution (also representable as the function SUBJPRED) of the expressive function of its parts – the thought that grass is green. This is similarly the case with logically complex sentences such as 'If grass is green, then chlorophyll is green.' Here the representationalist says that the meaning is the sentence's purporting – via the systematic contribution (representable as the function COND) of the representational purport of its parts – to represent the fact that, if grass is green, then chlorophyll is green. And the

ideationalist has a parallel story: this sentence's meaning is its conventionally expressing – via the systematic contribution (also representable as the function COND) of the expressive function of its parts – the thought that, if grass is green, then chlorophyll is green. And so on.<sup>13</sup>

So the general strategy available to the ideationalist for capturing the systematicity of the meaning of whole claims through an isomorphism with the structure of thoughts is much like the representationalist's strategy of appealing to an isomorphism with the structure of what is represented. Given that compositional systematicity is precisely the stumbling block for any form of expressivism that signs up to a representationalist plus nihilist account of the theory of meaning, ideationalism would seem to be a very good general theory of meaning for expressivists to opt for instead. If they do so, it seems that they can provide a uniform account of the meaning of ethical and non-ethical claims by saying that all claims mean what they do by virtue of the conventional word-idea expressive function of their parts and the systematic contribution of these parts to the conventional sentence-thought expressive function of whole sentences.

### 3. A problem

So far, I've argued that the most natural general account of meaning for an expressivist to endorse is an ideationalist account, which explains linguistic meaning in terms of word-idea expressive function. However, those who favor representationalism in the general theory of meaning are unlikely to see the ideationalist's switch from word-world representational purport to word-idea expressive function as making much progress in accounting for meaning. For, even if it's true that sentences mean what they do in virtue of the mental states they conventionally express, the representationalist will insist that this is true only because mental states have the content that they have in virtue of how they represent the world as being. That is, from the representationalist's point of view, ideationalism will not look like a free-standing account of meaning but only a 'dogleg' through ideas/thoughts (mentally construed) on the way to representations of the world. And, in so far as this is correct, the representationalist will insist that it's the notion of representational purport, rather than expressive function, that is doing the fundamental explanatory work. And, moreover, whatever explanation the ideationalist has of the systematic compositionality of meaning in terms of the compositional structure of the content of thoughts it's derivative of the representationalist's explanation in terms of the compositional structure of the facts represented.

An ideationalist might argue in response that words are mere scribbles or sounds until they are *imbued* with their representational purport. Then, the ideationalist could argue that what imbues some scribbles and sounds with representational purport is, quite generally, their incorporation in a human



practice via the generation of expressive conventions linking scribbles and sounds to human ideas (i.e. parts of human thoughts). The idea is that scribbles and sounds that represent parts of the world do so only because they are linked via conventional rules to ideas and thoughts, which themselves represent parts of the world. More abstractly, the claim is that the notion of word-world representational purport could not even be attached to words without first appealing to the notion of word-idea expressive function. Hence, although the former notion can attach to words, the latter notion is more explanatorily fundamental. This, I take it, is what many ideationalists in the philosophy of language would say in response to the representationalist's criticism that their theory is a mere dogleg. And, moreover, it preserves the primacy of the structure of the content of thoughts in explaining the systematic compositionality of meaning.

However, it's unclear whether this line of response can work for the expressivist who wants to give an ideationalist explanation of the systematic compositionality of meaning while underwriting the view that ethical claims and non-ethical claims play a *distinctive* expressive role. More specifically, it seems that an expressivist who endorses ideationalism as a general theory of meaning will nonetheless have to hold that, while ethical and non-ethical claims both mean what they do in virtue of the thoughts they express, ethical thoughts do not have their contents in virtue of the representational purport attached to them via human expressive practices. Otherwise, the expressivist hasn't really denied the representationalism implicit in non-naturalism, naturalism, and fictionalism. To appreciate this, notice how these anti-expressivists can, as far as their debate with the expressivist goes, accept that ethical claims mean what they do in virtue of their distinctive expressive role. They'll say that ethical claims express ethical thoughts and non-ethical claims express non-ethical thoughts. However, that's not yet a metaethically interesting contrast. For, by the same token, physical claims express physical thoughts, biological claims express biological thoughts, economic claims express economic thoughts, etc.

Thus, expressivists need an expressive contrast drawn not in terms of the differing contents of the mental states expressed but in terms of the different kinds of mental states expressed. More specifically, it seems that expressivists need to cash out their core thesis that ethical claims mean what they mean in virtue of playing a distinctive expressive role by arguing that non-ethical claims express *beliefs*, while ethical claims express something like desires, intentions, or plans. This is a familiar distinction from Humean psychology of motivation, which is often characterized in terms of a difference in direction of fit.<sup>14</sup> Beliefs function like *maps* in that they are meant to fit the way the world is, while desires, intentions, and plans function more like *orders* in that they're meant to get the world to fit them. However, what this makes apparent, in light of our discussion of general theories of meaning, is that the expressivist cannot endorse a standard form of ideationalism that

accepts a strict isomorphism between the semantic content of sentences and the mental content of the ideas they conventionally express. For, although he's happy to say that a non-ethical sentence, such as 'Tormenting is loud,' means what it does in virtue of expressing the *belief* that tormenting is loud, he must deny the structurally isomorphic explanation of the meaning of an ethical claim, such as 'Tormenting is wrong.' This, on his view, expresses a desire, intention, or plan.<sup>15</sup>

However, it should now be clear that this spells trouble for the sort of expressivist I've been discussing in this section. For the main advantage of adopting ideationalism in the theory of meaning rather than the representationalism plus nihilism adopted by early expressivists was that ideationalism can explain the semantic content of whole sentences as a systematic function of the semantic content of their parts and the ways that they combine. However, if the expressivist now says that ethical and non-ethical claims both mean what they do in virtue of expressing thoughts, but the relevant way in which these thoughts differ is not directly in their contents but in their nature as beliefs or desires/intentions/plans, this undermines the ideationalist's ability to explain the meaning of whole sentences as a *systematic* function of the meaning of their parts and the deep grammar of ways they are combined.

That's all a bit abstract. Perhaps an example will help to make the point more concrete. Recall the logically complex sentences:

'If grass is green, then chlorophyll is green'  
'If tormenting is wrong, then torturing is wrong.'

If we accept the expressivist's idea that non-ethical thoughts are beliefs, which aim to fit the world, then the ideationalist account of the meaning of the first sentence amounts to the following. The logically simple sentence 'Grass is green' means what it does in virtue of expressing the belief that grass is green, whose content can be viewed as a systematic function (SUBJPRED) of the combination of the ideas of grass and greenness into a thought, and similarly with the logically simple sentence 'Chlorophyll is green.' Then, the conditionalized combination of these two sentences means what it does in virtue of expressing the conditionalized belief whose content can be viewed as a systematic function (COND) of the thoughts expressed by its parts, that is, the belief that, if grass is green, then chlorophyll is green. That's compositional systematicity of meaning at its finest.

What does the expressivist say about the meaning of the second conditionalized sentence? Well, compositional systematicity encourages us to break it into its logically simple parts as before: 'Tormenting is wrong' and 'Torturing is wrong.' However, according to the expressivist these are not the conventional expression of ethical beliefs (as the anti-expressivists say); rather, they're the expression of desires, intentions, or plans.<sup>16</sup> The

problem is that, although desires, intentions, and plans may be thought to have contents, it's clear that their contents are not isomorphic to the sentences that express them. If 'Tormenting is wrong' expresses a desire, intention, or plan, then surely it's not the desire that tormenting is wrong, or the intention that tormenting is wrong, or the plan that tormenting is wrong. Rather, it's something like the desire that people not torment, or the intention to stop people from tormenting, or the plan not to torment.

That already means that the expressivist cannot give a standard ideationalist explanation of the meaning of logically simple sentences as a systematic function of their parts determined by their deep grammar. The needed isomorphism between the contents of thoughts and the contents of sentences must break down in the ethical case, for the expressivist. The situation just gets worse when we consider logically complex ethical sentences. What thought does the second conditionalized sentence blocked out above express, according to the expressivist? Even if we agree that it does express a desire, intention, or plan, it's clear that the content of this desire, intention, or plan is not isomorphic to the content of the conditionalized sentence. It's not obviously even intelligible to desire, intend, or plan that, if tormenting is wrong, then torturing is wrong.

One way out of this problem is suggested by Blackburn's (1998, pp. 71–4; 2002) talk of 'being tied to a tree.' His idea is that logically complex claims should be seen, quite generally, as committing one to a predictable pattern of beliefs and attitudes. This would be consistent with the ideationalist idea that all sentences mean what they do in virtue of thoughts they express only if it is interpreted as the view that all logically complex claims – both ethical and non-ethical – express a complex desire-like disposition, which, depending on their combination with different claims that one endorses, leads one to have further thoughts and endorse further claims that express them. For example, on this kind of view, a claim such as 'If Grass is green, then chlorophyll is green' would be said to express not the belief that if grass is green then chlorophyll is green but something like the desire, intention, or plan that (i) if one thinks that grass is green then one thinks that chlorophyll is green and (ii) if one thinks that chlorophyll is not green, then one does not think that grass is green.

The problem I see with this way of respecting the systematicity of meaning across similar logical contexts is that it misplaces the distinction between claims that describe the world and claims that don't describe but express desires, intentions, or plans. For, on it, *any* logically complex claim comes out as non-descriptive. For example, 'Grass is green' would be said to describe the world in virtue of expressing a belief, but 'It is not the case that grass is not green' would be said to express a desire, intention, or plan that is characterizable as 'being tied to a tree.' I don't know if that is how Blackburn intends his view to be interpreted (he speaks of 'commitments' that are specifiable in terms of what other mental states it 'makes sense' to hold in combination). However, I think this is the most natural way to

interpret the idea within an ideationalist account of meaning. As we'll see below, there is an inferentialist tradition in the theory of meaning that may subvert the distinction between descriptive and non-descriptive sentences, and another way to interpret Blackburn's idea is in inferentialist terms,<sup>17</sup> but within an ideationalist version of expressivism I think it would be quite awkward if the only sentences thought to express beliefs and so to describe the world were atomic non-ethical sentences such as 'Grass is green.'

For these reasons, it seems to me that the only way an expressivist, who is committed to ideationalism as a general theory of meaning, can avoid this problem is to return to the core ideationalist idea that all claims express the same kind of mental state, whose differing contents explain differences in meaning. However, that is inconsistent with the expressivist's aspirations in metaethics to draw an ontologically and psychologically interesting contrast between ethical and non-ethical discourse in a plausible way.<sup>18</sup>

#### 4. A different way

In Section 1 I suggested that two different forms of expressivism emerge when we query the view's underlying assumptions in the theory of meaning. Early versions of expressivism seem to be committed to semantic representationalism (about the non-ethical) plus semantic nihilism (about the ethical), whereas later versions of expressivism seem to be committed to a form of anti-representationalism called ideationalism. I've pointed to the problems almost everyone agrees undermine early versions of expressivism, and I've rehearsed a related problem for the later versions of expressivism. In this section I want to argue that this is not enough to strike anti-representationalist accounts of the meaning of ethical claims off the list of theoretical options in metaethics. And I want to do so by again focusing on foundational issues in the theory of meaning that I believe should underpin the metaethical debate about the meaning of ethical claims.

Ideationalism is not the only anti-representationalist theory of meaning one finds in the philosophy of language. Indeed, in their presentation of 'Old Testament' theories of meaning, which are founded on word-world relations, it's clear that Fodor and Lepore mean to contrast these not with ideationalist theories but with *inferentialist* or *conceptual-role* theories of meaning, which are founded instead on the inferential/conceptual connections between words.<sup>19</sup> These 'New Testament' theories are anti-representationalist in that their most basic explanation of what constitutes the meaning of words and sentences doesn't rest on word-world representation relations. But they're not ideationalist forms of anti-representationalism, since they also don't take word-idea/thought expressive function as the most basic. Rather they focus on the inferential (and perhaps more broadly conceptual) connections between words and (more importantly) between the sentences they can be used to compose.

I think such theories of meaning are very promising, but I won't argue in general for them here, as my purpose is not to defend a general theory of meaning but to point out its near absence<sup>20</sup> in the metaethical debate about the meaning of ethical claims and to argue that it has some advantages over expressivism. To this end, however, it is necessary to sketch enough of an inferentialist theory of meaning to contrast it with representationalist and ideationalist theories sketched above.

Recall that representationalists hold that a declarative sentence means what it does in virtue of what fact it can be used to represent, and ideationalists hold that a declarative sentence means what it does in virtue of what thought it can be used to express. Thus, for example, the sentence 'Grass is green' would be said, by the representationalist, to mean what it does in virtue of representing the fact that grass is green, and, by the ideationalist, to mean what it does in virtue of conventionally expressing the thought that grass is green. The inferentialist position, in contrast, can be put in terms of the notion of an inferentially articulable commitment. The idea is that a sentence means what it does in virtue of what commitment it can be used to undertake.

Such commitments are not conceived ontologically, that is, as commitments to the existence of some fact in the final ontological reckoning. This is what makes the view a form of anti-representationalism. But neither are they conceived psychologically, as some part of an agent's mind. This is what makes the view a form of anti-representationalism that is different from ideationalism. There may be connections to commitments conceived ontologically or psychologically, but, in the first instance, the commitments I have in mind are conceived in terms of the inferential entitlements and obligations that undertaking them carries.<sup>21</sup>

The idea is that one can gain entitlement to a commitment (e.g. to grass's being green) by undertaking another commitment (e.g. to grass's being full of chlorophyll and chlorophyll's being green); and one can be obligated by one commitment (e.g. to grass's being green) to undertake another commitment (e.g. to grass's being colored). The inferentialist contends that the meaning of a claim is constituted by its place in a network of these sorts of connections. Thus, in contrast to representationalism and ideationalism, inferentialism says that the claim 'Grass is green' means what it does in virtue of the fact that it can be used to undertake a commitment to grass's being green; and this commitment is constituted by the network of inferential entitlements and obligations that commitment carries.

It's important to appreciate that this style of semantic explanation offers a completely generic formula for explaining the meaning of a claim, which extends to any degree of logical complexity. So, for instance, the inferentialist will also say that the conditionalized sentence 'If grass is green, then chlorophyll is green' means what it does in virtue of the inferentially articulable commitment it can be used to undertake. In this case, it's a commitment

to chlorophyll's being green in case grass is green. But if the sentence were instead 'If grass is red, then chlorophyll is purple,' we'd get the structurally isomorphic result: it means what it does in virtue of being usable to undertake a commitment to chlorophyll's being purple in case grass is red. The general theory, then, is that a claim means what it does in virtue of the inferentially articulable commitment it can be used to undertake.

I want to flag and try to put aside two questions about this general theory of meaning before exploring the different path I think it offers in the meta-ethical debates about the meaning of ethical claims.

First, a representationalist might wonder: isn't the commitment to something's being a certain way best understood in representationalist terms? This question points to a vexed issue about the notions of being, fact, and representation that I won't delve into deeply here. But I do want to recognize that one way to understand the notion of a commitment to something being a certain way (e.g. a commitment to grass being green) is in terms of the notion of representing the world as containing the fact that this thing is that way (e.g. the fact that grass is green). However, either this is a trivial and ontologically neutral reformulation, or it gets one very deep into questionable ontological commitments. For instance, most philosophers are committed to its being impossible that  $1 + 1 = 3$  and to its being unlikely that the sun will rise an hour later tomorrow. However, in so far as ontology goes, it's up for debate whether these commitments entail commitment to the world containing the fact that  $1 + 1 = 3$  is impossible and the fact that it is unlikely that the sun will rise an hour later tomorrow. For commitment to something being impossible might plausibly be thought to be about not merely what's a fact in this world but what could and couldn't be a fact, as we sometimes say, in other unreal but possible worlds. And commitment to something being unlikely might plausibly be thought to be not about what will in fact happen but about the strength of someone's evidence for thinking that something will happen. There are, of course, realist views of impossibility and improbability, but they raise naturalist qualms even more than realist views about what's ethically right and wrong. So, I think there is considerable theoretical room to resist the idea that commitment to something being a certain way is best thought of as an ontological commitment in implicitly representationalist terms.

Second, someone impressed with truth-conditional theories of the compositional systematicity of meaning might ask: what does the inferentialist say about the meaning of subsentential parts of sentences? In the case of representationalism and ideationalism, a 'bottom-up' explanation of meaning seems to be in the offing. For example, the representationalist could say that a predicate means what it does in virtue of the property, the having of which would make the predicate true of something. And the ideationalist could say that a predicate means what it does in virtue of the concept, the falling under which would make the predicate true of something. Then facts can be viewed as built (at least partially) out of properties, while thoughts can be viewed as built (at least partially) out of concepts. Again, there are

vexed issues here that I will not delve into deeply. However, I think it's important to appreciate that, although the inferentialist's explanation of meaning starts at the level of full claims, which are the proper relata of inferential/conceptual relations, this doesn't preclude a 'top-down' explanation of the meaning of subsentential parts of sentences. For example, the account could represent the meaning of predicates as sets of inferentially articulated commitments. The meaning of 'is green', for instance, could be seen as the set of inferentially articulable commitments one gets by substituting into 'x is green.' This is by no means the end of the story, but it should be clear enough that, by generalizing across claims with equivalent subsentential parts, we can represent the meaning of these subsentential parts by a process of top-down abstraction.<sup>22</sup>

Having flagged and put aside those two difficult questions, I now want to consider how the inferentialist theory of meaning applies to the meta-ethically interesting case of ethical claims. The discussion above suggests that an inferentialist will say that a sentence such as 'Tormenting is wrong' means what it does in virtue of the fact that it can be used to undertake an inferentially articulable commitment to tormenting being wrong. What it means for this commitment to be inferentially articulable is the same as before. It carries with it a network of entitlements and obligations to other commitments. So, it's ultimately the network of these inferential relations that constitutes the claim's meaning.

This style of explanation extends to logically complex ethical sentences in the expected way. So, for instance, the claim 'If tormenting is wrong, then torturing is wrong' means what it does in virtue of the inferentially articulable commitment it can be used to undertake. In this case, the commitment will be a commitment structurally isomorphic to the commitments undertaken with the conditional claims above: 'If grass is green, then chlorophyll is green' and 'If grass is red, then chlorophyll is purple.' That is, it's the commitment to torture being wrong in case tormenting is wrong.

So far, it seems that representationalism, ideationalism, and inferentialism – considered as general theories of meaning – have importantly different but parallel formulas for explaining the meaning of claims in a way that respects the compositional systematicity of meaning. They're explained in terms of the fact represented, the thought expressed, or the commitment undertaken. And, as long as there is a systematic function between the deep grammar of the claim and the structure of the fact, thought, or commitment, the explanation can respect the compositional systematicity of meaning on which early expressivists floundered.

Above, however, I argued that more recent expressivists, who assume an ideationalist order of semantic explanation, flounder in a more subtle way on the compositional systematicity of meaning. This is because their meta-ethical convictions force them to violate the systematic connection between the deep grammar of some claims and the structure of the thought they

express. In order to fund the purported difference in expressive role between ethical and non-ethical claims, contemporary expressivists have to tell a different story about ethical claims and their logical parts from the standard story about non-ethical claims and their logical parts. For, in order to challenge the semantic assumption underlying non-naturalism, naturalism, and fictionalism, they have to defend an anti-representationalist account of the meaning of ethical claims. But, in order to do this, they draw a distinction between thoughts that are like maps in representing the world and thoughts that are like commands in guiding one through the world, which in turn leads to the idea that the thoughts expressed by ethical claims are not beliefs but, rather, something like desires, intentions, or plans. But it turned out to be implausible to think that the content of desires, intentions, or plans bears the same systematic structural relation to the deep grammar of the sentences that express them as the content of beliefs. Would similar metaethical convictions force an inferentialist to make a similarly problematic move?

I don't think so. If we conceive of the commitments that constitute a claim's meaning not in ontological or psychological terms but in inferential terms, then we already have a general theory of meaning that challenges the semantic assumption underlying non-naturalism, naturalism, and fictionalism. It seems that this cannot be undermined by an analogue of the dogleg argument against ideationalism sketched above. For inferentially articulable commitments do not have the content they do in virtue of the facts they represent the world as having. Rather, it's precisely the point of calling them 'inferentially articulable' that their content is articulated in terms of inferential/conceptual connections to other claims and not in terms of their representational connection to facts putatively in the world.

So far as that goes, however, inferentialism may seem to undermine the contrast between ethical and non-ethical claims that was supposed to carry two dialectical advantages for anti-representationalist metaethical views such as expressivism over representationalist views such as non-naturalism, naturalism, and fictionalism. In my view, that's only partially right, which we can appreciate by reconsidering these putative dialectical advantages from the point of view of inferentialism.

The first putative dialectical advantage mentioned above was that the expressivist's way of being an anti-representationalist provides the basis for a novel and compelling explanation of the distinctive connection between ethical discourse and action. So stated, this is ambiguous between the idea that ethical claims seem to be distinctively connected to their authors' being motivated to act in certain ways, and the idea that ethical claims seem to be distinctively connected to the justification or legitimization of certain actions from their authors' point of view. I think that both of these ideas can, when properly understood, be better captured by the inferentialist than by the expressivist.



Early expressivists were sometimes animated by the thought that one who sincerely makes an ethical claim will be at least somewhat motivated to act in its accord. However, this form of judgment-internalism is clearly too strong. Often ethical claims are not about anything that their author expects to have a bearing on how he or she acts (e.g. claims about the rights and wrongs of international policy or about what some historical figure did), or they are about *pro tanto* considerations that their author considers outweighed (e.g. about what would most please one's mother), or they are abstract enough to leave one cold (e.g. that one ought to act in such a way that one could will one's maxim as universal law), or they simply do leave one cold because of akrasia or some other disconnect between conviction and motivation. So, whatever form of judgment-internalism is true, it must be weak enough to allow motivations and sincere ethical claim-making to diverge. But, if that is the case, then I think the inferentialist can say two plausible things about judgment-internalism. First, perhaps sincere ethical claim-making is indeed better correlated with having the sorts of desires, intentions, and plans that explain motivation, but the exceptions to this psychological correlation indicate that it shouldn't dictate our explanation of the meaning of the relevant claims. We can recognize that certain sorts of ethical claims conventionally express desires, intentions, and/or plans without holding that they mean what they do in virtue of expressing these things. Second, if ethical commitments are inferentially connected to the core concepts deployed in practical reasoning in a way that doesn't hold for all other commitments – as I will suggest below – then it wouldn't be surprising if agents generally, though defeasibly, acquired motivation to act in ways justified by (at least some of) their ethical commitments. Perhaps the reason we expect one who claims that giving to charity is morally required to be motivated to give to charity is that this claim expresses a commitment that inferentially supports the practical commitment to give to charity; and we generally expect people to be defeasibly motivated to do what they're committed to in virtue of the commitments undertaken by their claims.

The other way to disambiguate the idea of a distinctive connection between ethical discourse and action is about justification or legitimization of action rather than directly about motivation. I think this idea is even more easily captured by the inferentialist than by the expressivist. To see this, notice that the expressivist has to say that it's something about the desire, intention, or plan expressed by ethical claims that explains why related actions are justified or legitimate, from the agent's point of view. Whether one finds that compelling will depend on whether one thinks desires, intentions, and plans can play this justifying or legitimating role. Perhaps they can, but, depending on how passively acquired such states are, one might also reasonably worry that they can't play that justificatory role until the agent endorses them. In contrast, the inferentialist can build the notion of a justifying or legitimating connection into his account of the

meaning of ethical claims. The details can be worked out differently for different ethical claims depending on how tight one thinks the connection is. But, as an example, consider a claim of the following form: 'I ought, all things considered, to  $\phi$ .' It would not be implausible to think that part of the inferential articulation of the commitment undertaken by this claim construes it as obligating its author to a further directly practical claim that we might express as 'I shall  $\phi$  when the appropriate time comes.' And we could then capture the apparent contrast with non-ethical claims by pointing out that they do not carry similar inferential connections to the core 'I shall' type claims involved in practical reasoning.

So, I think the inferentialist can actually do better than the expressivist in capturing the apparently distinctive practicality of ethical discourse in comparison to non-ethical discourse. If right, this means that the first dialectical advantage claimed for expressivism over non-naturalism, naturalism, and fictionalism actually counts more in favor of inferentialism than expressivism. But what about the second dialectical advantage? This had to do with the expressivist's stance in the theory of reality towards ostensible ethical properties and facts. By adopting expressivism, one is able to maintain a kind of naturalist-inspired anti-realism about the ethical, which obviates the need for certain controversial sorts of explanations crucial to non-naturalism, naturalism, and fictionalism. Can an inferentialist claim the same advantage?

In one sense, the answer is clearly 'yes,' since inferentialism – at least as far as it's been sketched here – is a thoroughgoing anti-representationalist theory. Unlike the ideationalist, who says that claims mean what they do in virtue of the thoughts they express but then goes on to grant that some thoughts (i.e. beliefs) have the content they have in virtue of how they represent the world as being, the inferentialist doesn't have to appeal to the notion of representing the world as being a certain way even as a secondary explanatory notion in his theory. This means that there is a way to be an inferentialist that generalizes the expressivist's naturalist-inspired anti-realist stance towards the ethical into a general anti-realist stance across the board.

However, I suspect that many metaethicists would view that way of validating a naturalism-inspired anti-realist stance towards the ethical as subverting the relevant debate in the theory of reality rather than capturing the anti-realist motivation for expressivism. That is to say that, if the reason ethical claims don't commit us ontologically to ethical facts is that no claims ever commit us ontologically, then the issue of ontological commitment and naturalism-inspired anti-realism is spurious.

Even if it is spurious, that still leaves inferentialism looking better than expressivism with respect to issues about the compositional systematicity of meaning and the apparently distinctive practicality of ethical discourse. And it would appear better than naturalism, non-naturalism, and fictionalism, in so far as they make supposedly spurious ontological claims. However,

I think a more subtle form of inferentialism wouldn't spurn the ontological debate, but would rather seek to reconstruct it in directly inferentialist terms. In characterizing inferentially articulable commitments above, I said that they are not conceived *in the first instance* as ontological commitments. That is to say, they are not constituted by which facts they commit one admitting into one's ontology. On this way of thinking of things, one can be committed to the impossibility that  $1 + 1 = 3$ , the unlikelihood that the sun will rise an hour late tomorrow, and the wrongness of torture, without being committed *ontologically* to the existence of some piece of reality that corresponds to these commitments. However, that doesn't rule out the possibility that something else would commit one ontologically, and that something else could make a difference between the types of commitments just canvassed and commitments to things like grass being green and grass being full of chlorophyll, about which we may want to be more realist.

The question of what else could commit one ontologically is a difficult question in meta-ontology, and there are vexing related issues about how best to understand the notions of nature, observation, and explanation. I won't be able to address these issues here, but I do want to explain how a historically prominent idea might serve as a placeholder for a more fully worked-out account.

What I have in mind is the idea that ontological commitment tracks with commitment, not to something being true, but to something being part of the best natural explanation of what we can observe. The rough idea is that our theory of reality is an implicit and incomplete attempt to explain our actual and potential observations of the world. If that's right, and an inferentialist wants to be a realist about a claim such as 'Grass is full of chlorophyll,' then she can say that this claim is implicitly explanatory. That is to say that part of what it obligates one to, inferentially, is a certain explanatory claim – that is, that grass's being full of chlorophyll explains why we can observe certain things about grass. What's important is that there's theoretical space to go in the other direction as well. If an inferentialist wants to be an anti-realist about a claim such as 'It's unlikely that the sun will rise an hour late tomorrow,' then she can say that this claim is not similarly explanatory in its inferential implications. Perhaps it doesn't commit its author ontologically to facts about what has the property of being unlikely, but rather commits its author practically to treating certain future contingents as settled.

Whichever way this contrast is refined, I think we can begin to see theoretical space for the inferentialist to reconstruct the realism/anti-realism debate in metaethics in directly inferentialist terms. If an inferentialist wants to be an ethical realist, then he'll argue that the commitments undertaken when one makes an ethical claim are implicitly explanatory. That is, his account of the inferential implications of this claim will be similar to the one just sketched for the claim 'Grass is full of chlorophyll.' However, if

an inferentialist wants instead to be an ethical anti-realist, then he'll have to argue that the commitments undertaken when one makes an ethical claim are not implicitly explanatory in this way. That is, his account of the inferential implications of this claim will be similar to the one just sketched for the claim 'It's unlikely that the sun will rise an hour late tomorrow.'

I think something like this provides a more nuanced version of inferentialism that is able to capture not only the idea that there is a distinctive connection between ethical discourse and action, but also the prospects of a naturalism-inspired form of anti-realism about the ethical. Since these were the two dialectical advantages expressivists claim over traditional anti-expressivists, and since expressivism faces a problem with the systematicity of its form of semantic explanation that is not faced by the inferentialist, I think this means that inferentialism has better prospects than expressivism as an anti-representationalist theory in metaethics.

## 5. Conclusion

The guiding thought of this paper is that one's metaethical account of the meaning of ethical sentences should fit with one's account of the meaning of sentences in general. In metaethics, one finds non-naturalist, naturalist, fictionalist, and expressivist accounts of the meaning of ethical claims. And, in the theory of meaning, one finds representationalist theories, which take word-world representational purport as their fundamental explanatory notion, ideationalist theories, which take word-idea expressive function as their fundamental explanatory notion, and inferentialist theories, which take claim-claim inferential connections as their fundamental explanatory notion. So, my guiding thought led me to ask how those more specific metaethical theories fit with these more general theories of meaning. In the case of non-naturalism, naturalism, and fictionalism, the answer is relatively straightforward: they are implicitly forms of representationalism (perhaps allowing for the possibility of a 'dogleg' through mentalistic notions). However, in the case of expressivism, the answer was not so straightforward. The crux of this paper has been searching for this answer, articulating problems with it, and using those problems to motivate a kind of metaethical theory that is not on the standard maps of theoretical options in metaethics. This is the theory that fits with or is an application of a more general inferentialist theory of meaning.

I don't take myself to have defended or even articulated this theory fully. My aim here was more modest. I want inferentialism to be considered one of the viable theoretical options in metaethical debates about the meaning of ethical claims. To this end, I've tried to place it as a form of anti-representationalism that differs from expressivism in important respects. First, by appealing to the notion of *an inferentially articulable commitment undertaken* rather than to *a thought conventionally expressed* by a

claim, I think inferentialism has resources to respect the compositional systematicity of meaning, which both early and later versions of expressivism lacked. Second, I think inferentialism retains old resources and brings new resources to explaining the distinctive connection between ethical discourse and action. Third, there is room in the inferentialist theory to reconstruct the realism/anti-realism distinction in a way that doesn't depend on different directions of fit with the world that different types of mental states (beliefs or desires) might be thought to have.

That being said, however, I'd like to close by voicing a worry about inferentialism in metaethics. Even if the sort of account I've pointed to does to some extent work to explain the meaning of ethical claims, one might think that it does so by merely pushing the metaethical question back a level. Inferentialism explains the meaning of ethical claims in terms of the inferentially articulable commitment they can be used to undertake, but what is its account of the meaning of claims about the entitlements and obligations involved in a commitment's inferential articulation? For instance, the claim 'One who is committed to grass's being green is committed to grass's being colored' is not an ethical claim, but it seems to be a normative claim. Doesn't it raise all of the same metaethical issues as ethical claims? It does, but I think it's already some advance if we're able to locate the metaethical issue about meaning and normative character of ethical claims within a broader meta-normative issue about normative claims more generally. However, I think there is also a deep and difficult question about the genesis and nature of these norms. Here, I think more work needs to be done to determine to what extent the metaethical debate will simply be re-engaged at a more fundamental level, or to what extent pursuing it at that level affects the relative attractiveness of the various theoretical options.<sup>23</sup>

## Notes

1. Cf. Shafer-Landau (2003, ch. 3), Parfit (forthcoming: ch. 24). Distinguishing what is non-natural from what is natural is notoriously difficult, especially since non-naturalism is usually thought to involve the denial that moral properties are supernatural as well as the denial that they are natural. See Ridge (2008) and Sturgeon (2009) for useful discussion. Another way, then, to organize theories into the first two camps would be to distinguish between *reductionists* who see moral properties as reducing to some other sort of property, and *anti-reductionists* who deny this. The differences between these ways of organizing the kinds of metaethical views are not important for my present purposes, which are more focused on the semantic assumptions of these theories.
2. Boyd (1988), Sturgeon (1985; 2003), Brink (1986; 1989), Railton (1989), Smith (1994a, ch. 5–6).
3. Mackie (1977, ch. 1) defends a form of Error-Theory, which is the progenitor to modern fictionalism. Cf. Joyce (2001), Kalderon (2005).
4. Blackburn (1984; 1992; 1998), Gibbard (1990; 2003), Timmons (1999), Ridge (2006).

5. A view that does not show up explicitly on my list is the sort of metaethical constructivism defended by Korsgaard (2003) and Street (2008). I remain unsure of its distinctiveness, but I also don't know where to slot it in. See Fitzpatrick (2005) and Hussain and Shah (2006) for useful discussion of constructivism. Response-dependent views such as McDowell (1985) are also hard to place, but that is because it's hard to know whether or not the facts about response dependence to which he appeals are to be thought of as natural facts. In any case, the location of views such as Korsgaard's and McDowell's won't matter for the critical discussion of expressivism to follow.
6. This is consistent with a later attempt to reclaim whatever talk of description is part of ordinary ethical discourse, as the quasi-realist does (Blackburn, 1993; 1998; Gibbard, 2003). It's just that, at the more basic theoretical level, the expressivist doesn't think the idiom of description can serve in the best account of the meaning of ethical claims.
7. He is perhaps omitting the verificationist theories that were prominent at the beginning of the twentieth century, but I think he is right that *contemporary* philosophers of language are often divided into these two broad camps.
8. Admittedly, spelling out the notion of explanatory priority here is not straightforward. For all theorists of meaning will want to recognize that meaning is conventional, in the sense that it's only because of regularities of use that some arbitrary sign/sound-designs have meaning while others don't. However, the truth-theorist sees that as pragmatic background and takes the notion of a truth-condition to be semantically more basic than the notion of a rule of correct use, whereas a use-theorist thinks that we can explain why sentences have the truth-conditions that they have only via semantic rules of use.
9. It would, of course, be possible instead to endorse a non-realist but non-deflationary conception of truth, such as pragmatism or coherentism about truth. However, as theories of truth, these are quite implausible, although pragmatist and coherentist ideas have found their way into general philosophical views naturally allied with deflationist or minimalist accounts of truth. See Wright (1992), Rorty (1995) for relevant discussion.
10. This may involve, as a corollary, a connection to mental states. More specifically, the anti-expressivists mentioned above may want to say that ethical claims serve to express beliefs – but they'll want to understand the contents of these beliefs in representationalist terms. I return to this issue in Section 2.
11. Schroeder (2008, ch. 2) and Blome-Tillmann (2009) make similar points.
12. Schroeder (2008, ch. 3–5) makes a similar point.
13. Compare Schroeder (2008, ch. 4–5). The representationalist's and ideationalist's functions COND and SUBJPRED are not strictly identical but isomorphic. The important point is that, whatever one's favored order of semantic explanation, there's a uniform semantic function between parts and wholes for sentences with the same deep grammatical form.
14. See, especially, Smith (1994a, ch. 4) and citations therein.
15. Some expressivists may be happy with the ordinary practice of calling these states 'beliefs,' but they'll have to insist that their underlying nature is different from descriptive beliefs. See Gibbard (2003, pp. 181–3), Ridge (2009).
16. Again, at one level, the expressivist may insist that there is no problem in calling these 'beliefs,' of a sort. But, in order to maintain his distinctive position in metaethics, he'll have to argue that, at a deeper explanatory level, they're not on a par with descriptive beliefs but more like desires, intentions or plans.

17. See Blackburn (2006, p. 247) where he seems to interpret himself in these terms.
18. There may be *implausible* ways to draw the contrast that are inconsistent with ideationalism. One of these would be to extend recent hybrid views of ethical claims (e.g. Ridge, 2006) to claim that *all* claims express a hybrid belief–desire state. Another comes from Schroeder (2008, ch. 4–9) who argues that the best way for an expressivist to respect the compositional systematicity of meaning is to treat all claims as expressing a state he calls *being for*, which one can take towards different contents. However, Ridge agrees that it would be problematic if we had to say that a descriptive claim such as ‘Grass is green’ means what it does in virtue of expressing a hybrid belief–desire state (personal communication); and the remainder of Schroeder’s book involves drawing out many implausible consequences for the view that all claims express states of being for.
19. See, for instance, Sellars (1969; 1974), Rosenberg (1974), Harman (1982), Block (1987; 1993), Brandom (1994; 2000; 2008).
20. Wedgwood (2001; 2007, ch. 4–5) has championed a conceptual-role account of the meaning of ethical terms. This bears some resemblance to the sort of inferentialism I will go on to sketch, except in the important respect that Wedgwood sees his theory as resolutely realist and so it is in a sense deeply representationalist. As mentioned above, Blackburn (2006) now interprets his own previous views in inferentialist terms.
21. This is consistent, I believe, with the thought that one hasn’t *fully* accounted for meaning until one has explained the conventional expression relations that stand between public language and private thoughts and the representation relations that stand between language/thought and the world. It’s just that the inferentialist thinks that inferential-role is most fundamental for understanding semantic content, which means that, in so far as these other notions are relevant to meaning, they must be conceived ultimately in inferential terms rather than vice versa.
22. See Brandom (2008, ch. 5) for discussion of the top-down strategy for capturing compositional systematicity in an inferentialist framework.
23. I’d like to thank Michael Ridge and the audiences at the University of Glasgow and the University of Sydney for helpful feedback on material for this paper.