Cognitivist Expressivism
Terry Horgan & Mark Timmons
University of Arizona

[For *Metaethics After Moore*, T. Horgan & M. Timmons (eds.). Oxford University Press]

Expressivism is a position in metaethics that is a descendant of noncognitivism—a view that was perhaps the dominant metaethical theory for about 40 years, between 1935 and 1975.[1] The basic insight of the noncognitivists was that language can play a dynamic as well as a descriptive role in interpersonal interaction, and that moral discourse is a prime example of the dynamic use of language. According to one dominant strain of noncognitivism, *emotivism*, championed by A. J. Ayer and C. L. Stevenson, moral judgments function primarily to express one's feelings about some object of evaluation. In an interpersonal context, such expressions of feeling typically serve the dynamic function of influencing the attitudes of others. The other main variety of noncognitivism, *prescriptivism*, developed by R. M. Hare, took dynamic, imperatival utterances as a model for moral discourse, and thus emphasized the directive, action-guiding element of such discourse. Noncognitivism did come in for its share of criticism, depending on the version under scrutiny, but the bottom line seemed to be that this kind of view appears incompatible with what Allan Gibbard calls the 'objective pretensions' of moral thought and discourse, including the idea that moral judgments seem to be beliefs with assertible, truth-apt content. Emotivists and prescriptivists, because they took moral language to express noncognitive attitudes, were forced to explain away such pretensions. So, for instance, Carnap (1935: 25) held that 'a value statement is nothing else than a command in misleading grammatical form'. But noncognitivist attempts to explain away various deeply embedded features of moral thought and discourse have seemed implausible and indeed unnecessary to many moral philosophers.

In the late 1970s and on through the 90s, some moral philosophers have been tempted by moral realism, thinking that with the help of various developments in philosophy of mind and philosophy of

---

[1] Ayer's *Language, Truth and Logic* was published in 1936 and contained the first widely influential presentation of a noncognitivist metaethical theory. As Mary Warnock noted in her *Ethics Since 1900*, 'It is part of the measure of the importance of [Ayer's] book that no sooner was it published than it seemed that emotivists were everywhere. They had not been converted by the book; it was their creed already' (1960: 84).

language from the 1960s and early 70s, one could countenance moral properties and facts and yet remain faithful to philosophical naturalism—the dominant metaphysic of the times. Also, metaethical history lately has been returning to its Moorean roots with some moral philosophers bolding defending versions of non-naturalism.[2] But just as the past 100 years of metaethics has seen realism go and come back again, those working in the tradition of Stevenson and Hare (ourselves included) have devoted time and effort into reviving the spirit, if not the letter, of older noncognitivist positions. Recent work along these lines includes Simon Blackburn's quasi-realism, Gibbard's norm-expressivist view, and our own position, here labeled *cognitivist expressivism*. In some of our previous writings, we have either individually or collaboratively tried to make progress in articulating and defending our view. This paper is another installment in a series of works devoted to this project.

1. Preview of coming attractions

In the days of noncognitivism, the idea that moral judgments are not primarily descriptive of moral properties and facts (nondescriptivism), and the idea that moral judgments do not express beliefs (noncognitivism) were taken to be mutually entailing. Nondescriptivism and noncognitivism were a package deal. And the views *are* mutually entailing if one embraces the following claim, which we call the *semantic assumption*:

SA   All cognitive content (i.e., belief-eligible, assertible, truth-apt content) is descriptive

content. Thus, all genuine beliefs and all genuine assertions purport to represent or

describe the world.[3]

---

[2] We have in mind Audi (2004) and Shafer-Landau (2003).

[3] For present purposes we use 'represent' and 'describe' (and their variants) interchangeably. (This use of 'describe' is at work in metaethics when the label 'descriptivism' is deployed; it is wider in scope that the use in which 'descriptive' and 'normative' operate as contrast-terms.) Beliefs and assertions that 'purport to represent or describe the world' include ones with theological content, and can also include ones about abstract entities like numbers; so 'world' does not just mean 'physical world' or 'spatiotemporal world' or 'world of concrete particulars'.

The semantic assumption has been widely taken for granted in metaethics; it has framed much of the philosophical debate, and has constrained the range of options on the standard menu of competing positions. But this deeply entrenched piece of orthodoxy, we maintain, is false. Its falsity would mean that some metaethical space opens up (potentially anyway) for the combination of nondescriptivism and cognitivism.

We defend just such a view, claiming that although moral judgments are genuine beliefs, their overall content is not descriptive content. Moral judgments count as beliefs, despite being non-descriptive, because they possess enough of the key, generic, phenomenological and functional features of belief (as well as satisfying the relevant platitudes governing the concept of belief) to qualify as genuine beliefs. We defend the claim that moral judgments are genuine beliefs in section 6 below.

In our previous writings, we have used different labels for our position, sometimes calling it 'assertoric nondescriptivism', sometimes calling it 'nondescriptivist cognitivism'. Gibbard (1990: 7-8) uses the term 'expressivism' to refer to metaethical views that take the primary role of moral judgments to be expressive of attitudes that do not purport to represent or describe some moral reality. So expressivism is committed to nondescriptivism about moral judgments and utterances. Expressivism thus subsumes old-time versions of noncognitivism. But it leaves open the possibility of a cognitivist construal of moral thought and discourse. Because the term 'expressivism' has taken hold, we are here calling our view *cognitivist* expressivism.

Cognitivist expressivism is very similar in spirit to Blackburn's more recent presentation of his quasi-realism—it is a metaethical project that embraces an austere irrealist moral metaphysics and yet attempts, in its semantic construal of moral terms and the concepts, to account for the deeply embedded assumptions of moral thought and discourse. The main differences between our view and Blackburn's have to do with philosophical execution.[4]

---

4 Blackburn (1984: 167-71) refers to any philosophical view that regards the judgments of some discourse 'expressive' as opposed to descriptive, an 'expressive theory' about that discourse. So his quasi-realism is a form of expressivism. And in more recent writings (1996) he explicitly distances his ethical expressivism from forms of noncognitivism. See also Blackburn (1998).

In this paper we will not spend time explaining why we reject versions of moral realism, moral constructivism, moral relativism, the error theory, and noncognitivism. We have done that elsewhere.[5] Rather, we plan to articulate and defend our evolving metaethical view in a way that re-packages key ideas from our prior writings while also going beyond our previous work in two important respects. First, we will dwell on matters of moral phenomenology—the "what-it's-like-ness" of experiences involving moral judgment; we will argue on one hand that this phenomenology supports the cognitivist contention that moral judgments are genuine beliefs, and on the other hand that such cognitive phenomenology also comports with the denial that the overall content of moral judgments is descriptive. Second, we will offer a more detailed account than we have provided before of how to accommodate certain crucial generic features of the psychological role of belief-states (and corresponding features of the interpersonal role of moral utterances)—features involving the embeddability of moral content within logically complex beliefs and sentences, and inferences employing beliefs and sentences with embedded moral content.

We begin with an insight to be found in G. E. Moore's metaethical reflections and, with Moore's guidance, we then proceed to develop and partially defend our view.

2. Moore on methodology

In *Principia Ethica*, Moore famously defended the idea that goodness is a 'simple, indefinable, unanalysable object of thought' (1903: 72). He also thought that 'real' definitions of terms—definitions that reveal the essential nature of their referent—are possible only when the term to be defined refers to something complex. Since the property of goodness is simple, having no parts, Moore claimed that 'good' cannot be defined.

The most important sense of 'definition' is that in which a definition states what are the parts which invariably compose a certain whole; and in this sense 'good' has no definition because it is

---

[5] For a very brief overview of the problems for such theories with references to our other work, see Horgan &

simple and has no parts. It is one of those innumerable objects of thought which are themselves

incapable of definition, because they are the ultimate terms by reference to which whatever *is*

capable of definition must be defined. (1903: 61)

Often, the term, 'irreducible', is used in connection with the idea that in some important sense it is not

possible to analyze or define fundamental moral concepts and the terms that express them. Put in these

terms, Moore thought that because the most fundamental concept in ethics, namely, goodness, refers to

something simple, the concept (and the term expressing the concept) is irreducible.[6]

We do not accept Moore's moral realism, and we do not accept his view that a moral term like

'good' refers to a property. But we do adhere to what we call 'Moore's methodological maxim': moral

terms and concepts are irreducible. However, to say this does not mean that that nothing philosophically

illuminating can be said about them (in addition to their being irreducible). R. M. Hare, for instance, had

quite a lot to say about the semantics of moral terms and concepts, which, if true, is very illuminating.

However, in making use of imperatives to understand moral language, Hare insisted that 'it is no part of

my purpose to "reduce" moral language to imperatives' (Hare, 1952: 2).[7] Rather than offer a reductive

analysis of moral terms and concepts which would, in effect, express those terms and concepts in some

sort of nonmoral idiom, Hare proposed to understand terms like 'ought' by exploring the 'logical

behavior' of these words in ordinary language via similarities between such moral language and

imperatives.

Of course, by taking imperative sentences as a model for the 'logical behavior' of moral terms and

concepts, Hare did deny in effect that moral judgments are genuine beliefs and that moral utterances are

genuine assertions. On his imperatival model, the declarative grammatical form of moral sentences is

misleading, since declarative sentences normally are employed to make assertions and to express beliefs.

Timmons (2004).

[6] It is worth noting that Moore himself thought that talk of what is intrinsically good is equivalent in meaning to talk about what *ought to exist* (1903: 33-4, 68, 166). So, roughly, the kind of reduction Moore was anxious to rebut would involve defining basic moral terms and concepts by using nonmoral terms and concepts.

[7] See also Hare, 1952: 180-1.

So his metaethical position can be called *weakly* reductive in its use of non-assertoric, non-declarative, sentences as a model of moral thought and language, even though it eschews the *strongly* reductive claim that moral utterances are synonymous with, or semantically interchangeable with, imperatives.

Like Hare, we propose to explore the meaning of moral terms like 'ought' by considering how such terms function in thought and language. But we plan to give more weight than did Hare to moral thought (as opposed to moral language), and we will give specific emphasis to matters of phenomenology. Also, again like Hare, we will argue that sentences with certain distinctive grammatical features are a useful model for understanding moral thought and discourse (viz., sentences in a specific sort of formal language, described in section 6). But our linguistic model will not be 'reductive' in even weak senses, because it fully accommodates the cognitivist claim that moral judgments are beliefs and moral utterances are assertions. In eschewing any sort of 'reductive' account of moral terms and concepts, we will be respecting Moore's (anti-reductive) methodological maxim.

In articulating and defending our view, we set for ourselves three main tasks that we will proceed to take up in order in the following sections.

Task I: Describe some key generic features of beliefs, and argue that moral judgments exhibit these features.

Task II: Explore some key distinctive characteristics of moral judgments in particular, as contrasted with ordinary nonmoral beliefs.

Task III: Set forth a theoretical account of belief that simultaneously (i) treats moral judgments as genuine beliefs, (ii) treats moral judgments as not descriptive in their overall content, (iii) accommodates the key distinctive characteristics of moral judgments in a way that renders these features consistent with the claim that moral judgments are genuine beliefs, and (iv) accommodates the key generic features of belief in a way that is consistent with the denial that the overall content of moral beliefs is descriptive.

In the course of pursuing the first two tasks, certain specific challenges will emerge that will need to be faced in addressing the third task: some of the distinctive features of moral judgments threaten the idea that these judgments are genuine beliefs, whereas some of the generic features of belief exhibited by moral judgments threaten the idea that these judgments are non-descriptive in their overall content. These challenges will be noted as the first two tasks are pursued, and will be addressed in the course of dealing with the third task.

By completing these tasks we will provide a presumptive case in favor of cognitivist expressivism, but there remain further issues and challenges that our view must meet. Late in the paper, we consider some of the most pressing of these, and we briefly explain how our view attempts to meet them.

3. Terminological preliminaries.

Before proceeding, it will be useful to make some explicit remarks about matters of terminology. We begin with some observations about our use of 'moral *judgment*'. First, we use this term in a metaethically neutral way to refer to those psychological states whose contents are expressible by a moral sentence. Thus, calling this sort of state a judgment leaves open whether it is a belief, a desire, an intention, or some other psychological state. Below we argue that moral judgments are most plausibly understood as beliefs. Second, like Mandelbaum (1955: 46), our use of the term is intentionally broad in another way: what we are calling moral judgments need not be psychologically inferential; they might be psychologically spontaneous as when one just 'sees' that some action is obligatory. Third, 'judgment' allows for process/product ambiguity in its usage, i.e., between an episode of judging and being in the psychological state resulting from a judging episode. (Likewise for 'belief'.) Context should make clear how we are using the term.

We have been using, and will continue to use, the currently widespread term 'content' in connection with moral judgments and utterances. We discuss content of various kinds: the *overall* content of a judgment or utterance (the content of the 'that'-clause employed to describe the state, or the sentence

employed to express it), *cognitive* content (the kind of content that is belief-eligible and truth-apt), *descriptive* content (the kind of content that represents, or constitutes, a way the world might be), and *non-descriptive* content.

Although 'content'-talk is extremely natural and convenient in metaethics and in other branches of philosophy (as was 'meaning' talk in earlier decades), we stress that in relation to cognitivist expressivism, this terminology needs to be taken with a metaphysical grain of salt. On our account, talk of non-descriptive content is to be understood as not really positing any such *items* as overall contents or cognitive contents; likewise for generic talk of overall content and of cognitive content, construed as encompassing non-descriptive as well as descriptive content. (We will take up 'descriptive content' presently.) Rather, such talk is both syncategorematic and pleonastic.

It is *syncategorematic* in the sense that one can use such talk only in a fairly restricted, fairly specific, range of syntactic-grammatical contexts. One can talk about a psychological state or a declarative sentence as 'having cognitive content', about its 'having cognitive content that is not descriptive', and so forth. One can even use definite descriptions like 'the cognitive content of the belief', but only in certain specific kinds of sentential contexts (e.g., contexts like 'The overall content of his belief is that Jones ought to apologize'.) But such talk is not rightly construed as positing some *entity* that *is* a non-descriptive cognitive content. On our picture, there is no such entity.

Generic content-talk is *pleonastic* in this sense: it normally functions as a way of saying something that could pretty much be said some other way (although the other way needn't be outright *synonymous* with the original way). For example, saying that moral judgments, and the utterances expressing them, 'have cognitive content' is often, in effect, a way of saying that the judgments are genuine beliefs and that the utterances are genuine assertions. Likewise, saying that a psychological state has 'overall content' is often, in effect, a way of saying that it is a state describable via a mentalistic characterization employing a 'that'-clause construction.

Although such syncategorematic, pleonastic usage is very convenient, one can employ it as we do to facilitate communication and exposition even if one denies, as we do, that there are any such in-the-

world *items* as non-descriptive contents. (Compare: One can sincerely utter the sentence 'He has a loud voice', and this sentence can be true, even though one's usage does not incur any ontological commitment to such in-the-world items as *voices*.) The claim that this terminology is 'ontologically lightweight' will receive theoretical justification from within our own metaethical position, as articulated below.

Talk of *descriptive* content, on the other hand, is a different matter, from our theoretical perspective. In this paper we will assume, at least provisionally, that there *are* such items as descriptive contents—ways the world might be, and or ways that some thing (or things) might be. Below we will invoke two kinds: (i) *propositional* descriptive contents, expressible in English via closed sentences like 'Snow is white', and (ii) *non-propositional* descriptive contents, expressible in quasi-English via open sentences like 'x is white'. We will leave it open how exactly to construe such items metaphysically.[8]

We now proceed to our three tasks.

4. Task I: key generic features of belief

We begin with some prototypical, generic, features of belief—features associated respectively with the phenomenology, the semantic assessability, and the functional role of beliefs. We begin with the former. By 'phenomenology of belief' as used here, we mean to refer to the sort of subjective, what-it-is-like experiential aspect of occurrent beliefs, something that unfortunately has largely been ignored by contemporary philosophers of mind.[9] In making various observations about the phenomenology of belief,

---

[8] Perhaps even talk of *descriptive* content ultimately should be construed as pleonastic, syncategorematic, and not ontologically committal. But even if this turns out to be so, such talk is still importantly different metaphysically from talk of non-descriptive content: sentences and psychological states 'with descriptive content' have in-the-world *truth-makers* of some sort, whereas utterances and psychological states 'with non-descriptive content' do not. That difference would need to be respected and reconstructed, within any ontological approach that treats talk of descriptive content as being syncategorematic and pleonastic itself.

[9] That occurrent beliefs have a distinctive phenomenology was recognized by Hume who thought this aspect of beliefs was crucial for understanding the difference between an occurrent belief with a certain content and various non-belief states having the same content—states such as entertaining without believing. Hume distinguishes between the 'ideas' (roughly contents) and the manner of considering those ideas, and explains:

> [B]elief consists not in the nature and order of our ideas, but in the manner of their conception, and in their feeling to the mind. I confess 'tis impossible to explain perfectly this feeling or manner of conception. We make use of words, that express something near it. But its true and proper name is *belief*, which is a term that every one sufficiently understands in common life. And in philosophy we can go no farther, than assert, that it is something *felt* by the mind, which distinguishes the ideas of the judgment from the fictions of

we will be discussing ordinary nonmoral descriptive beliefs that most clearly purport to represent or describe some aspect of the world. But we will also be discussing occurrent moral ought-judgments, calling attention to those phenomenological features that they share with nonmoral descriptive beliefs and which (together with considerations of semantic assessability and functional role) we claim qualify them as genuine beliefs.

The phenomenology of belief

Here, then, is a list of five interrelated features of what we are calling the phenomenology of occurrent belief. This what-it-is-likeness typically involves: (1) psychologically 'coming down' on some issue, in a way that (2) classifies (sometimes spontaneously) some 'object' of focus as falling under some category, where one's classificatory coming down is experienced (3) as involuntary, (4) as a cognitive response to some sort of consideration that is experienced (perhaps peripherally in consciousness) as being a sufficient reason for categorizing as one does, and (5) as a judgment that is apt for assertion and hence is naturally expressible in public language by a sentence in the declarative mood. Each of these elements of phenomenology requires comment.

Occurrent belief involves the experience of 'coming down' on some issue, where this manner of coming down may be preceded by some amount of reflection, as when one is looking up at the October night sky and taking a close look at a very bright celestial object and, after staring for a moment, comes to believe that it is Mars. But very often belief comes about as an automatic response to one's surroundings, as when one catches a glimpse of a passing car while watering the lawn. In this kind of case, by spontaneously and unreflectively taking it to be a car, one spontaneously psychologically comes down with respect to the object-kind instantiated by a moving object passing through one's field of vision.

---

imagination. It gives them more force and influence; makes them appear of greater importance; infixes them in the mind; and renders them the governing principles of all our actions. (1739: 629).

Hume, of course, attempts to use the features of force and vivacity to explain the nature of belief. Although we agree with Hume that belief has a distinctive phenomenology, we do not follow him in trying to make sense of this phenomenology (and associated functional role) in terms of force and vivacity.

Turn now to all-things-considered occurrent ought-judgments. One obvious feature of so judging is that one experiences this activity as a matter of psychologically 'coming down' on whatever issue is under consideration. As with descriptive beliefs, coming down with respect to some moral issue may be preceded by deliberation. After listening to various opposing opinions about the morality of same-sex marriages, Jones finds himself accepting the view that such marriages are not morally wrong and that therefore a government ought not to make such unions illegal. Jones, we are supposing, is initially undecided about this issue, but later and in light of reflecting on various facts about same-sex marriages, finds himself 'coming down' on one side of the issue. Of course, many of our moral judgments, like descriptive beliefs, are formed spontaneously: they are triggered immediately in experience. Harman's (1977: 4) famous example of seeing some hoodlums igniting a cat and, without deliberation, coming to think that what they are doing is wrong, is a case in point. In such cases, one's moral judgment is experienced as a fairly unreflective and immediate 'coming down' morally with respect to what one sees. Indeed, here people often talk about just 'seeing' an action as wrong.

Clearly, making decisions and forming intentions is likewise experienced as coming down with respect to some issue. Since intentions are not beliefs (we are supposing), the experience of psychologically coming down is not alone a sufficient phenomenological indicator of belief. The kind of psychological coming down that seems fairly distinctive of belief, as opposed to intention-forming and the like, involves experiencing perceived or contemplated items as falling under categories. Things get *sorted* experientially, when one comes down on a matter in the belief-wise way. In the case of descriptive beliefs, things get sorted experientially into descriptive categories, via category-concepts that purport to represent some worldly object-kind or some property. Likewise, in the case of a moral judgment that some action ought not to be done, a particular act-type or act-token gets sorted experientially into the category of those actions that one is not to perform. We claim that the sorting/categorizing aspect is central to the phenomenology of belief.

Another salient phenomenological feature of beliefs is their involuntariness. One looks out the window and spontaneously and involuntarily believes that the sun is shining. Sometimes reflection

precedes one's involuntarily coming down on some issue. After mulling over various bits of evidence concerning a mechanical failure, a trained mechanic involuntarily comes to believe that the failure was due to a faulty oil pump. After inspecting the evidence, she just sees what must have caused the failure. Regarding ought-beliefs, we have already mentioned Harman's case of unreflectively and spontaneously coming to have a moral belief about the hoodlums, as well as a case in which Jones comes to have a belief about same-sex marriages preceded by some amount of reflection.[10]

Related to the involuntariness of beliefs is the fact that they are experienced as possessing a kind of *rational authority*, consisting in their being grounded by *reasons*. In the case of ordinary nonmoral perceptual beliefs about objects and their properties, this experienced authority issues from one's perceptual experiences. In moving my hand across the desk, I feel its smoothness and spontaneously come to believe that it is smooth. The confidence I experience in having this belief is arguably part of the overall experience of occurrently believing and partly explains why involuntarily psychologically coming down, rather than just entertaining a corresponding hypothesis, is appropriate. Similarly, the experience of morally coming down on some issue includes experiencing this state as possessing a kind of rational authority—an authority grounded in factual considerations that are experienced as rationally grounding one's moral judgment. In matters of morality, it is not 'up to' an individual what to think about some matter of moral concern, just as it is not 'up to' someone what to think about some nonmoral factual matter of concern. Sometimes this feature of phenomenology is put in terms of the idea that judgments of moral obligation have an 'external' authoritative source, just as do nonmoral factual beliefs. It is by dwelling further on the phenomenology of moral authoritativeness that one can begin to distinguish ought-beliefs from nonmoral descriptive beliefs, which, as we shall see, leads to considerations of the distinctive functional role of such beliefs. We save further discussion of these matters for the next section.

Finally, the categorizing manner of psychologically coming down, in occurrent belief, is experienced as a mental affirmation that is *apt for assertion*, and thus is experienced as being naturally and

---

[10] The spontaneity involved in both moral and nonmoral expertise is a major theme in Dewey (1922), and is stressed in Dreyfus & Dreyfus (1990), and Varela (1992).

spontaneously expressible in public language by a sentence in the declarative mood.[11] Declarative-mood sentences are the standard public-language vehicles employed for the speech acts of asserting.

To sum up so far: There are a number of phenomenological features characteristic of belief as a distinct state type. These experiential features are possessed by moral judgments, and so there is good prima facie reason for claiming that such judgments are beliefs.

Prima facie, these very features can easily seem to require a construal of moral judgments and moral utterances as being descriptive in their overall content—or at least as *purporting* to be descriptive, even if there are no in-the-world moral properties or facts. How, one might wonder, can one make good sense of the classificatory, involuntary, reason-based, coming-down-ish, phenomenology of moral judgments—including the experience of such judgments as being apt for assertion—except by supposing that their overall content is descriptive in purport? Addressing this challenge will be part of Task III.

Semantic assessability

It is grammatically permissible, and also common in practice, to ascribe truth and falsity to moral judgments and statements. Such judgments and statements thus appear to be *semantically assessable*. Furthermore, truth ascription seems entirely natural and appropriate, given the belief-ish features of moral phenomenology lately noted. If one definitively 'comes down' on the apartheid issue by finding oneself having formed a judgment expressible as 'Apartheid is wrong', and one expresses this moral judgment by asserting that apartheid is wrong, then one will also think, and will be ready to assert, that it's *true* that apartheid is wrong.

Semantic assessability presents an obvious challenge to moral irrealists: making sense of truth ascription and falsity ascription within a general approach that treats moral judgments, and the declarative sentences expressing them, as having overall content that is not descriptive content. We need to deal persuasively, within our cognitivist expressivism, with the line of thought that says, 'Well, if moral judgments lack descriptive, way-the-world-might-be, content, then their content cannot *correspond*

---

[11] This aspect of belief is a case where the phenomenon we are pointing to is plausibly understood as involving both

or *fail to correspond* with how things really are; so, such judgments cannot *really* be either true or false, even though people often apply the terms "true" and "false" to them in ordinary discourse'. Addressing this challenge too will be part of Task III.

Functional role

Beliefs are often characterized functionally as being psychological states that combine in distinctive ways with other psychological states to rationally-inferentially yield further content-appropriate states as well as action. This is at least a partial characterization of their typical functional role, which is clearly illustrated in cases where one's belief combines with one's desires leading one to form intentions (assuming that intentions represent a distinct kind of psychological state) — both long term and short term. Intentions in turn lead to action. My desire to eat a taco for lunch together with my belief that Taco del Cielo is around the corner (as well as a host of other content-appropriate beliefs and desires) combine to yield an intention to go to Taco del Cielo in a few minutes.

The *generic* rational-inferential functional role of beliefs is a matter that needs to be approached with some care, in our view. Although *descriptive* beliefs do typically generate intentions and actions only in combination with conative states like desires, *moral* beliefs — if such there be — might very well have somewhat different prototypical functional roles in human mental life. (More on this in section 6.) However, one crucially important and fully generic aspect of the rational-inferential role of beliefs is this:

> The overall content of any given belief can occur as an *embedded content-constituent* of other, logically more complex beliefs; and beliefs with such embedded content-constituents interact with other beliefs to generate new beliefs via logical relations among their contents (specifically, via logical-entailment relations).

We will call this feature *logical embeddability*. A familiar example is the following piece of moral reasoning:

---

phenomenological and functional role features.

If it's wrong to steal, then it's wrong to get your little brother to steal.

It's wrong to steal.

Therefore, it's wrong to get your little brother to steal.

Logical embeddability has often been posed as a challenge to non-cognitivist positions. If what one is really doing in making a moral utterance is something like expressing an attitude of disapproval or issuing a command, then how exactly can one make sense of utterances in which moral content occurs embedded, like the first premise in the argument just above? And how can one make sense of the judgments expressed by such utterances?

Whatever one might think of the prospects for dealing with this challenge within traditional noncognitivism, our own challenge is a different one, viz., to make sense of the relevant, logically complex psychological states as full-fledged *beliefs*, and to make sense of logical-inferential relations involving such beliefs, without supposing that moral belief-content is descriptive content.

This challenge is closely related to the preceding one about semantic evaluability. Inferential connections among beliefs (and among sentences) are normally explained in terms of truth and falsity: if the beliefs that are the premises of an argument are true, then the belief that is the conclusion must be true (in the case of logical entailment) or must be made likely to be true by the premises (in the case of cogent inductive reasoning). Accommodating logical embeddability thus goes hand in glove with accommodating truth-aptness. Yet another burden for Task III.

This completes our first task of exploring the key generic features of belief while at the same time calling attention to the fact that moral judgments exhibit these features. As we have been saying, this creates at least a strong prima facie case for genuine moral belief—while also posing various challenges needing to be addressed by our cognitivist expressivism. We now turn to various phenomenological and associated functional role characteristics that are distinctive of moral judgments, with specific attention to moral ought-judgments. (Hereafter we will explicitly call these judgments *beliefs*.)

5. Task II: the distinctive nature of ought-beliefs

As noted already, beliefs generally, and moral beliefs in particular, are experienced as grounded by a kind of rational authority. Talk of 'rational authority' while suggestive is quite vague, at least so far. But we can begin to illuminate this idea by dwelling further on matters of moral phenomenology.

In characterizing the experienced authority attendant to first-person ought-beliefs, we follow the lead of Maurice Mandelbaum who, in his 1955 *The Phenomenology of Moral Experience*, distinguished what he called 'direct' from 'removed' judgments of moral obligation. Direct moral beliefs (judgments) are first-person ought beliefs, formed in some context in which the agent experiences a 'felt demand' on her own behavior, while removed ought-beliefs are about what someone else (including one's past self) ought to do or ought to have done. For the time being, we will focus on direct ought-beliefs, making passing references to removed ought-beliefs. Later, in section 8, we will return to the latter type of ought-belief.

Mandelbaum characterizes the felt demand involved in direct moral ought-beliefs as a complex phenomenon involving the experience of an *origin* and a *direction*. In judging that I ought to perform some action—that the action is morally obligatory upon me—I experience a felt demand placed upon my choices and consequent action *issuing from* the circumstances that I confront (at least as I believe them to be). Thus, phenomenologically, one experiences the demand as having an origin that is 'external' to oneself. It is this element of the kind of felt demand characteristic of judgments of moral obligation that distinguishes moral demands from the felt demands that are experienced as having one's own 'internal' desires or preferences as their origin. So, whereas one's own desires may issue forth in a kind of 'internal' felt demand for their satisfaction by the agent who has them, the demandingness associated with a judgment of moral obligation is 'external', coming from, as it were, the morally relevant facts of the circumstances in which one finds oneself on some occasion. This aspect of one's moral experience constitutes, then, the particular manner in which moral ought-judgments are experienced as being grounded in 'objective' reasons. Such reasons are factual considerations confronting the agent that she

16

takes to be morally significant. These considerations are experienced as themselves demanding a certain course of action; they are, as Mandelbaum puts it, the 'origin' in the phenomenology of felt demandingness.

The so-called direction of a felt demand has to do with whom, relative to the judger, the obligation is directed against. In direct ought-judgments, the obligation is directed against the judger herself; this is what Mandelbaum calls a 'reflexive demand'. By contrast, removed judgments of moral obligation are experienced as directed against someone other than the individual making the judgment. This element of the phenomenology of the experience of direct ought-beliefs brings us to issues having to do with their distinctive functional role.

Certain kinds of moral belief—direct ought-beliefs in particular—typically play a motivationally 'hot' functional role in human mental life: they have motivational force in and of themselves, apart from any pre-existing desires or other 'pro-attitudes'. The thought that first-person ought-beliefs are more directly action-oriented than are ordinary nonmoral descriptive beliefs is what inspires various forms of ethical internalism. This thought seems right and important, even though a proper characterization of internalism—one that allows, for instance, for the conceptual possibility of an 'amoralist' with moral beliefs that have no motivational force at all—is a delicate matter.

Also, even though ought-beliefs are typically hot cognitions and thus need not combine with a pre-existing desire in order to provide sufficient motivation to action, a sensible internalism should allow that at least in atypical cases, such ought-beliefs can work in concert with desire to produce action. Suppose, for instance, that normally Joe's judging that he ought to do something is sufficient to move him to action. However, on some occasions, where Joe is perhaps suffering from mild depression, his having a direct ought-belief will only move him if he has certain content-appropriate desires which serve to provide a supplemental motivational spark strong enough to move Joe to do his duty. Perhaps Joe has a strong desire to maintain a certain self-image, so strong that even in a state of mild depression his focusing on this desire of his, together with his belief that failing to do what he ought would seriously damage his self-image, moves him to act. So, even though direct ought-beliefs have as part of their typical

17

functional role a direct motivational force independent of any pre-existing desires, such psychological states are still capable of combining with prior desires to yield intention and subsequent action.

To sum up: we saw in the previous section that general phenomenological considerations support the claim that moral judgments are genuine beliefs: they involve an involuntary, categorizing, way of psychologically coming down on some issue of moral concern, on the basis of considerations that are experienced as rationally requiring the judgment — where this judgment is experienced as truth apt and hence as naturally expressed in thought and language by sentences in the declarative mood. Moreover, the overall contents of moral judgments can occur as embedded constituents in logically complex judgments, which then can combine with other judgments to logically-inferentially generate new judgments. These are key generic features of belief. The presumptive case for moral belief is thus strong.

However, moral judgments and, in particular, direct ought-judgments exhibit some distinctive phenomenological and functional role characteristics too; in particular, typically they are motivationally hot. So yet another challenge to be addressed as part of task III is to explain how such judgments, despite sometimes being motivational all by themselves, nonetheless can be full-fledged beliefs. This brings us to our third task — the task of sketching a framework for belief that can accommodate what we have been saying about belief in general and moral belief in particular and which treats moral beliefs as non-descriptive in their overall content.

6. Task III: a framework for belief[12]

Earlier we noted that one deeply entrenched assumption of metaethics and philosophical inquiry generally — the *semantic assumption*, as we call it--is that *all* beliefs are descriptive beliefs. We reject this assumption. Having made our prima facie case for moral belief, we turn now to the task of providing a general framework for understanding belief (and also assertion) which incorporates beliefs (and

---

[12] This and the following two sections derive partly, but with significant refinements, from sections I and III-V of Horgan & Timmons (2000a).

assertions) whose overall content is not descriptive—i.e., does not purport to represent the world as being a certain way.

Two logically fundamental belief types: is-commitment and ought-commitment

We begin with the logically most basic kinds of belief, leaving embeddability matters until later. On the account we recommend, a belief is a certain kind of *commitment state*—an *affirmatory* commitment—with respect to a descriptive content that we call the belief's *core* descriptive content. There are two logically fundamental belief types: *is*-commitment with respect to a core descriptive content, and *ought*-commitment with respect to a core descriptive content. For example, the belief that *it is the case that Bush is U.S. president*, and the belief that *it ought to be the case that Bush is U.S. president*, are respectively an is-commitment and ought-commitment vis-à-vis the same core descriptive content, viz., *that Bush is U.S. president*. (In the case of the ought-commitment, the core descriptive content, *that Bush is U.S. president*, differs from the belief's overall content, *that it ought to be the case that Bush is U.S. president*; and this overall content is not itself descriptive. In the case of the is-commitment, on the other hand, the core descriptive content coincides with the overall content, viz., *that Bush is U.S. president*.)[13]

These two species of belief, involving two distinct ways of mentally affirming a core descriptive content, are both sui generis: neither type of mental state is reducible to the other, and neither type is reducible to some kind of non-belief state such as an attitude of approval or a mental state appropriately expressible via an imperative sentence. The way to understand the nature of the beliefs we are calling ought-commitments is not by trying to reduce them to something else, but rather by delineating their key features—both the generic features they share in common with descriptive beliefs (is-commitments), and the distinctive features that set them apart from descriptive beliefs. This is just what we have lately been doing, in the course of addressing Task I and Task II.

---

[13] The 'that'-clause 'that Bush is U.S. President' expresses a *way the world might be*, as does the 'that'-clause 'that Gore is U.S. President'. Neither clause says that the world *is* the way expressed (or that the world *ought* to be the way expressed), because 'that'-clauses by themselves don't have assertoric force.

Recall that Task III was formulated as follows: Set forth a theoretical account of belief that simultaneously (i) treats moral judgments as genuine beliefs, (ii) treats moral judgments as not descriptive in their overall content, (iii) accommodates the key distinctive characteristics of moral judgments in a way that renders these features consistent with the claim that moral judgments are genuine beliefs, and (iv) accommodates the key generic features of belief in a way that is consistent with the denial that the overall content of moral beliefs is descriptive. Part (i) is addressed by what we have just said, in conjunction with our treatment of Task I in section 4 above: a moral judgment is an ought-commitment with respect to a core descriptive content, and ought-commitments are a species of belief because they possess the key generic features of belief described in section 4.

Concerning part (ii), the crucial point is that ought-commitment is a fundamentally different kind of affirmatory mental stance toward a core descriptive content than is-commitment. An ought-commitment is not a mental state whose overall content is descriptive, representing a way the world might be; hence it is not a state of mentally affirming that the world *is* such a descriptively-represented way. To construe moral beliefs in this manner is to mistakenly assimilate them to descriptive beliefs, i.e., to is-commitments. Rather, an ought-commitment is a *distinct* kind of mental affirmation vis-à-vis a core descriptive content. Although there is a certain temptation to assimilate ought-commitments to non-belief states of the sort expressible linguistically by non-declarative utterances like 'That Bush is U.S. President, boo!', or 'U.S. citizens, do not elect Bush as U.S. president!', this temptation should be resisted—and can be, once one gives up the semantic assumption. Ought-commitment is a sui generis type of mental state, while also being an irreducible species of belief. Although the overall content of ought-commitments is non-descriptive, nevertheless these states exhibit the key generic features that qualify them as beliefs.

Concerning part (iii), the motivationally 'hot' psychological role typically played by first-person ought-judgments now gets smoothly accommodated. Although this feature makes trouble for the idea that moral judgments are beliefs insofar as one assumes that all beliefs have overall content that is descriptive, it makes no trouble at all for us, because our framework rejects the semantic assumption SA and treats ought-commitments as a distinctive species of belief whose overall content is non-descriptive.

20

Motivational hotness is an important aspect of what *constitutes* ought-commitment in typical cases—although we think that an adequate moral psychology also should allow for cases in which the typical motivational force of moral belief is outweighed, or suppressed, or even silenced altogether by other psychological states of the agent.

Part (iv) of Task III will require more extensive treatment. We will segment the discussion into several subsections that address respectively the three challenges noted earlier in discussing tasks I and II: accommodating the phenomenology of belief, accommodating truth-aptness, and accommodating the key generic functional-role feature of belief, viz., inferential embeddability.

Accommodating the phenomenology of belief

In section 4 we described some key generic phenomenological features of belief, and we argued that moral judgments exhibit these features. Is the possession of these characteristics consistent with the contention that moral judgments are ought-commitments whose overall content is non-descriptive?

Indeed it is. One can experience an occurrent ought-commitment as an involuntary, classificatory, coming-down state vis-à-vis some descriptive content, even though the overall content of this state is not descriptive. The phenomenologically classificatory aspect of this coming-down state, expressible in language via moral terminology (e.g., via 'ought'), need not be a matter of experiencing oneself to be mentally attributing some putatively in-the-world moral *property* to some act, agent, or state of affairs. Rather, it can perfectly well be the experiential manifestation of the specific *mode of affirmatory commitment* that the agent now instantiates with respect to the given descriptive content—viz., ought-commitment.

Likewise, an occurrent ought-commitment can be based psychologically upon descriptive considerations that are experienced as *rationally grounding* this commitment state independently of the morally judging agent's pre-existing desires, even though the overall content of the commitment-state is not descriptive.

Furthermore, an occurrent ought-commitment can be experienced as apt for assertion by virtue of its involuntarily classificatory phenomenology and its experienced authority, despite not being

descriptive in its overall content. So, since the declarative grammatical mood is the appropriate linguistic

vehicle for assertion, and since a sincere assertion is normally a belief-expressing speech act, it is no

surprise that occurrent ought-commitments are experienced as psychological states appropriately

expressible linguistically via declarative sentences.

In short, the generic phenomenological features in virtue of which moral judgments count as

beliefs can perfectly well be present even if, as we maintain, moral beliefs are ought-commitments whose

overall content is non-descriptive. These phenomenological features do not, in and of themselves, build

descriptivity into moral judgment.

It remains possible even so, however, that the phenomenology of moral judgment does include

descriptivity, even though the phenomenological features described in section 4 do not *themselves* entail

it. That is, it remains possible that the full what-it's-like of classificatory coming-down, in moral

judgment, includes the experience of *predicating a putative in-the-world moral property*. Well, does it in fact

include such an experiential dimension?

Introspection, we submit, yields no ready or obvious answer to this question. Moral belief is

experientially much like descriptive belief, to be sure: both kinds of state involve the experience of

involuntary, classificatory, coming-down that is based upon considerations as grounding-reasons. But

beyond this patent and powerful phenomenological similarity in the two kinds of belief, can one also

detect introspectively that moral belief is clearly like—or clearly *unlike*—descriptive belief with respect to

descriptivity per se? It seems not, either way. Rather, the presence or absence of descriptivity as an aspect

of the phenomenology of moral belief is a subtle question about which introspection does not deliver a

confident judgment.

This being so, wider theoretical considerations weigh in on the matter. For one thing, there is no

particular reason why moral phenomenology *should* include descriptivity, given its functional role in

human cognitive economy and given the sociological role of moral discourse in human social intercourse.

On the contrary: since there are powerful theoretical reasons for denying the existence of in-the-world

moral properties and facts, the presence of descriptivity within moral phenomenology would constitute a

built in experiential *error* with respect to the nature of external reality. Since descriptivity would be a gratuitous and erroneous aspect of moral phenomenology, probably it is not really an aspect of moral phenomenology at all.

Another theoretical consideration in support of this conclusion is the fact that first-person moral judgments typically are motivationally hot, despite qualifying as genuine beliefs. Descriptivity, though, would not comport smoothly with this distinctive motivational role. Why should a belief concerning putative in-the-world facts of any kind, qua *factual* belief, be intrinsically motivating? Surely it comports better with the motivational role of moral judgments to say that they are beliefs of a different, *non-factual*, kind: viz., non-descriptive ought-commitments.

These theoretical considerations are not conclusive, admittedly. The question whether moral phenomenology includes an aspect of descriptivity strikes us as ultimately empirical—albeit an issue whose investigation would need to be methodogically very subtle, and might well need to incorporate especially careful introspection. But meanwhile, two key points need emphasis.

First, we take it that the dialectical burden is on those who would claim that moral phenomenology *does* include descriptivity; they need to make a case for this claim, given that simple introspection delivers no clear verdict either way.

Second, even if it should turn out that descriptivity is indeed an aspect of moral phenomenology, this would not be terribly damaging to our metaethical position. Although we would then be forced to claim that there is an erroneous *element* in moral phenomenology, this would not undercut our contention that moral judgments already qualify are beliefs anyway, for reasons independent of their descriptivity — the reasons set out earlier in this paper. Nor would it undercut our claim that moral judgments are a distinctive species of belief, viz., ought-commitments. Qua ought-commitments, moral judgments would be a sui generis kind of belief already, even if they turned out *also* to be erroneous is-commitments that

mistakenly predicate putative in-the-world moral properties. And they would still play a crucial and legitimate action-guiding role in human psychology.[14,15]

One final point. Moral phenomenology may very well be susceptible to influence by higher-order beliefs about the nature of morality itself. Certainly many people believe that there are objective moral facts—a belief that can easily be instilled, for instance, through the persistent intertwining of religious instruction with moral education. For those who believe (perhaps only implicitly) in objective moral facts, there may well arise a derivative kind of moral phenomenology—induced by the interaction of this higher-order belief with the more universal aspects of moral experience—that does include descriptivity. But even if such erroneous moral phenomenology sometimes occurs by virtue of the permeating effects of false beliefs about the metaphysics of morals, we contend, the more fundamental, more universal, kind of moral experience does not include an aspect of phenomenological descriptivity.

Accommodating semantic assessability

The concepts of belief, assertion, and truth are interconnected by a battery of platitudes. For example, to sincerely assert some claim is to express one's belief regarding that claim; to believe a claim is to take that claim to be true; and so forth. Since, on our view, moral judgments are genuine beliefs and moral utterances are genuine assertions, our position must be able to make sense of attributions of truth and falsity to moral judgments and utterances. Since we advocate a robust form of irrealism in ethics, we claim there are no in-the-world moral facts that could serve as truth-makers for moral beliefs and

---

[14] J. L. Mackie's version of error theory is much more philosophically problematic than would be the kind of error theory just described, because Mackie in effect construed moral judgments as being *only* descriptive beliefs whose contents involve putatively in-the-world moral properties. If that's all there is to being a moral belief, and if there are no such properties, then it becomes hard to see why moral thought and moral discourse are not so hugely mistaken that they should be abandoned altogether. But if moral beliefs are also ought-commitments with respect to core descriptive contents, then moral beliefs and moral assertions have important, legitimate, and indeed indispensable psychological and sociological roles to play even if they also include an erroneous aspect of descriptivity.
[15] Also worth noting is that moral realists too could accept our contention that moral judgments are ought-commitments, and could accept that ought-commitments are a distinctive, sui generic species of belief. A moral realist could claim that a moral judgment is *both* an is-commitment with respect to its overall content *and* an ought-commitment with respect to its core descriptive content. This approach would have the advantage, for moral realists, of allowing them to acknowledge the internalist aspects of moral judgment. For further discussion of this theme, with explicit attention to the reasons for preferring our own cognitivist expressivism to such a view, see Horgan and Timmons (2000a), especially section V.

assertions. Moreover, since we are nondescriptivists about moral thought and discourse, we claim that moral beliefs and assertions lack overall descriptive content and so we maintain that they are not in the business of purporting to represent or describe the world: we are not error theorists. So the challenge for us is to make sense of truth in ethics.

Our fundamental contention, in addressing this challenge, is that truth ascriptions to statements and judgments with moral content are *morally engaged* semantic appraisals—i.e., appraisals in which semantic evaluation is 'fused' with moral evaluation. These truth ascriptions thus are not *descriptive*, because the overall content of the first-order judgments and utterances to which they are applied is not descriptive. That the notion of truth should be employable in a non-descriptive, morally engaged way is to be expected (given our irrealist construal of moral concepts), since ordinary uses of the truth predicate normally operate in accordance with schema T. Since first-order moral judgments and utterances have overall content that is not descriptive, the accompanying truth ascriptions governed by schema T inherit this same feature.

Although the point just made suffices to accommodate semantic assessability within our cognitivist expressivist position, we think there is quite a lot more to say about the notion of truth in general and about its various uses with respect to matters moral. Here we will briefly sketch some further views of ours on these matters that are developed at greater length elsewhere (cf. Horgan 2001 and other papers cited therein; Horgan and Timmons 2002, forthcoming; Timmons 1999 chapter 4). We mention these additional claims partly because we think they speak to various potential concerns concerning our position, and partly because we think they smoothly situate the morally engaged form of truth-ascription within an independently plausible general conception of the how the notion of truth operates.

First, although in many contexts it is appropriate to employ the truth predicate in a morally engaged way in which one's truth ascriptions run in tandem with one's first-order moral beliefs and assertions, there are also contexts in which it instead becomes appropriate to engage in morally *detached* semantic assessment. Under this usage, truth is a matter of correspondence to the world, and falsity is a matter of non-correspondence: a judgment or assertion is true or false if and only if it has objective truth

conditions, and otherwise it lacks truth value. (It is true if the world satisfies those truth conditions—this is correspondence—and it is false if the world fails to satisfy them—non-correspondence.) When the notion of truth is being employed in the morally detached, correspondence-requiring manner, the proper thing to say about moral judgments and assertions is that they are neither true nor false.

Second, first-order moral judgments and assertions, and likewise morally engaged truth ascriptions, are typically *categorical* in content. Although they are made from within a morally engaged stance in which one brings one's own moral standards to bear, they are not implicitly relativized to those standards themselves. Relativism, about first-order moral claims and/or about truth ascriptions to such claims, seriously misconstrues their content. In effect, relativism of this sort conflates morally engaged and morally detached usage, by mistakenly treating engaged usage as a form of detached usage that incorporates implicit relativization to some specific set of moral standards.

Third, we advocate a general approach to concept-world and language-world relations, and to the notion of truth, that we call *contextual semantics*. Some leading ideas are these. (1) Truth is *semantically correct affirmability*, under contextually operative semantic standards; falsity is semantically correct deniability, under such standards. (2) Numerous concepts and terms, including the concept of truth itself, are governed by contextually variable semantic standards of correct affirmability/deniability — where contextual variation can occur not only across different subject matters, but even within thought or discourse about a specific subject matter (e.g., morals). (3) Sometimes the contextually operative semantic standards are *tight*, i.e., these standards conspire with how things are in the world to render a given judgment or statement correctly affirmable or correctly deniable. (4) Thought and discourse governed by semantically tight standards is descriptive. (5) Uses of the truth predicate (and the falsity predicate) governed by semantically tight standards express correspondence (or non-correspondence) to the world. (6) Sometimes the contextually operative semantic standards are not tight. (7) Thought and discourse governed by semantically non-tight standards is non-descriptive in its overall content. (8) Uses of the truth predicate (and the falsity predicate) governed by non-tight standards do not express

correspondence or (or non-correspondence) to the world; when applied to moral thought and discourse, such truth-ascriptions are fused semantic/normative appraisals.[16]

Fourth, given the general framework of contextual semantics, it would be far too crude to say that the detached and the engaged forms of semantic appraisal involve 'different concepts of truth' or 'different meanings of the word "true"'. Contextual variation in operative semantic standards, both at the first-order level and at the level of truth ascription, is a much more subtle matter. The same concept and the same meaning *persist* across such variation despite identity-preserving changes from one context to another, in something like the way a single person persists through time even while undergoing identity-preserving changes.

Logical embedding: accommodating the generic functional role of belief

How is logical embedding to be explained, within cognitivist expressivism? We will now sketch the leading ideas of the account we recommend. We provide further elaboration of technical details, plus further commentary, in the appendix.

Suppose that an agent, Tom, has a logically complex belief—say, the belief that *either Dick is cheerful or Harry ought to apologize*. This belief does not contain an 'embedded is-commitment' or an 'embedded ought-commitment', whatever those might be. In holding this belief, Tom is not is-committed to the content *that Dick is cheerful*, and is not ought-committed to the content *that Harry apologizes*. Rather,

---

[16] Another key thesis of contextual semantics is that tight semantic standards often operate in such a way that a statement can be correctly affirmable under such standards—i.e., true—even though there are no in-the-world objects or properties answering directly to the statement's singular terms, predicates, or unnegated existential quantifications. Such *indirect* correspondence, as we call it, does involve in-the-world truth-makers (and falsity-makers), but the truth-making conditions need not include items in the world answering directly to the given statement's referential apparatus. (For instance, the statement 'Mozart composed 27 piano concertos' can be true even if the objectively correct ontology does not include such items as piano concertos.) Indirect correspondence is very important with respect to the ontological commitments of statements governed by tight semantic standards. But for present purposes it is secondary, because we claim that the semantic standards at work in the engaged use of moral concepts are non-tight; this means that moral statements lack even the *indirect* kind of correspondence (or non-correspondence) to reality. For a treatment of truth and objectivity with substantial similarity to our own contextual semantics, see Wright (1992). Wright, however, seriously flirts with the idea of reducing truth to an idealized form of *epistemically* warranted affirmability that he calls 'superassertibility', whereas we maintain that semantically correct affirmability cannot be reduced to any form of epistemically warranted affirmability however idealized (cf. Horgan 1995, 1996).

the belief state is a *logically complex commitment state*, one that obtains with respect to these core descriptive contents jointly. It is a *disjunctive* commitment state.

Crucial and essential to such a logically complex commitment state, on the view we are here suggesting, is its distinctive *constitutive inferential role* in the cognitive economy of a cognitive agent (insofar as the agent does not exhibit lapses in logical competence). Such a state is one that is poised to interact with other potential beliefs to inferentially generate yet further beliefs that are inferentially "in the offing." For instance, if Tom occurrently has the disjunctive belief about Dick and Harry, and Tom also occurrently has the belief that Dick is not cheerful, these these two beliefs together should inferentially generate (insofar as Tom does not exhibit a rationality failure) the occurrent belief that Harry ought to apologize. (Note that such inferential processes have a phenomenological aspect too; there is *something that it's like* to consciously recognize such logical connections, and there is something that it is like for such inferences to occur in conscious experience.) The constitutive inferential role of logically complex beliefs also includes combining with other beliefs to yield certain *implicit*, logically grounded, further commitment states—even if these remain implicit and perhaps fail to be psychologically operative in the agent.

Constitutive inferential role is a matter of logical consequence relations among various beliefs. So we need to provide a way of construing the logical consequence relation among belief-commitments of the various kinds: is commitments, ought-commitments, and logically complex commitments. We will do so in three steps: first, describing a formal language whose logical syntax overtly models the various types of belief-commitment posited by our account; second, providing formal semantics for sentences of this language, including a definition of the relation of logical consequence; and third, explaining how this formal semantics comports with our above-described treatment of the semantic assessability of moral beliefs and assertions. (This three-step account is sketched in this section and then developed more fully in the appendix.)

Step 1. The formal language we propose employs two affirmatory-force operators, **I[ ]** and **O[ ]**, whose respective analogues in English are the constructions 'It is the case…' and 'It ought to be the

case…' The familiar atomic formulas of predicate logic here are construed not as sentences but rather as *closed non-sentential formulas*. The natural analogues in English are 'that'-clauses. So, for example, letting Pb symbolize 'that Bush is U. S. President', the formal sentence **I[**Pb**]** says *It is the case that Bush is U. S. President*, whereas **O[**Pb**]** says *It ought to be the case that Bush is U.S. President*. Grammatically, the operators **I[   ]** and **O[   ]** are thus sentence-forming (and more generally, sentential *formula* forming) operators.

Logically complex non-sentential formulas, including open ones containing free-variable occurrences, are constructable within the formal language in the standard ways described in predicate logic. Because they all are *non-sentential* formulas, however, the closed ones (i.e., those without free-variable occurrences) are not sentences. A sentence (or open sentential formula) results from application of the syntactic operation of *inserting* a non-sentential formula into the bracketed slot of **I[   ]** or **O[   ]**. For example, letting 'Fx' symbolize 'x is a Fraternity member' and and 'Gx' symbolize 'x takes out the Garbage', the formal version of the sentence *It ought to be the case that some fraternity-house resident takes out the garbage* results from insertion of the closed nonsentential formula (∃x)(Fx & Gx) into the bracketed slot of the operator **O[   ]**, to yield the sentence **O[**(∃x)(Fx & Gx)**]**. Although non-sentential formulas that can be inserted into the bracketed slots of the operators **I[   ]** and **O[   ]** can be arbitrarily complex, they all have *descriptive* content (insofar as the formal language is semantically interpreted). Closed non-sentential formulas have *propositional* descriptive content, and open ones have *non-propositional* descriptive content.

Turn now to logically complex commitment-types, like the disjunctive one involved in Tom's belief that *either Dick is cheerful or Harry ought to apologize*. Within the formal language, such logically complex commitments are explicitly reflected by complex *sentence-forming operators*, which are built from the primitive operators **I[   ]** and **O[   ]** by way of operator-forming connectives and operator-forming quantifiers. For vividness, we use different symbols for these connectives and quantifiers than for the connectives and quantifers that are used to construct logically complex non-sentential formulas; we also use boldface for all these operator-forming connectives and quantifiers. There is a whole hierarchy of

logically complex commitment-types, corresponding to the various logically complex sentence-forming

operators.

For instance, Tom's disjunctive belief that *either Dick is cheerful or Harry ought to apologize* involves

a logically complex commitment-type expressible by the complex, disjunctive, sentence-forming operator

(I[ ] o O[ ]). If the closed non-sentential formulas 'Cd' and 'Ah' respectively symbolize *that Dick is*

cheerful and *that Harry apologizes*, then Tom's disjunctive belief about Dick and Harry is formally

expressible by the sentence that results from respectively inserting these formulas into the operator's

respective slots, thus: **(I[**Cd**] o O[**Ah**])**. Ontologically, Tom's belief is a logically complex commitment-

state of type **(I[ ] o O[ ])**, with respect to the pair of propositional contents *that Dick is cheerful* and *that*

Harry apologizes.

For an example of a complex operator with quantification, suppose that Tom believes that *there is*

a specific fraternity member who ought to take out the garbage. This belief involves a logically complex,

quantificational, commitment-type whose structure is reflected by the sentence-forming operator **(Σ)(I[**

] ∧ O[ ]). A formal sentence expressing Tom's belief results from inserting a variable into the existential-

quantificational slot of this complex operator and inserting open sentential formulas into the operator's

bracketed slots, thus: **(Σx)(I[**Fx**] ∧ O[**Gx**])**. Ontologically, Tom's belief is a logically complex commitment-

state of type

(Σ)(I[ ] ∧ O[ ]), with respect to the pair of non-propositional descriptive contents *that x is a fraternity*

member and *that x takes out the garbage*.

Step 2. In the appendix we employ the idea of a *valuation*, i.e., an assignment of the truth values T

and F to some (but not necessarily all) of the sentences and closed non-sentential formulas of the formal

language, and we then use this idea to define the relation of logical consequence.

Step 3. Nothing in the formal semantics set forth in the appendix requires that the notion of truth

employed in the definition of valuation be understood as operating in the morally detached

'correspondence' manner. On the contrary, insofar as the formal language is construed as an interpreted

language rather than an uninterpreted formal calculus, the truth values assigned to sentences by a given

valuation can perfectly well be those that reflect a given agent's morally *engaged* truth assessments. (These

run in tandem with the agent's morally engaged first-order beliefs, in accordance with schema T.) Thus,

the definition of logical consequence in the appendix can likewise be understood as reflecting logical

relations among an agent's various morally engaged beliefs (including logically complex beliefs), and as

also reflecting the implicit belief-commitments logically generated by these beliefs.

The notion of valuation is defined so as to allow for valuations in which certain sentences and/or

nonsentential formulas are assigned neither T nor F. This is because, insofar as a valuation reflects a given

agent's non-moral and moral beliefs, certain sentences might be ones whose overall content the agent

holds neither true nor false (even when using the truth predicate in a morally engaged way), but instead

is agnostic about.

On the other hand, there will also be a valuation that assigns truth values in accordance with a

morally detached correspondence-usage of 'true', and that furthermore assigns truth and falsity based on

whether or not a given sentence actually corresponds to how things are or not (rather than on the basis of

any particular agent's *beliefs*). In such a valuation, all sentences of the form **O[**A**]** will be assigned neither

T nor F.

In short: In the formal language we have described, there is a hierarchy of sentence-forming

operators with logical structures that explicitly reflect the various kinds of commitment, both simple and

logically complex, that constitute distinct types of belief. The non-sentential formulas, corresponding to

'that'-clauses in natural language, have descriptive content but not assertoric force. A sentence,

constructed by inserting the respective members of a sequence of non-sentential formulas into the

respective bracketed slots of a sentence-forming operator, reflects the ontological structure of the belief it

expresses: the belief is a certain type of commitment-state with respect to a sequence of propositional or

non-propositional descriptive contents—where the commitment-type is expressed by the relevant

sentence-forming operator, and the respective descriptive contents are expressed by the closed and/or

open non-sentential formulas inserted into this operator's respective bracketed slots. The notion of a

valuation—a logically permissible assignment of the truth values T and F—can be defined for this formal

language, and the relation of logical consequence can be defined using this notion. This definition of logical consequence comports well with the contention that truth ascription to sentences whose overall content is non-descriptive is a matter of morally engaged semantic evaluation in which the evaluator's moral and semantic standards are fused. The definition of logical consequence thereby comports well with cognitivist expressivism.

In addressing Task III, we have described belief as *affirmatory commitment* with respect to one or several core descriptive contents. There are two logically fundamental belief-types: is-commitment and ought-commitment. Although an ought-commitment with respect to a core descriptive content is indeed a species of belief, its overall content is non-descriptive; nevertheless, it does have the involuntary, classificatory, coming-down phenomenology that is distinctive of occurrent beliefs. Truth ascription to beliefs and assertions with moral content is a morally engaged fusion of moral and semantic evaluation, and thus comports with the fact that moral content is not descriptive. There are also logically complex belief-types: kinds of affirmatory commitment expressible by logically complex sentence-forming operators in the formal language we have proposed as modeling the ontological structure of beliefs. Beliefs bear logical consequence relations to one another, whether or not they have overall content that is descriptive.[17]

We have mainly been dwelling on matters of moral psychology, arguing that moral judgments are beliefs whose overall content is non-descriptive. Similar points can be made about moral utterances. These are assertoric speech acts that play a distinctive sociolinguistic role—a role in interpersonal dynamics. An assertion, we claim, is best understood as a *stance-taking* speech act, an act through which one overtly comes down on some issue and thereby expresses an is-commitment, or an ought-commitment, or a logically complex commitment with respect to one or more core descriptive contents. In so doing one positions oneself within the context of socio-linguistic interaction, vis-à-vis that core content. A stance is therefore an orientation that one occupies in an interpersonal situation. An ought-stance in particular is typically an action-guiding stance with respect to some core descriptive content.

Some ought-stances are more directly action-guiding than others, but they are all distinctively action-oriented—just as removed ought-beliefs are still still tied to action, although less directly so than direct ought-beliefs. Moreover, just as moral beliefs typically involve a responsiveness to reasons, so engaging in a moral stance-taking speech act normally involves occupying a sociolinguistic role involving the preparedness to *give* reasons for one's moral stance on some issue. For instance, by asserting that Jones ought not to lie to his neighbor, one thereby takes a moral stance in which one signals one's willingness to engage in interpersonal reason-giving with respect to one's own ought-commitment, and to defend one's commitment against objections, or else give up one's commitment.

So sincere utterances of declarative sentences with moral content are speech acts of assertion, even though the overall content of such sentences is not descriptive. Of course, sincere utterances of sentences with descriptive overall content are assertions too. But it should be noted that token descriptive sentences also can be construed as making *belief-independent* assertions (as we will call them)—i.e., assertions that are independent of any specific asserter(s). For instance, an inscription on a subway wall of the sentence 'There are no weapons of mass destruction in Iraq!' asserts—in and of itself, apart from its author(s)—that there are no weapons of mass destruction in Iraq. According to cognitivist expressivism, however, sentence tokens whose overall content is not descriptive cannot rightly be construed as making belief-independent assertions. If an inscription of 'Abortion ought to be outlawed!' is written on a subway wall, then it can be correctly regarded as an assertion only insofar as one interprets it as expressing an ought-commitment of some person(s), known or unknown.

7. More on methodology

Having completed our three main tasks in setting out our metaethical view, let us pause to reflect a bit more on matters of metaethical methodology. Then, following this section, we will briefly consider various additional tasks that a view like ours must eventually tackle.

---

[17] In the appendix we explain how the descriptive/non-descriptive distinction applies to beliefs involving logically complex commitments, and to the sentences expressing such beliefs.

Recall from section 2 that we take a page from Moore in refusing to 'reduce' moral thought and discourse to any other type of discourse. We can now be more explicit in explaining what we are refusing to do and in explaining what we think can be done by way of illuminating such thought and discourse.

First, we refuse to set forth truth conditions for moral beliefs and assertions except, of course, non-substantive ones such as, for example, 'Slavery is wrong' is true if and only if slavery is wrong. The typical expectation in giving some sort of semantically illuminating set of truth conditions for a form of statement is (in effect) to set forth a set of non-trivial substantive truth-makers for such claims. On our view, it is a mistake to suppose that there are substantive truth-makers for moral beliefs and assertions; our view is robustly irrealist in this respect.

Second, we refuse to engage in any sort of paraphrase of moral 'ought' judgments that would effectively 'reduce' them to some other type of judgment, or to some combination of those. We have in mind any sort of 'analysis' that would construe an 'ought' judgment of the form 'S ought to do A' as equivalent in meaning to (say): 'S, do A!' directed toward oneself or another.

But these denials do not mean that we are quietists about the possibility of illuminating the meaning of moral thought and discourse. Like Hare, quoted earlier in section 2, we think that proper illumination comes from understanding the distinctive point and purpose of moral 'ought' judgments in moral thought and discourse as well as their associated phenomenology. So, there is a methodological component to our brand of cognitivist expressivism that we may formulate as follows:

> Moral thought and language does not admit of any kind of semantic 'reduction'; rather it is sui generis. Moreover, a proper understanding of such thought and discourse involves understanding the distinctive phenomenology and associated functional roles of such psychological states and sociolinguistic speech acts.

In articulating the key elements of our view, we have been illustrating our Moorean nonreductive methodology.

8. Remaining tasks — brief progress report

There is a battery of challenges that any view like ours must meet. In the remainder of this paper we can only indicate our current thinking about how to go about meeting these challenges: dealing with them fully will have to be left as remaining tasks. The challenges we have in mind concern: (1) extending our account of direct ought-beliefs to other types of moral belief, (2) moral progress, (3) reasons, (4) moral seriousness, (5) lingering adherence to the descriptivist conception of belief. We now proceed to take them up in order.

Extending the account

In order to fully defend cognitivist expressivism, we need to extend our account of direct ought-beliefs to encompass removed ought-beliefs, beliefs about value (goodness), and moral beliefs that employ 'thick' moral concepts such as benevolence, courage, and malice. For present purposes, we will only consider removed ought-beliefs, again following the lead of Mandelbaum.

According to Mandelbaum, removed judgments of moral obligation, like direct judgments, are a response to one's experiencing some action as 'fitting' vis-a'-vis the circumstances confronting an agent. But as distinct from direct judgments, removed judgments are: (1) made from a third person, spectator's point of view, (2) typically expressed by sentences of the form 'S ought to do/have done A' (where 'S' might refer to one's past self as well as to another person), and (3) are connected with an agent's motivation relatively indirectly and are thus further removed from an agent's motivation to act accordingly.

Phenomenologically, and in contrast to direct judgments, there is obviously no felt reflexive demand, upon the individual engaged in this kind of judging, to perform or not perform the action whose performance or avoidance is judged to be obligatory. However, such judgments often ground attitudes of disinterested approval or disapproval toward the action being evaluated. Here is how Mandelbaum describes this basic contrast:

Removed moral judgments, as we have seen, involve attitudes which may be denominated as "selfless" or "detached"; they are "contemplative" rather than being states of the self. On the other hand, in a direct moral judgment the element of reflexive demand evokes emotion; this emotion, like fear or anger, is experienced as a state of the self and is directly related to action. Thus, the stirredupness and pressure which are present in direct moral judgments have no counterpart in removed moral judgments. In the latter we approve or disapprove, or we may feel admiration, disgust, contempt, or loathing; but even when these stronger affective states are present they appear as by-products of our acts of moral apprehension, and not as direct manifestations of what are experienced to be motivational forces. (127)

So with removed 'ought' judgments, in contrast to direct ought-commitments, the element of felt reflexive demand is absent and thus the motivational role of the judgment differs from direct moral judgments.[18]

So direct and removed ought-beliefs differ in that beliefs of the former type involve a felt reflexive demand and whose role is directly action-guiding, whereas the latter type of belief lacks these characteristics. Further exploration of the distinctive phenomenology and functional role of these types of commitment-state cannot be undertaken here. Nevertheless, both types of ought-commitment state are genuine beliefs—they both involve involuntarily coming down on some matter of moral concern in which they are typically grounded in reasons and thus possess a kind of felt authority.

Moral progress

---

[18] However, this is not to say that removed moral judgments do not have an important tie to motivation. Hare (1952), for instance, maintained that because moral judgments are grounded in reasons and thus commit one to a moral principle which expresses those reasons, in making a removed moral judgment, one is committing oneself to act in a certain way were one to face the circumstances in which the agent being judged is placed. In this way such judgments are at least indirectly action-guiding.

For a metaethical descriptivist-realist, intellectual moral progress is made when one's moral beliefs come to better approximate the moral facts. For a descriptivist-constructivist, intellectual moral progress is made when one's moral beliefs come to better accord with the moral norms that would be accepted by individuals under certain idealized conditions. Indeed, for a cultural moral relativist, there can be individual moral progress in the sense that one's moral beliefs come to better approximate the moral norms (either actual or ideal, depending on the version of relativism) of her culture. But if moral belief and assertion are not to be understood as purporting to describe or represent substantive moral facts, then how can one make sense of genuine moral progress? Furthermore, if there is no metaphysical anchor for moral thought and discourse, then why take it seriously, why not construe moral dispute and discussion as being more like disputes over fashion in clothes and matters of taste generally?

These challenges focus on our irrealist moral metaphysics, and we consider them to be among the most difficult for any robust moral irrealist. Here, then, is an indication of how we would respond to these challenges, though they certainly deserve more attention than it is possible to give them here.

On our view, moral progress is not to be understood as a matter of bringing one's beliefs into closer proximity to some realm of moral facts. Instead, we propose thinking of moral progress as something to be judged from within a committed moral outlook: when one makes judgments about moral improvement, one does so from an engaged moral perspective. So, for example, in judging that one's current view about the morality of euthanasia is better (more correct) than one's former view about this issue, one is not simply registering the fact that one's view on this matter has changed—mere change is not equivalent to progress! Rather, one is making, based on what one experiences as an improved understanding of the morally relevant aspects of euthanasia (including various types of cases this practice covers), a moral judgment—a judgment backed by reasons. In other words, judgments about moral progress are morally engaged themselves, to be understood according to the general metaethical picture we have been sketching.

Reasons

According to our characterization of moral belief and assertion, such commitment-states are typically grounded in what one experiences as reasons for the belief or assertion in question. We are irrealists about moral properties and facts, but what about moral reasons? Is our view committed to realism about moral reasons? If it is, then our overall metaphysical view countenances normative properties and facts—properties and facts having to do with reasons. But then, why be irrealists about normative properties and facts such as goodness and rightness? On the other hand, if we are reasons irrealists, how do we propose to make sense of such claims as, 'The fact that her lie was motivated by pecuniary self-interest is a reason for concluding that her lie was wrong'?

Here, we embrace the spirit if not the letter of C. L. Stevenson's way of dealing with this issue on behalf of his version of noncognitivism. According to Stevenson, claims about some nonmoral fact R constituting a reason for or against some moral judgment are themselves moral claims made from within a morally engaged outlook. Here is how Stevenson put the point:

> So, the general situation is this: when we claim that the factual reason, R, if true, would justify or help to justify the evaluative conclusion, E, we are in effect making another value judgment, E', of our own—the latter serving to evaluate the situation that we shall have if the facts of the case include those that R purports to describe. (1963: 89)

Thus, an inquiry into those nonmoral factual considerations that serve as good reasons for accepting or denying moral statements is what moral thinking is all about. And various normative moral theories have been proposed (versions of consequentialism, deontology, virtue ethics, and so forth) in answer to questions about reasons in ethics. Moral reasons claims, then, express substantive moral beliefs and are themselves to be understood (depending on the type of reasons claim in question) in the general manner set forth in our version of cognitivist expressivism.

Moral seriousness

How, on our view, can make sense be made of the fact that people take their moral views with utter seriousness? If there are no moral facts to which our moral beliefs and assertions must answer, then why care so much about morality?

In response, we think that the sort of challenge being posed is best construed as a *moral* challenge: why *ought* people to take their moral views seriously? And the appropriate response is to give moral reasons—reasons that, for instance, will likely appeal to the important role that morality plays in people's lives. Such moral reasons are not hard to find. And here again, our understanding of the challenge is to take it as appropriately dealt with from within an engaged moral outlook.

Beliefs as descriptive

A final challenge worth mentioning will come from those who continue to embrace the semantic assumption we described in section 1, and who insist that genuine, full-fledged, *beliefs* must be psychological states whose overall content is descriptive.

We have two points to make in response. First, we have offered a battery of *arguments* supporting both the claim that moral judgments are genuine beliefs and the claim that the overall content of a moral judgment is non-descriptive, and we have offered *responses* to various challenges that such a position faces. Given these arguments and responses, there is a substantial burden of proof upon those who wish to insist nonetheless that real beliefs must be descriptive. The fact that the semantic assumption has traditionally been so widely accepted does not alone suffice to justify its acceptance.

Second, even if it should turn out that genuine beliefs really must be descriptive in their overall content, a variant of our cognitivist expressivism would still be available, and would constitute a significantly novel, non-traditional, version of noncognitivism. This variant position, which might be called *quasi-cognitivist expressivism* [19], would deny that ought-commitments are a species of belief, but otherwise it would look very much like our own view. It would embrace the claims (1) that these states are sui generis and irreducible, (2) that they share with beliefs certain key phenomenological features

---

[19] This label was suggested by Uriah Kriegel.

involving involuntary, reason-based, classificatory 'coming down' on an issue, (3) that they are subject to morally engaged semantic assessment, and (4) that there are logical consequence relations among (i) ought-commitments, (ii) beliefs whose overall content is descriptive, and (iii) logically complex commitment-states.

In section 2 we pointed out that traditional forms of noncognitivism are at least *weakly* reductive, by virtue of their reliance on certain non-declarative grammatical constructions (e.g., universal imperatives) as providing a putative model of moral discourse—even though some prominent advocates of such theories (e.g., Hare) repudiated the kind of *strong* reductionism that asserts the outright synonymy or semantic equivalence of moral utterances with such non-declarative utterances. Quasi-cognitivist expressivism, on the other hand, is not reductive even in the weak sense, because it shares with our own position a reliance on certain *declarative* grammatical constructions as providing a model of moral discourse and of the ontological structure of moral judgments—viz., sentences containing the ought-operator **O[ ]**, in the formal language we have described. Eschewing weak reductionism is thus another important respect in which the fallback view we call quasi-cognitivist expressivism is similar to our own position and is different from traditional forms of noncognitivism.

9. Conclusion

We have covered quite a lot of ground, sketching a general phenomenology of belief, developing a general framework of belief (and assertion) that treats some beliefs and assertions as having overall content that is not descriptive, and arguing that there is reason to construe moral beliefs and assertions as nondescriptive. A cognitivist version of expressivism thus emerged. After saying a bit more about our methodology, we concluded by considering a battery of likely challenges often raised against expressivist views in metaethics, and briefly explained how our view answers them. We maintain that our brand of

cognitivist expressivism is superior to its metaethical competitors and is well worth developing further.[20]

[21]

Appendix
We here describe in detail the formal language we propose, with sentences whose syntactic structure models the ontological structures that we claim are possessed by various kinds of belief. We also set forth formal semantics for this language, including a definition of the logical consequence relation. We then add a number of observations by way of commentary.

Syntax:
Primitive Symbols
> Non-boldface: name letters, predicate letters, individual variables, connectives ~, v, &, ⊃, ≡, quantifier symbols ∀ and ∃, parentheses, commas.
> Boldface: operator symbols **I**, **O**, **Bel**, connectives ¬, **o**, ∧, →, ↔, quantifier symbols ∏ and **Σ**, parentheses, brackets.

Non-sentential formulas (nsf's):
> If P is an n-ary predicate letter or a complex n-ary predicate, and $T_1,…,T_n$ are terms (not necessarily distinct) each of which is a name letter or a variable, then $PT_1,…,T_n$ is a non-sentential formula. (Complex n-ary predicates are defined below.)
> If A is a non-sentential formula and X is a variable, then ~A, (∀X)A, and (∃X)A are non-sentential formulas.
> If A and B are non-sentential formulas, then so are (A v B), (A & B), (A ⊃ B), and (A ≡ B).
> Nothing else is a non-sentential formula.
> If F is a non-sentential formula with no free variable-occurrences, then it is **closed**. Otherwise it is **open**.

Sentential-formula forming operators (sff operators):
> **I[ ]** and **O[ ]** are sff operators. (The slots in these operators are *bracketed* slots.)
> If Ω is an sff operator, then **¬Ω**, **(∏)Ω**, and **(Σ)Ω** are sff operators. (The indicated slots are *quantificational* slots.)
> If **Ω** and **Δ** are sff operators, then so are **(Ω o Δ)**, **(Ω ∧ Δ)**, **(Ω → Δ)**, and **(Ω ↔ Δ)**.
> Nothing else is an sff operator.

Sentential formulas:
> If **Ω** is an sff operator containing n quantificational slots and m bracketed slots, and *Q* is a sequence of n variables (not necessarily distinct), and *F* is a sequence of m non-sentential formulas (not necessarily distinct), then **Ω**//(*Q*,*F*) is a sentential formula. (Notation: **Ω**//(*Q*,*F*) is the expression obtained by inserting the respective elements of *Q* into the respective left-to-right quantificational slots in **Ω** and inserting the respective elements of *F* into the respective left-to-right bracketed slots in **Ω**.)
> Nothing else is a sentential formula.
> A sentential formula containing no free variable-occurrences is a **sentence**.

Predicate-forming-operator forming operators (pfo forming operators):
> **Bel[ ]** is a pfo forming operator.
> **Bel( )[ ]** is a pfo forming operator.
> Nothing else is a pfo forming operator.

Predicate-forming operators (pfo's):

  If Ω is a sff operator, then **Bel[Ω]** and **Bel( )[Ω]** are predicate-forming operators.

  Nothing else is a predicate-forming operator.

Complex predicates:

  If Ω//(Q,F) is a sentence, then **Bel[Ω]**//(Q,F) is a complex 1-place predicate.

  If Ω//(Q,F) is a sentential formula containing free occurrences of all and only the distinct
    variables $X_1,…,X_n$, then **Bel($X_1,…,X_n$)[Ω]**//(Q,F) is a complex (n+1)-ary predicate. (The
    operator **Bel( )[ ]** binds the initial occurrences of $X_1,…,X_n$ in **Bel($X_1,…,X_n$)[Ω]**//(Q,F),
    and also binds all occurrences of these variables that are free within Ω//(Q,F).)

  Nothing else is a complex predicate.

Semantics:

A ***non-sentential valuation*** **N** is an assignment of the truth-values T and F to some (but not necessarily
  all) closed non-sentential formulas, in conformity with the following conditions.

  No closed non-sentential formula is assigned both T and F by **N**.

  If A is a closed non-sentential formula, then

    N assigns T to A iff **N** assigns F to ~A.

    N assigns F to A iff **N** assigns T to ~A. (So **N** assigns neither T nor F to A iff **N** assigns
      neither T nor F to ~A.)

  If A and B are closed non-sentential formulas, then

    N assigns T to (A v B) iff either **N** assigns T to A or **N** assigns T to B.

    N assigns F to (A v B) iff **N** assigns F to both A and B.

    N assigns T to (A & B) iff V **N** assigns T to both A and B.

    N assigns F to (A & B) iff either **N** assigns F to A or **N** assigns F to B.

    N assigns T to (A ⊃ B) iff either **N** assigns F to A or **N** assigns T to B.

    N assigns F to (A ⊃ B) iff **N** assigns T to A and **N** assigns F to B.

    N assigns T to (A ≡ B) iff either **N** assigns T to both A and B or **N** assigns F to both A and
      B.

    N assigns F to (A ≡ B) iff either **N** assigns T to A and F to B, or **N** assigns F to A and T to
      B.

  If A is a non-sentential formula in which the only free variable-occurrences are occurrences of X,
    then

    N assigns T to ($\forall$X)A iff for every name letter N, **N** assigns T to A(X/N). (Notation:
      A(X/N) is the result of replacing every free occurrence of X in A by N.)

    N assigns F to ($\forall$X)A iff for some name letter N, **N** assigns F to A(X/N).

    N assigns T to ($\exists$X)A iff for some name letter N, **N** assigns T to A(X/N).

    N assigns F to ($\exists$X)A iff for every name letter N, **N** assigns F to A(X/N).

A non-sentential formula A is a *non-sentential consequence* of a set of non-sentential formulas {$B_1,…,B_n$} iff
  (1) every non-sentential valuation that assigns T to each of $B_1,…B_n$ also assigns T to A, and (2)
  every non-sentential valuation that does not assign F to any of $B_1,…B_n$ also does not assign F to A.

A *valuation* **V** is a pair < **N**,**S**> such that **N** (the *non-sentential element* of **V**) is a non-sentential valuation
  and **S** (the *sentential element* of **V**) is an assignment of truth values to some (but not necessarily
  all) sentences in conformity with the following conditions:

  No sentence is assigned both T and F by **S**.

  If A is a closed non-sentential formula, then

    S assigns T to the sentence **I[A]** iff **N** assigns T to A.

    S assigns F to the sentence **I[A]** iff **N** assigns F to A.

S assigns T to the sentence **O[A]** only if for every closed non-sentential formula B that is a non-sentential consequence of A, **S** assigns T to **O[B]**.

S assigns F to the sentence **O[A]** only if for every closed non-sentential formula B such that A is a non-sentential consequence of B, **S** assigns F to **O[B]**.

If $\Omega//(Q,F)$ is a sentence, then

S assigns T to $\neg\Omega//(Q,F)$ iff **S** assigns F to $\Omega//(Q,F)$.

S assigns F to $\neg\Omega//(Q,F)$ iff **S** assigns T to $\Omega//(Q,F)$.

If $\Omega//(Q_1,F_1)$ and $\Delta//(Q_2,F_2)$ are sentences, then

S assigns T to $(\Omega \text{ o } \Delta)//(Q_1,F_1,Q_2,F_2)$ iff either **S** assigns T to $\Omega//(Q_1,F_1)$ or **S** assigns T to $\Delta//(Q_2,F_2)$. (Notation: $(\Omega \text{ o } \Delta)//(Q_1,F_1,Q_2,F_2)$ is the expression obtained by inserting the respective elements of Q_1 and F_1 into the respective quantificational and bracketed slots of the Ω segment of $(\Omega \text{ o } \Delta)$, and likewise inserting the respective elements of Q_2 and F_2 into the slots of the Δ segment of $(\Omega \text{ o } \Delta)$.)

S assigns F to $(\Omega \text{ o } \Delta)//(Q_1,F_1,Q_2,F_2)$ iff **S** assigns F to both $\Omega//(Q_1,F_1)$ and $\Delta//(Q_2,F_2)$.

S assigns T to $(\Omega \wedge \Delta)//(Q_1,F_1,Q_2,F_2)$ iff **S** assigns T to both $\Omega//(Q_1,F_1,)$ and $\Delta//(Q_2,F_2)$.

S assigns F to $(\Omega \wedge \Delta)//(Q_1,F_1,Q_2,F_2)$ iff either **S** assigns F to $\Omega//(Q_1,F_1)$ or **S** assigns F to $\Delta//Q_2,F_2)$.

S assigns T to $(\Omega \rightarrow \Delta)//(Q_1,F_1,Q_2,F_2)$ either **S** assigns F to $\Omega//(Q_1,F_1)$ or **S** assigns T to $\Delta//Q_2,F_2)$.

S assigns F to $(\Omega \rightarrow \Delta)//(Q_1,F_1,Q_2,F_2)$ iff **S** assigns T to $\Omega //(Q_1,F_1)$ and **S** assigns F to $\Delta//(Q_2,F_2)$.

S assigns T to $(\Omega \leftrightarrow \Delta)//(Q_1,F_1,Q_2,F_2)$ iff either **S** assigns T to both $\Omega //(Q_1,F_1)$ and $\Delta//(Q_2,F_2)$ or **S** assigns F to both $\Omega //(Q_1,F_1)$ and $\Delta//(Q_2,F_2)$.

S assigns F to $(\Omega \leftrightarrow \Delta)//(Q_1,F_1,Q_2,F_2)$ iff either **S** assigns T to $\Omega //(Q_1,F_1)$ and F to $\Delta//(Q_2,F_2)$. or **S** assigns F to $\Omega //(Q_1,F_1)$ and T to $\Delta//(Q_2,F_2)$.

If X is a variable and $\Omega//(Q,F)$ is a sentential formula in which the only free variable-occurrences are occurrences of X, then

S assigns T to $(\prod X)\Omega//(Q,F)$ iff for each name letter N, **S** assigns T to $\Omega//(Q,F)///(X/N)$. (Notation: $\Omega//(Q,F)///(X/N)$ is the expression obtained from $\Omega//(Q,F)$ by replacing all free occurrences of the variable X by the name letter N.)

S assigns F to $(\prod X)\Omega//(Q,F)$ iff for some name letter N, **S** assigns F to $\Omega//(Q,F)///(X/N)$.

S assigns T to $(\Sigma X)\Omega//(Q,F)$ iff for some name letter N, **S** assigns T to $\Omega//(Q,F)///(X/N)$.

S assigns F to $(\Sigma X)\Omega//(Q,F)$ iff for each name letter N, **S** assigns F to $\Omega//(Q,F)///(X/N)$.

A valuation **V** *assigns* T (or F) to a sentence or non-sentential formula Φ iff either the sentential element or the non-sentential element of **V** assigns T (or F) to Φ.

A valuation is *complete* iff it assigns T or F to every closed non-sentential formula and every sentence. Otherwise it is *partial*.

A sentence or closed non-sentential formula *A* is a *logical consequence* of a set of sentences or closed non-sentential formulas $\{B_1,\ldots,B_n\}$ iff (1) every valuation that assigns T to each of $B_1,\ldots,B_n$ also assigns T to *A*, and (2) every valuation that does not assign F to any of $B_1,\ldots B_n$ also does not assign F to *A*.

Commentary:

1.	We can now state how the descriptive/non-descriptive distinction applies to sentences generally, including sentences constructed by insertion of non-sentential formulas (nsf's) into logically complex sentence-forming operators. A sentence *A* has descriptive overall content just in case there is a partial valuation **V** such that (1) for every closed nsf B, **V** assigns neither T nor F to **O[B]**, and (2) **V** assigns T or F to *A*. The idea is that there's some valuation that assigns a truth value to the sentence *A* while also assigning no truth value to any ought-sentence.

2.	The semantics involves a substitutional rather than objectual treatment of the quantifiers. This is for simplicity, but one could instead formulate the semantics model-theoretically with an objectual construal of the quantifiers, rather than the truth-theoretic way with substituational quantifiers.

3.	Truth values are assigned by a valuation not only to sentences but also to closed nsf's, because the latter too have propositional descriptive content and hence can be true or false. Likewise in English, a that-clause can be true or false.

4.	As explained already, some sentences can be assigned neither T nor F by a valuation. Falsity conditions thus needed to be built into the notion of valuation, rather than a stipulation that F is assigned to any sentence or closed nsf not assigned T.

5.	We depart from familiar approaches to formal semantics, which construe a valuation as an assignment of truth values just to the logically simplest sentences and then define 'truth in a valuation' recursively. This is because we take it that not all sentences have their truth values uniquely determined by the truth values assigned by a valuation to logically simpler ones. For instance, suppose that for every name symbol N, a valuation **V** assigns F to every sentence **O[GN]**, where 'Gx' symbolizes 'that x takes out the garbage'. Such a valuation might yet assign either T or F to the sentence **O[(∃x)Gx]**. Even though no particular individual is such that *he/she* ought to take out the garbage, perhaps nonetheless it ought to be the case that *somebody* takes out the garbage; or perhaps not. Logically, both possibilities remain open.

6.	The fundamental semantic principles governing the operator **I[ ]** are the ones saying that a valuation assigns T (or F) to **I[A]** iff it assigns T (or F) to A. The idea is that a given descriptive content should be assigned the same truth value by a valuation as is assigned to the sentence that makes an is-the-case assertion with respect to that content. As we pointed out in addressing semantic assessability in section 6, a valuation can serve either of two roles: first, reflecting an agent's beliefs, and second, describing the truth values the sentences possess when they are construed as belief-independent assertions. Consider these two construals of a valuation, in turn. First, if an agent is is-committed with respect to the content expressed by A, i.e., has a belief of type **I[ ]** with respect to A, then a valuation reflecting the agent's beliefs will assign T both to A itself and to **I[A]**. Second, if sentences of the type **I[ ]** are being construed as making *belief-independent assertions* (rather than belief-expressive assertions), then again a sentence **I[A]** should be assigned T by **V** just in case **V** also assigns T to A itself. If a certain descriptive content is true under a valuation, then so is the sentence that is-asserts that content; and conversely.

7.	The fundamental semantic principles governing the operator **O[ ]** are the ones saying that **V** assigns T to **O[A]** only if for every B that is a non-sentential consequence of A, **V** assigns T to **O[B]**; and that **V** assigns F to **O[A]** only if for every B such that A is a non-sentential consequence of B, **V** assigns F to **O[B]**. It is because these are *if/then* constraints, rather than biconditionals, that the truth values assigned by a valuation to statements of the type **O[A]** are not uniquely determined by the truth values assigned to logically simpler statements. For instance, if for every N, **V** assigns F to **O[FN]**, then **V** may assign either T or F to **O[(∃x)Fx]**. (See the example in comment 5 above.)

8.	The fundamental semantic principles governing the operator **O[ ]** are quite weak. One could consider strengthening them in various ways, even while retaining their *if/then* form. But we doubt that there is adequate theoretical motivation for doing so. Also, a general reason to avoid stronger principles is the need to avoid various well known deontic paradoxes. Consider, for instance, this candidate semantic principle: If **V** assigns T to both **O[A]** and **O[(A ⊃ B)]**, then **V** assigns T to **O[B]**. Building this closure principle into the formal semantics would allow the generation of the "contrary to duty imperative paradox." Suppose, for example, that Andy steals the money, that he ought to be punished for

doing so, but that that he is otherwise undeserving of punishment. Letting 'Sa' and 'Pa' respectively symbolize 'that Andy steals the money' and 'that Andy is punished', these four claims are all true: **O[**Pa**]**, **O[**~Sa**]**, **(O[**(~Sa ⊃ ~Pa)**]**, **(O[**~Pa**] → ¬O[**Pa**])**. Given the closure principle under consideration, these claims would jointly entail the contradiction **(O[**Pa**] ∧ ¬O[**Pa**])**.

9.   Iteration of the operator symbols **I** and **O** and is not permitted in this formal language, as we have specified its syntax. Allowing iteration of **I** would be entirely pointless, as far as we can tell. If desired, however, one could permit iteration of **O** by modifying the syntax to say that if **Ω** is an sff operator then so is **OΩ**. (One might also modify the formal semantics too, to lay down certain constraints on the how a valuation's assignment of T or F to a sentence of type **OΩ** must be related to the truth values the valuation assigns to various other sentences. Alternatively, one might not build in any such constraints; perhaps logic alone—or anyway, *non-modal* logic alone—does not impose any formal constraints upon the iteration of 'ought'.) Although it is not obvious that moral thought and moral discourse really need iterated 'ought' constructions, a prima facie case can be made for this claim. Suppose, for instance, that Andy has stolen the money, and let 'Sa' and 'Pa' respectively symbolize 'that Andy steals the money' and 'that Andy is punished'. Arguably, although **O[**Pa**]** is true, **OO[**~Pa**]** is also true (because **O[**~Sa & ~Pa**]** is true).

10.   Suppose that a sentence is constructed (by insertion of non-sentential formulas, plus perhaps insertion of variables and prefixing of quantifiers) from a *logically complex* sentential-formula forming operator, rather than being constructed directly from either of the logically simple sff's **I[ ]** or **O[ ]**. The assertoric force of such a sentence is borne by the *entire logically complex sff operator*, and not by its *proper constituents*. Likewise, for a belief expressed by such a sentence, the belief's affirmatory force is borne by the entire logically complex commitment-type of which the belief is an instance. For example, letting 'Cd' and 'Ah' respectively symbolize 'that Dick is cheerful' and 'that Harry apologizes', the assertoric force of the disjunctive sentence **(I[**Cd**] o O[**Ah**])** is borne by the entire operator **(I[ ] o O[ ])** from which this sentence is constructed, and not by either of this operator's proper constituents **I[ ]** or **O[ ]**. Likewise, the affirmatory force of the *belief* that either Dick is cheerful or Harry ought to apologize is borne by the entire logically complex commitment-type of which this belief is an instance—a commitment-type expressed by the complex operator **(I[ ] o O[ ])**.

11.   In the definition of a valuation, the clauses governing logically complex sentential formulas work semantically exactly the same way they would work if these sentential formulas were constructed syntactically in the more familiar manner—viz., directly from logically simpler sentential formulas via connectives and quantifiers (rather than by insertion of nsf's into the slots of logically complex operators). But it bears reiteration that the purpose of our own syntactic construction rules is this: the syntactic structure of the sentences thereby reflects the ontological structure (on our account) of the relevant beliefs expressible by these sentences. The key point to appreciate is that formal semantics does not *mandate* construing logically complex sentences as being built up syntactically in ways familiar from standard symbolic logic; on the contrary, the logical syntax can perfectly well be as we have specified.

12.   The formal semantics is naturally construed as yielding two complementary classes of logical relations, applicable to two distinct domains of relata. First is the domain of *propositions* and *non-sentential formulas*. A proposition can be construed as a *set of ways the world might be*, and is expressible linguistically by a closed non-sentential formula. (Such a formula *expresses* the proposition without *asserting* it.) Propositions, and the closed non-sentential formulas that are their linguistic vehicles, bear logical relations to one another characterizable in terms of the notion of a non-sentential valuation—notably, the relation of non-sentential consequence. Second is the domain of *beliefs*, *assertions*, and *sentences*—with sentences being the linguistic vehicles for expressing beliefs and for making assertions. Beliefs, assertions, and sentences bear logical relations to one another (notably the relation of sentential consequence) characterizable in terms of the notion of valuation. Not all beliefs and assertions are ones whose overall content is *propositional* content; this is reflected in the fact that the notion of valuation is broader than that of sentential valuation.

13.   Consider beliefs of logically complex type—i.e., beliefs other than those belonging to the two logically simple commitment-types, is-commitments and ought-commitments. For any belief B of

logically complex type: if B is descriptive (and thus is expressible by a descriptive sentence —cf. comment 1), then B is logically equivalent to some belief B* that is a logically simple is-commitment (and thus is expressible by a descriptive sentence constructed by inserting some closed non-sentential formula into a single occurrence of the operator '**I[   ]**'). Take, for instance, the belief that *either it is the case that Dick is cheerful or it is the case that Harry apologizes*, expressible symbolically via the descriptive sentence **(I[**Cd**] o I[**Ah**])**. This belief is logically equivalent to the belief that *it is the case that either Dick is cheerful of Harry apologizes*, expressible symbolically via the descriptive sentence **I[**(Cd v Ah)**]**. The former belief is a commitment of logically complex type **(I[   ] o I[   ])** with respect to the pair of descriptive contents *that Dick is cheerful* and *that Harry apologizes*, whereas the latter belief is a logically simple is-commitment with respect to the logically complex descriptive content *that either Dick is cheerful or Harry apologizes*.

14.   For any belief B of logically complex type: if B is *not* descriptive (i.e., is not expressible by a descriptive sentence), then B is *not* logically equivalent to any is-commitment. Take, for example, the belief that *either it is the case that Dick is cheerful or it ought to be the case that Harry apologizes*, expressible symbolically via the sentence **(I[**Cd**] v O[**Ah**])**. This belief, which is a commitment of logically complex type **(I[   ] v O[   ])** with respect to the pair of descriptive contents expressed respectively by the non-sentential closed formulas 'Cd' and 'Ah', is not logically equivalent to any is-commitment whatever.

15.   The asymmetry revealed in points 13 and 14 is to be expected, given our irrealism about the ontology of morals: an is-commitment (or an ought-commitment) can only obtain with respect to a *way the world might be*, and according to our moral irrealism, only *descriptive* that-clauses express genuine ways the world might be. (All closed non-sentential formulas of the formal language correspond to *descriptive* that-clauses of natural language.)

16.   The formal language includes belief-attributing predicates. Once again, we have specified the syntax so that these constructions too reflect, in their syntactic structure, the ontological structure of the beliefs they attribute. For each kind of commitment-state that is a type of belief, including the logically complex ones, there is a predicate-forming operator whose structure mirrors that belief-type. A belief-attributing predicate expressing a specific belief is constructed by inserting, into the operator's bracketed slots, nsf's expressing the respective descriptive contents that are the descriptive-content relata of the given belief. For instance, Tom's belief that *either Dick is cheerful or Harry ought to apologize* involves a commitment-type expressed by the predicate-forming operator **Bel[(I[   ] o O[   ])]**. Letting 'Cd' and 'Ah' respectively symbolize 'that Dick is cheerful' and 'that Harry apologizes', a formal predicate expressing Tom's belief-property is obtained by inserting into the operator's respective bracketed slots the closed sentential formulas expressing the two propositional-content relata of the belief, thus: **Bel[(I[**Cd**] o O[**Ah**])]**. So, letting 't' refer to Tom, the logically complex non-sentential formula *that Tom believes that either Dick is cheerful or Harry ought to apologize* is symbolized this way: **Bel[(I[**Cd**] o O[**Ah**])]**t. Likewise, Tom's belief that that *there is a specific fraternity member who ought to take out the garbage* involves a commitment-type expressed by the predicate-forming operator

Bel[(Σ)(I[   ] ∧ O[   ])]. A formal predicate expressing the belief-property is obtained by inserting a variable into the operator's quantificational slot and inserting into its bracketed slots open sentential formulas expressing the two non-propositional descriptive-content relata of the belief: **Bel[(Σx)(I[**Fx**] ∧ O[**Gx**])]**. Thus, the non-sentential formula *that Tom believes that there is a specific fraternity member who ought to take ought garbage* becomes: **Bel[(Σx)(I[**Fx**] ∧ O[**Gx**])]**t.

17.   The point of the belief-attributing predicates employing the variable-binding operator **Bel(   )[   ]** is to accommodate 'de re' belief constructions, in addition to de dicto ones. So, letting 'r' and 'o' refer to Ralph and Ortcutt respectively and letting 'S' symbolize 'is a spy', the de dicto belief predicate 'believes that Ortcutt is a spy' is symbolized as **Bel[I[**So**]]**, and the de dicto non-sentential formula *that Ralph believes that Ortcutt is a spy* is symbolized as **Bel[I[**So**]]**r—whereas the de re belief predicate 'believes of Ortcutt that he is a spy' is symbolized as **Bel(x)[I[**Sx**]]**, and the de re non-sentential formula *that Ralph believes of Ortcutt that he is a spy* is symbolized as **Bel(x)[I[**Sx**]]**r,o.

18.   Our treatment of the logical syntax of de dicto and de re belief predicates is adapted from a proposal of Quine (1960, 1970), further discussed and developed in Horgan (1989). Quine, though, construed his approach as a way of dispensing with so-called "propositional objects" of belief, whereas here we are instead construing beliefs as relations linking a believer to one or more descriptive contents—

viz., the contents expressed by the non-sentential-formulas that get inserted into the bracketed slots of the complex predicate-forming operator expressing a particular belief-type. Thus, Ralph's de dicto belief is a relation between Ralph and the propositional descriptive content expressed by 'that Ortcutt is a spy'; and Ralph's de re belief is a relation linking Ralph both to Ortcutt and to the non-propositional descriptive content expressed by 'that x is a spy'.

References

Audi, R. 2004. *The Good in the Right: A Theory of Intuition and Intrinsic Value*. Princeton, N. J.: Princeton University Press.

Ayer, A. J. 1936. *Language, Truth Logic*. 2nd edition. New York: Dover.

Blackburn, S. 1984. *Spreading the Word*. Oxford: Oxford University Press.

Blackburn, S. 1996. 'Securing the Nots'. In W. Sinnott-Armstrong & M. Timmons (Eds.), *Moral Knowledge? New Readings in Moral Epistemology*. New York & Oxford: Oxford University Press.

Blackburn, S. 1998. *Ruling Passions*. Oxford: Oxford University Press.

Carnap, R. 1935. *Philosophy and Logical Syntax*. London: Kegan Paul, Trench, Trubner & Co.

Dewey, J. 1922. *Human Nature and Conduct.* Carbondale & Edwardsville: Southern Illinois University Press.

Dreyfus, H. & S. Dreyfus. 1990. 'What is Morality: A Phenomenological Account of the Development of Ethical Expertise'. In D. Rassmussen (Ed.), *Universalism vs. Communitarianism*. Cambridge, MA: MIT Press.

Gibbard, Allan. 1990. *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.

Hare, R. M. 1952. *The Language of Morals*. Oxford: Oxford University Press.

Harman, G. 1977. *The Nature of Morality*. Oxford: Oxford University Press.

Horgan, T. 1989. 'Attitudinatives'. *Linguistics and Philosophy* 12: 133-65.

Horgan, T. 1995. 'Critical Study: Crispin Wright, *Truth and Objectivity*'. *Nous* 29: 127-38.

Horgan, T. 1996. 'The Perils of Epistemic Reductionism'. *Philosophy and Phenomenological Research* 56: 891-97.

Horgan, T. 2001. 'Contextual Semantics and Metaphysical Realism: Truth as Indirect Correspondence'. In M. Lynch (Ed.), *The Nature of Truth: Classic and Contemporary Perspectives*. Cambridge, MA: MIT Press.

Horgan, T. & M. Timmons. 2000a. 'Nondescriptivist Cognitivism: Framework for a New Metaethic'. *Philosophical Papers* 29: 121-53.

Horgan, T. & M. Timmons. 2000b. 'Copping Out on Moral Twin Earth'. *Synthese* 124: 139-152.

Horgan, T. & M. Timmons. 2002. 'Conceptual Relativity and Metaphysical Realism'. *Philosophical Issues* 12: 74-96.

Horgan, T. & M. Timmons. 2004. 'Morality without Moral Facts'. In J. Dreier (Ed.), *Contemporary Debates in Moral Theory*. Oxford: Blackwell Publishers.

Horgan, T. & Timmons, M. Forthcoming. 'Expressivism, Yes! Relativism, No!'.

Hume, D. 1739. *A Treatise of Human Nature*. Oxford: Oxford University Press.

Mandelbaum, M. 1955. *The Phenomenology of Moral Experience*. Glenco, IL: The Free Press.

Moore, G. E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.

Quine, W. V. 1960. *Word and Object*. Cambridge, MA: MIT Press.

Quine, W.V. 1970. *Philosophy of Logic*. Englewood Cliffs, N.J.: Prentice-Hall.

Shafer-Landau, R. 2003. *Moral Realism: A Defense*. Oxford: Oxford University Press.

Stevenson, C. L. 1963. *Facts and Values*. New Haven & London: Yale University Press.

Timmons, M. 1999. *Morality without Foundations: A Defense of Ethical Contextualism*. New York: Oxford University Press.

Varela, F. 1992. *Ethical Know-How*. Stanford: Stanford University Press.

Warnock, M. 1960. *Ethics After 1900*. London: Oxford University Press.

Wright, D. 1992. *Truth and Objectivity*. Cambridge, MA: Harvard University Press.