

Penultimate version. See the official version in *Oxford Studies in Metaethics*, Vol. 8, Ed. Russ Shafer-Landau, Oxford University Press, 2013, pp. 195–222.

# Truth Conditions and the Meanings of Ethical Terms<sup>\*</sup>

Alex Silk

a.silk@bham.ac.uk

**Abstract** This paper motivates and develops what I call a *condition semantics* for moral terms. According to condition semantics, moral sentences conventionally distinguish among moral standards (or test whether a moral standard meets a certain condition) just as ordinary factual sentences conventionally distinguish among possible worlds (or test whether a possible world meets a certain condition). This point is captured formally within an extension of a familiar truth-conditional paradigm. The resulting analysis improves upon its main competitors: invariantism and contextualism. The framework of condition semantics also offers a perspicuous way of posing various classical ethical and metaethical questions — e.g., concerning relativism, expressivism, and judgment internalism. This can motivate clearer, better motivated answers and suggest new ways the dialectic may proceed.

**Keywords:** Condition semantics • Invariantism • Contextualism • Relativism • Expressivism  
• Semantic competence • Truth-conditions • Attitude ascriptions

---

<sup>\*</sup>For helpful discussion and comments, thanks to Matthew Chrisman, Jan Dowell, Billy Dunaway, Andy Egan, Dmitri Gallow, Dan Greco, Tristram McPherson, Sarah Moss, David Plunkett, Peter Railton, Dan Singer, Eric Swanson, and audiences at the University of Michigan Graduate Student Working Group and the Eighth Annual Wisconsin Metaethics Workshop. Thanks especially to Allan Gibbard for extensive discussion and comments on previous drafts.

# 1 Truth conditions in ethics and metaethics

Other things being equal, it would be nice to have a semantics that differentiates the truth conditions of (1) and (2).

- (1) Torturing babies for fun is wrong.
- (2) No more than four colors are needed to color a map in a plane so that no adjacent regions share the same color. (the Four Color Theorem)

On the face of it, (1) and (2) carry different information; as you know, (1) and (2) are about totally different things. Similarly, it is not counterintuitive that semantic competence with (1) requires a certain non-trivial ability; not just anyone counts as understanding (1).

Unfortunately, it is hard to see how we can capture these two intuitions if certain familiar positions about truth-conditional content and moral truths are right.

## TRUTH-CONDITIONAL CONTENT

The semantic content of a declarative sentence is, or at least determines, its truth conditions (represented as a set of possibilities).

## NECESSARY MORAL TRUTHS

If the proposition expressed by a moral sentence like (1) is true (false) at all, it is true (false) necessarily — i.e., true (false) at all possible worlds.

## MORAL INVARIANTISM

Moral sentences semantically express (or determine) ordinary possible worlds propositions and are interpreted with respect to an invariant moral view or standard, namely, the correct one.

It follows from these claims that a moral sentence like (1) has the set of all possible worlds as its truth conditions. Assuming that mathematical truths like (2) are necessary truths, the truth conditions of (1) and (2) will be equivalent. Similarly, the above positions seem to imply that semantic competence — of the sort that requires an ability to correctly distinguish among possibilities — with moral sentences requires no more than an ability to “divide” the space of possibilities into  $W$  and  $\emptyset$ . (More on this requirement for semantic competence below.) I am not endorsing these familiar positions just yet. The point here is simply that these ways of thinking which we may be drawn to adopt — and which often *are* adopted — lead to a puzzle. We will be revising them in due course. (For concreteness I will tend to couch

our discussion in terms of specifically moral language, though much of it applies to normative language more generally.)

This puzzle is, of course, a familiar one from possible worlds semantics, often known as *the problem of logical omniscience*. (By calling it a “problem” I am not suggesting that progress has not been made on it.) But it seems particularly problematic in the moral case. If the problem of how (2) and ‘ $2 + 2 = 4$ ’ say the same thing is a real problem, the problem of how (1) and (2) say the same thing is even worse.

A number of options are available in reply. First, one might deny that there is a puzzle: “Sure,” one might say, “(1) and (2) have the same truth conditions, but this isn’t a problem since they express different propositions.” One might cash out this sort of reply in terms of structured propositions, impossible worlds, or a model theory that treats propositions as primitives. Or one might employ a Stalnakerian diagonalization strategy and say that though (1) and (2) semantically express the same proposition, and have the same truth conditions, we reinterpret them as having relevantly different asserted contents in conversation. For the moment I want to put these responses to the side.<sup>1</sup> It will be instructive to examine whether we can find a theory that differentiates the truth conditions of sentences like (1) and (2). So let’s assume as a working hypothesis there *is* a puzzle — that, other things equal, it would be preferable to distinguish the truth conditions of (1) and (2).

Granting that there is a puzzle, one might respond by giving up NECESSARY MORAL TRUTHS. One might do so by denying that moral sentences are true or false. But this is a high cost; even contemporary heirs of emotivism grant that moral sentences have truth values. Instead one might say that though moral sentences like (1) have truth values, their truth values vary across possible worlds. But this is also a tough bullet to bite. It seems eminently plausible that no matter how the world might be or might have been, if anything is right or wrong at all, torturing babies for fun couldn’t not have been wrong. At least some moral truths seem to be necessary.<sup>2</sup> (Moving to epistemically possible worlds won’t help: For all I know, no matter how the world could have been, torturing babies for fun couldn’t not have been wrong. The reader may feel free to substitute her favorite candidate necessary moral truth.)

Alternatively, one might respond by denying TRUTH-CONDITIONAL CONTENT. But this is also a tough bullet to bite. Even if propositions need to be individuated

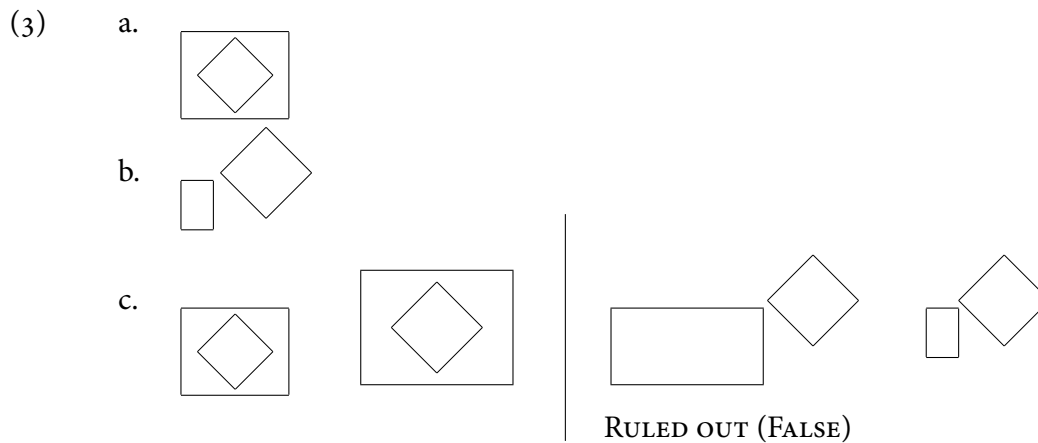
---

<sup>1</sup>For one thing, there is as yet no consistent theory of structured propositions (for recent discussions see CROSS 2001a,b, DEUTSCH 2008). For worries with the propositions-as-primitives and impossible worlds approaches, see, e.g., CRESSWELL 2002. I return to diagonalization and the impossible worlds approach in §2.

<sup>2</sup>*Cf.*, e.g., MOORE 1903, p. 9; GIBBARD 2003, p. 30; PARFIT 2011a, p. 129, 2011b, pp. 307, 489–490, 517–518.

more finely than by their truth conditions, it is uncontroversial that propositions at least determine truth conditions. This is all we need to get our puzzle going.

With that said, a brief motivation of the picture of content I have in mind may be in order. I am sympathetic with the standard view of (semantic and mental) content according to which contents function to *distinguish among possibilities*. Semantic competence with a sentence — knowing a sentence’s conventional meaning (or at least its conventionally proffered content) — requires having a certain capacity: a capacity to correctly divide the space of possibilities, to correctly rule out a certain set of possibilities from some domain. Equivalently, it requires having the ability, when given a possibility or circumstance, to tell whether the sentence is true or false in it.<sup>3</sup> (Even if this is not the only requirement for semantic competence, it is surely an important one.) For a simplistic example, being semantically competent with ‘The diamond is inside the box’ requires having the following sort of capacity: a capacity to know that it is true if shown a picture like (3a); to know that it is false if shown a picture like (3b); and to correctly rule out a certain subset of situations from a domain, as in (3c).



<sup>3</sup>See, e.g., LEWIS 1970, CRESSWELL 1978, STALNAKER 1984, 1999, among many others. I do not speak to Davidsonian versions of truth-conditional semantics here. In what follows I will often speak of “knowing what an expression means,” or “understanding an expression,” or “being semantically competent with an expression.” If one wishes to eschew complications from “knowing what/how” and “understanding” talk, one may construe my discussion in terms of one important component of semantic competence. Our question is what the conventional contents of moral expressions are, at the relevant level of abstraction, where these contents serve as a rational reconstruction or model of speakers’ dispositions to interpret moral language; they represent one important component of what is grasped by language users, under certain idealizations, who are semantically competent with this fragment of the language. (By ‘conventional’ I mean, to a first approximation, “not calculable from the Cooperative Principle and the conversational maxims.”) Thanks to Sarah Moss for discussion.

It may turn out that the problems facing truth-conditional pictures of content are too great. For now, though, let's take it as our starting point, for it will help bring into relief important issues about the meanings of moral terms, issues which are independent of the possible worlds framework. We will end up revising the standard picture anyhow. So for the moment let's suppose that propositions determine truth conditions, which are represented by a set of possibilities, and let's see where this takes us. (All I need is this determination relation, though for simplicity I'll often identify propositions with their truth conditions.) If one likes, one can treat my inquiry as a conditional: *If it's right to think of meaning and content as determining truth conditions in roughly this way, what implications does this have for our understanding of the meanings of ethical terms?* As we will see in §3, approaching the semantics of ethical terms from this perspective will suggest a new, illuminating way of understanding various classic metaethical questions and positions. But I'm getting ahead of myself.

This leaves us with MORAL INVARIANTISM. On the face of it, MORAL INVARIANTISM seems attractive. By treating the moral standard relevant for interpretation as determined by and invariant across possible worlds, it seems to capture the objective purport of moral language, or how the “moral order... is... part of the fundamental nature of the universe (and... of any possible universe in which there were moral agents at all)” (Ross 1930, p. 18). And by treating moral sentences as expressing ordinary possible worlds propositions, it can make use of the powerful resources of truth-conditional semantics. Nevertheless I think we should reject MORAL INVARIANTISM.

There are good reasons to think we should not build substantive normative assumptions about what the correct moral theory is into the meanings of moral terms. For example, suppose classical utilitarianism is correct, and (1) is true iff torturing babies for fun does not maximize overall pleasure. Even if she is fully informed about all the non-normative facts, the non-utilitarian could coherently disagree with these “truth conditions” while knowing what (1) means. Instead she might think that (1) is true iff torturing babies for fun does not maximize one's self-perfections.

This point — that we should not build substantive normative views into our semantics for moral terms — is not a new one: Though it is sometimes underappreciated in metaethics,<sup>4</sup> it has a rich history in philosophy of language and semantics. For example, here are Stalnaker and Lewis:

[O]ne can and should answer questions about the semantics and prag-

---

<sup>4</sup>Though certain expressivists, as I understand them, have also urged as much (e.g., GIBBARD 1990, 2003).

matics of constructions involving the words *good* and *right* without solving philosophical problems about the nature of morality. (STALNAKER 1988, p. 151)

The semantic analysis tells us what is true (at a world) under an ordering. It modestly declines to choose the proper ordering. That is work for a moralist, not a semanticist. (LEWIS 1978, pp. 85–86)

And slightly modifying a point of Kaplan's in a related context: If we avoid confusing substantive normative and metaphysical matters with phenomena of the semantics of moral terms, "the result can only be healthy for all... disciplines" (1989, p. 537). We can put the condition as follows: *The meanings of moral expressions do not presuppose some particular view on how to live.*

As our discussion of the role of content suggests, MORAL INVARIANTISM fails to capture what is said or conventionally communicated in speakers' utterances of moral sentences, and what it takes to be semantically competent with a moral expression, to know what a moral expression means. We said that being semantically competent with a sentence requires, perhaps among other things, being able to divide in the right way the alternative possible ways the world might be or might have been. But if certain moral claims are necessarily true (if true) or necessarily false (if false), then "understanding" such claims claim involves trivially "distinguishing" among worlds, either ruling them all out or ruling none out.

But how can we reject MORAL INVARIANTISM while accepting a truth-conditional picture of content? How can we accept that the semantic (conventional) content of (1) is (or determines) its truth conditions while denying that the conditions in the world that would have to obtain for (1) to be true constitute (or determine) the content of (1)? I suggest that it is the phrase "conditions *in the world*" that is an important source of our confusion. Distinguish the following three questions:

#### EXTENSIONALITY

What property (or properties), if any, do all and only wrong actions (necessarily) have?

#### GROUNDING

What, if anything, grounds the fact that wrong actions are wrong? Fundamentally, what, if anything, makes it the case that wrong actions are wrong?

#### SEMANTICS

Given a moral sentence  $\phi$ , what would the correct moral view have to be like for  $\phi$  to be true?

EXTENSIONALITY is a substantive normative question. It is a question of what the correct normative ethical view is. For example, if classical utilitarianism is correct, all wrong actions have the property of not maximizing overall happiness. By contrast, GROUNDING is, plausibly, a metaphysical question. What grounds the wrongness of torturing babies for fun may be that the action instantiates a certain natural property — per the reductive or non-reductive naturalist — or some irreducibly normative non-natural property — per the non-naturalist. For example, according to certain forms of metaethical constructivism, the wrongness of an agent's torturing babies for fun is grounded purely by properties concerning that agent's evaluative attitudes. EXTENSIONALITY and GROUNDING are thus substantive normative and metanormative questions. But they are not, strictly speaking, semantic questions, questions relevant in our theorizing about conventional natural language meaning. Semantic competence with normative terms like 'wrong' does not require that one have an answer, let alone the correct answer, to these questions.

Instead, I suggest that — as the name suggests — it is SEMANTICS that is the properly semantic question, a question of the conventional meaning of 'wrong'. Semantic competence with an ordinary factual sentence requires knowing how the world would have to be for the sentence to be true. Semantic competence with a de se factual sentence requires knowing when and where one would have to be located in the world for the sentence to be true. And, by extension, semantic competence with a moral sentence — knowing what a moral sentence means — requires knowing how things would have to be morally for the sentence to be true, or what the correct moral view would have to be like, for the sentence to be true. If knowing a sentence's meaning requires being able to distinguish possibilities in the right way, it is plausible that knowing a moral sentence's meaning requires being able to correctly distinguish certain moral possibilities. (Again, this is not to deny that there are other aspects of the conventional and non-conventional meaning, broadly construed, of moral terms.) This suggests that we individuate the truth conditions of moral sentences more finely than by their possible worlds truth conditions. Following the lead of Stalnaker, Lewis, and Kaplan, we should replace MORAL INVARIANTISM with

#### MORAL NEUTRALITY

Moral terms must be interpreted with respect to some moral standard or other, but their meanings do not presuppose some particular view on how to live.

AS MORAL NEUTRALITY suggests, the sorts of truth conditions that reflect what it takes to know what a moral sentence means are conditions on how things might be

morally. By contrast, what conditions must obtain *in the world* for a moral sentence to be true is a substantive issue. (I use ‘substantive’ broadly to describe first-order normative issues and non-semantic metanormative issues.) Failure to distinguish the semantics and metaphysics of moral talk, or distinguish semantically relevant truth conditions from substantive normative and metaphysical grounding conditions, has contributed to the confusion over the broader metaethical implications of accepting — and reaping the compositional benefits of — a truth-conditional semantics.<sup>5</sup>

We now have the resources to outline the form of a response to our initial puzzle. MORAL NEUTRALITY, unlike MORAL INVARIANTISM, does not render the semantically relevant truth conditions of (1) and (2) equivalent, even when conjoined with NECESSARY MORAL TRUTHS and our truth-conditional picture of content. Even if, as a normative or metaphysical matter, some moral sentences like (1) are true in all possible worlds, those sentences need not be true according to all moral standards or views on how to live (*cf.* n. 11). Just as knowing what ‘Grass is green’ means requires being able to rule out worlds in which grass isn’t green (regardless of whether or not the actual world happens to be one of them), knowing what (1) means requires being able to rule out moral standards according which torturing babies for fun isn’t wrong (regardless of whether or not the correct moral standard happens to be one them).

This response brings a number of questions in its wake. For one, how are we to cash out MORAL NEUTRALITY? We will address this question in §2, but first I would like to address three concerns.

Worry: The problem of logical omniscience is, as noted above, a general problem. It also arises, arguably, in the case of names, natural kind terms, mathematical expressions, etc. Should we think that invariantism for all these types of expressions is incorrect too? Why not think that solutions to the problem of logical omniscience in these other domains will carry over to provide a successful alternative treatment of moral terms?

Reply: Nothing in what I have said here requires that all instances of the problem of logical omniscience be handled in the same way. I see no a priori reason to think that moral terms, names, etc. should all be treated uniformly. After all, they are a pretty multifarious lot. Perhaps they should all be given a non-invariantist semantics of the sort I will defend in §2, but perhaps not. In the case of moral terms, such a semantics, I will argue, improves upon its most prominent alternatives. Whether it

---

<sup>5</sup>The problem is pervasive. See, e.g., BLACKBURN 1984, pp. 194–195, 1998, ch. 4; SCHROEDER 2011, pp. 5–6.



improves upon *all* alternatives obviously waits to be seen.

Worry: Even if we should not build substantive assumptions about which normative ethical theory is correct into the semantics, the meanings of moral terms might still presuppose some particular view on how to live. Why can't the advocate of MORAL INVARIANTISM say that moral terms are interpreted with respect to the correct moral theory, *whatever it is*? The description 'the correct moral theory, whatever it is' would pick out some particular moral view, but would not build in any substantive assumptions about which one it is.

Reply: To say that moral terms are interpreted with respect to "the correct moral theory, whatever it is" supports MORAL NEUTRALITY, not MORAL INVARIANTISM. For it effectively treats the meaning of a moral sentences as, in some sense (see §2), a function from moral standards to truth values. Compare: The meaning of an ordinary factual sentence like *S* 'Grass is green' does not presuppose some particular view on which world is actual. Though *S* is, in a certain sense, interpreted with respect to "the actual world, whatever it is," this is just to say that the meaning of *S* is a function from worlds to truth values; knowing what *S* means requires knowing that *if* such-and-such world is actual, *S* is true (false). Similarly, knowing what a moral sentence like (1) means requires knowing that if such-and-such moral standard is the correct one, (1) is true (false). But this is just what MORAL NEUTRALITY says.<sup>6</sup>

Worry: If moral terms are interpreted with respect to a varying moral standard, does this imply that there are no absolute, objective moral truths? Or that moral matters are merely "matters of taste"? Or that moral disagreement is "faultless"?

Reply: No such conclusions follow from anything said thus far. MORAL NEUTRALITY is a soberly semantic thesis. A semantics is supposed to capture, at the relevant level of abstraction, an expression's conventional meaning. Building the correct moral standard into the conventional contents of moral terms fails to capture what it takes to be semantically competent with them. One cannot avoid the "core questions" in metaethics simply by denying MORAL INVARIANTISM.<sup>7</sup> Ethicists and metaethicists can still examine questions like the following: Are there any normative reasons that all agents have? What natural or non-natural property, if any, do all and only things that are *N* (for some normative predicate '*N*') necessarily have? If there is such a property, is it reducible? Do normative facts hold independently of the evaluative attitudes of the agents to whom they apply? For all I have argued there may be independent support for all sorts of normatively and metaphysically

---

<sup>6</sup>Thanks to Jan Dowell for pressing me on this issue. As we will see in §2, I take the intuitions motivating this reply as favoring a non-contextualist implementation of MORAL NEUTRALITY.

<sup>7</sup>This is contrary to the view expressed in SCHROEDER 2010, pp. 9, 13, 153, 211.

robust ethical and metaethical positions. But such support will need to be just that: *independent* (of semantics). Plausibly the meanings of moral terms do not presuppose verdicts on such lofty issues. At minimum, the way the semantics and other areas interact is going to be more complicated than we might have initially thought.


As we will see in §3, delineating issues often conflated in metaethics under the general heading of “meaning” will have significant payoffs. Among other things, it will suggest an alternative way of capturing the putative objective purport of moral language. More generally, it will help us articulate certain classical metaethical questions in a more perspicuous way, and clarify what is at stake in the dialectic between overall metaethical theories. Before addressing these issues, let’s see if we can make MORAL NEUTRALITY more precise.

## 2 Contextualism and condition semantics

Here, again, is our conclusion to be explained: Moral terms are interpreted with respect to some relevant moral standard but do not themselves presuppose some particular view on how to live. But what is the “relevant moral standard”?

Two familiar answers suggest themselves. To set the stage, let’s briefly get some terminology on the table. On a broadly Kaplanian view there are two senses of the “meaning” of an expression: its *character* and its *content* (KAPLAN 1989). *Character* is a function from contexts of use to contents and is constant across all literal uses of an expression; it represents an expression’s linguistic meaning, or what you know when you know what an expression means. *Content* is (or at least determines) an intension, a function from *indices* — or, in Kaplan’s terminology, “circumstances of evaluation” — to extensions. (More on “indices” below.) The content of a predicate is a property, and the content of a sentence is a proposition. The semantic content of a sentence is, roughly, what is conventionally communicated by the sentence in the relevant context, as determined by the language’s composition rules. The *context*

Figure 1: Context-index semantics (from LEWIS 1980, p. 96)



includes parameters that determine content from character, fixing the reference of

indexicals and other deictic expressions. The *index* includes parameters of evaluation that determine extension from content (e.g., the truth-value of a proposition). The index consists of those features of context that can be shifted by operators in the language (e.g., modals, attitude verbs). For Kaplan, the index is just a possible world, or perhaps a world-time pair. The extensions of expressions are thus defined relative to points of evaluation, consisting of *context* and *index*.

For example, in a context  $c$  where Sara is speaker, the sentence  $S$  ‘I like brownie sundaes’ expresses the proposition, or conventionally carries the information, that Sara likes brownie sundaes. Supposing Sara actually likes brownie sundaes, the proposition expressed by  $S$  in  $c$  is true in the actual world. But considering a possible world  $u$  in which Sara does not like brownie sundaes, the proposition expressed by  $S$  in  $c$  is false when evaluated at  $u$ . (In set-talk: the actual world, but not  $u$ , is in the proposition expressed by  $S$  in  $c$ .) This is all more-or-less standard fare from Kaplan (cf. STALNAKER 1970, KAMP 1971, LEWIS 1980).

This context-index framework (or “double indexing semantics,” as it is sometimes called) suggests two familiar answers to the question — call it ‘COORDINATE’ — of what moral standard is “relevant” for the interpretation of moral terms: (a) that it is supplied from the context, and (b) that it is supplied from the index. Call a view according to which the relevant standard is supplied from the context *contextualism*. Call a view according to which the relevant standard is supplied from the index *condition semantics*. (Reasons for this name will become evident below. For those familiar with the recent literature on contextualism and relativism about epistemic modals and predicates of personal taste,<sup>8</sup> it might seem that what I am calling ‘condition semantics’ is often called ‘relativism’. However, as I will suggest in §3, relativism is better understood as a particular type of condition semantics, one that takes a stand on certain further issues that are orthogonal to the points at issue between contextualism and condition semantics considered in this section. To make terminological matters worse, many positions in the ethics and metaethics literatures described as versions of “moral relativism” are really instances of contextualism on the current taxonomy.)

The debate over COORDINATE has centered around the problem of disagreement, the problem of accounting for how speakers uttering moral sentences  $\phi$  and  $\neg\phi$ , respectively, count as disagreeing, rather than, e.g., as making consistent claims about their respective moral standards. I would like to put this way of approaching COORDINATE to the side. (For one thing, I take it that a semantics need only focus on

---

<sup>8</sup>See, e.g., DEROSE 1991, EGAN ET AL. 2005, LASERSOHN 2005, EGAN 2007, MACFARLANE 2007, STEPHENSON 2007, YALCIN 2007, VON FINTEL & GILLIES 2008, DREIER 2009.

accounting for inconsistency or contradictoriness rather than disagreement, a plausibly pragmatic notion, but no matter: It is contentious whether either side has a distinct advantage in stating a notion of disagreement that accounts for all the data anyway.<sup>9</sup>) Instead I would like to revive the debate over COORDINATE by turning to a different issue, one more neatly tied to the semantics of normative language: How do contextualism and condition semantics fare in their respective abilities to capture what you know when you know what a moral sentence means, or what is required for being semantically competent with a moral sentence? On my preferred understanding of contextualism and condition semantics, both are moved by MORAL NEUTRALITY;<sup>10</sup> however, they intend to capture it in different ways. These differences generate different accounts of the contents of moral sentences. In the remainder of this section I will argue that, in light of our above discussion about the role of content in semantic theorizing, condition semantics's account of the contents of moral sentences better captures what we know when we know what moral sentences mean than contextualism's. (Again, I want to emphasize that we are considering specifically truth-conditional aspects of meaning, as described in §1. Given our purposes, I leave open whether the arguments below generalize to motivate condition semantics for personal taste predicates, epistemic modals, vague terms, etc.; however, I am inclined to think they do.)

Start with contextualism. According to contextualism, moral terms are interpreted with respect to a contextually relevant moral standard, where which particular moral standard is relevant varies across contexts. ' $\phi$  is wrong' is context-sensitive, expressing different propositions in different contexts, even if  $\phi$  contains no context-sensitive items itself. Roughly, (1) is true at a point of evaluation — a context-index pair — iff torturing babies for fun is wrong according to the relevant moral standard of the context.<sup>11</sup>

---

<sup>9</sup>See, e.g., MACFARLANE 2007, VON FINTEL & GILLIES 2008, DREIER 2009.

<sup>10</sup>*Pace* DREIER 1990, 2009, which views the main motivation (of at least contextualism) as being to explain judgment internalism.

<sup>11</sup>As is standard in descriptive semantics, the analyses given here and below are not intended to provide individuals who do not understand ethical terms with such an understanding; rather they are models, at the relevant level of abstraction, of speakers' dispositions to interpret moral language (*cf.* n. 3). Even so, the standard frameworks of ordering semantics and premise semantics (e.g., LEWIS 1973, VELTMAN 1976, KRATZER 1981, 1991) suggest a number of ways of giving an illuminating, non-trivial characterization of what it is for an action to be wrong according to a moral standard. One option is to treat a moral standard as a set of propositions  $P$  whose truth is required by morality, on some view of what morality requires, or as a set of propositions that describe some putative moral ideal (e.g., propositions like that no murders occur, that overall average happiness is maximized, that each individual maximizes her own self-perfections, etc.). To a first approximation, ' $\phi$  is wrong'

Does this contextualist proposal capture what you know when you know what (1) means? Or what you believe when you believe what (1) says? Start with mental content. Is believing that torturing babies is wrong a state of mind about one's moral standards? It seems not. Assume that believing what a sentence *S* says is a matter of believing the proposition expressed by *S*. First, if I ask you why to believe that torturing babies is wrong, I am asking for considerations regarding the action of torturing babies, not considerations regarding your, or whomever else's, moral views about torturing babies.

Second, suppose you offer George, your pet chimpanzee, a brightly colored plant to eat. Having grown up in the great outdoors, George recognizes that the plant is poisonous and recoils. You say:

(4) George knows he shouldn't eat the plant.

This seems true. But, contrary to the contextualist, it is implausible that (4) is true only if George knows that his eating the plant is incompatible with his, or anyone else's, normative views. (Obviously I am generalizing to types of normative language beyond the narrowly moral.) He's just a monkey. He might barely know he exists. Similarly, suppose you encourage Gabriel, your infant brother, to put his fingers into the electrical outlet. Gabriel, smart chap that he is, recoils; his mother has repeatedly scolded him not to do so. You say:

(5) Gabriel knows he shouldn't put his fingers into the outlet.

This seems true; you are attributing a certain normative belief to Gabriel. But it is implausible that (5) is true only if Gabriel has a belief about his, or anyone else's, normative views. After all, he's just a baby.<sup>12</sup>

In these ways contextualism seems to have trouble capturing what your state of mind is when you have a normative belief. It is (6c), not (6b), with which (6a) has an important semantic connection (though this is not to say they are equivalent).

- (6)
- a. Sara thinks torturing babies for fun is wrong.
  - b. Sara thinks torturing babies for fun is wrong according to her.
  - c. According to Sara's moral views, torturing babies for fun is wrong.

---

would mean that  $\neg\phi$  follows from every maximally consistent subset of *P*. Equivalently, the semantics could be couched in terms of orderings induced by *P* (LEWIS 1981). I continue to talk of "moral standards" simply for heuristic purposes. Jumping ahead, whereas the contextualist would treat *P* as supplied by the context, the condition semanticist would treat it as supplied by the index.

<sup>12</sup>Cf. YALCIN 2007, p. 997.

Whether one can represent or take a certain perspective on normative standards is independent of whether one can *have* a normative standard. Call this point ‘FIRST-ORDER STATES OF MIND’.

Return to linguistic content. Consider the context-dependency of words like ‘I’ and ‘every’. Knowing the linguistic meaning of ‘I’ requires being able to say, given a context of utterance  $c$ , that ‘I’ picks out the speaker of  $c$ . Knowing what ‘I like brownie sundaes’ means requires knowing, given a context  $c$ , that the speaker of  $c$  likes brownie sundaes. If you don’t know who the speaker is in the context, you don’t know what was said, or what proposition was expressed. You don’t know what possibilities have been ruled out, and whether your information state is compatible with the possibilities that have not been ruled out. Similarly with quantifiers. Knowing what ‘Everyone likes brownie sundaes’ means requires knowing that, given a context  $c$ , for all  $x$  in the relevant domain specified by  $c$ ,  $x$  likes brownie sundaes. If you don’t know what the contextually relevant domain is, you don’t know what proposition was expressed (every human being? every child?). You don’t know whether your information state is compatible with the possibilities that have not been ruled out. The same goes for other types of context-sensitive expressions — e.g., demonstratives, gradable adjectives, relational expressions (‘enemy’, ‘local’), neutral modal verbs (‘can’, ‘must’), adverbial quantifiers (‘usually’, ‘occasionally’) — but you get the idea.

Does the same hold with moral terms? It seems not. Suppose the contextually relevant moral standard for the interpretation of (1) is Sara’s. Intuitively, it doesn’t seem that knowing what (1) says in this context involves an ability to rule out the following possibilities: those in which torturing babies for fun isn’t wrong according to Sara. By contrast, knowing what ‘I like brownie sundaes’ says when uttered by Sara does seem to involve an ability to rule out the following possibilities: those in which Sara doesn’t like brownie sundaes.

This point is suggestive but perhaps not decisive. We saw that with context-dependent items like ‘I’ and ‘every’, among others, we can only retrieve the intended truth conditions given certain contextual cues; the semantic interpretation of such items is incomplete in an important sense. Is there a similar semantic incompleteness with moral terms? It seems not. I can determine the intended truth conditions of a sentence like (1) even if I have no idea who the speaker is or what the context is like. Similarly, consider (7).

(7) Nathan should be here by now.

I don’t know what (7) says unless I know what flavor of modality ‘should’ expresses.

Is ‘should’ to be given an epistemic reading? A deontic reading? But once I know that the context determines, say, a particular type of deontic reading — and I know the referents of ‘here’ and ‘now’ — I seem to know precisely what (7) says. And I know whether my information state is compatible with the possibilities it doesn’t rule out. Call this point ‘SEMANTIC COMPLETENESS’.

In these ways, the meanings of normative terms do not pattern with the meanings of context-dependent expressions. Knowing what ‘wrong’ means does not require knowing that, given a context  $c$ , ‘wrong’ picks out those actions that are wrong according to the relevant standard of  $c$ .<sup>13</sup>

It is obviously beyond the scope of the present paper to demonstrate clear disanalogies between moral terms and every kind of context-sensitive expression. But, first, my goal here is to help build a good inductive case that moral terms are not context-sensitive in the sense at issue. What is lacking is a contextualist analysis that meets all of the challenges canvassed. In certain semantically relevant respects, moral terms behave quite unlike pure indexicals, demonstratives, nominal and adverbial quantifiers, relational expressions, gradable adjectives, and (neutral) modal auxiliaries, all generally accepted as types of context-sensitive expressions. This gives us reason to investigate whether any non-contextualist analyses fare better. Second, the arguments presented here provide a case against moral terms being context-sensitive in a *lexically specific* way. This leaves open whether moral terms are context-sensitive in the broader sense concerning the conversational, discourse-level standard-setting discussed in LEWIS 1979. Ultimately I think that condition semantics for moral terms integrates better with this type of pragmatic picture of conversation. But that is a topic for another paper (see SILK 2012b).

These same phenomena that the contextualist has difficulty capturing support condition semantics. First let’s clarify the rough-and-ready explication of condition semantics from the beginning of this section. We started with a standard view of content according to which contents function to distinguish among possibilities. An ordinary factual sentence, on this view, places a condition on possible worlds;

---

<sup>13</sup>The contextualist might respond by treating the semantic content of a moral claim as a set of contexts, rather than as a set of worlds or indices (see SCHLENKER 2003, ANAND & NEVINS 2004, SHARVIT 2008; though *cf.* VON STECHOW 2003).

(i)  $\lambda c . \llbracket \phi \rrbracket^{c, w_c}$

On this analysis, the content of (1) is the set of contexts such that torturing babies for fun is wrong according to the relevant moral standard of that context. There may be reasons for introducing this kind of context shifting into the semantics. But insofar as doing so effectively integrates moral standards into the semantic content in the same way as worlds, this type of analysis appears to be structurally isomorphic to condition semantics.

it describes what property a world must have for it to be characterized by that sentence. But as we noted in our discussion of problems facing MORAL INVARIANTISM, sentences do not always simply characterize ways the world, narrowly construed, might be. Sentences can distinguish among ways things might be, broadly understood — for example, regarding how things might be morally, or what the correct moral view(s) might be. Condition semantics captures this feature of moral thought and talk. According to condition semantics, sentences (in context) place conditions on various sorts of items, conditions that the items in question must satisfy in order for them to be characterized by those sentences — hence ‘condition semantics’. Moral sentences place conditions on, or express properties of, moral standards just like ordinary factual sentences place conditions on, or express properties of, possible worlds. Just as ordinary factual sentences test whether a possible world meets a certain condition, moral sentences test whether a moral standard meets a certain condition. (If condition semantics is correct for predicates of personal taste, epistemic modals and adjectives, vague predicates, etc. — again, a point I leave open here — we would likewise understand sentences containing these expressions as placing conditions on standards of taste, credal states, standards of precision, etc.)

We noted in §1 that semantic competence with an ordinary factual sentence like ‘The diamond is in the box’ requires being able to say, given some situation  $w$ , whether the diamond has the property of being inside the box in  $w$ . This motivated the view that ‘The diamond is in the box’ has as its content a certain set of possible circumstances, those in which the diamond is in the box. For the condition semanticist, the picture is precisely parallel. Semantic competence with a moral sentence requires being able to correctly say whether, for any relevant moral standard  $i$ , if things are *that* way morally, or if *that* is the correct moral view, then the sentence is true. The contents of moral sentences give a rule for selecting a subset of moral standards from some domain — in the case of (1), a rule for selecting those moral standards according to which torturing babies for fun is wrong. So, if the standard possibility-distinguishing view of content is roughly on the right track, condition semantics seems to be its natural heir in the case of moral terms.

Condition semantics captures these points formally by interpreting moral terms with respect to a parameter of the index, rather than a parameter of the context. For the contextualist (and the invariantist), the index consists of a world parameter, and perhaps also a time parameter. However, there is nothing inherent in the Kaplanian framework that requires this restriction. As even Kaplan and Lewis acknowledge, the index can consist of any parameters that can be shifted by expressions that oper-



ate on content.<sup>14</sup> For the condition semanticist, the index also includes a parameter  $i$  that ranges over moral standards. Indices are thus treated as pairs  $\langle w, i \rangle$  of a world  $w$  and moral standard  $i$ . (I will systematically ignore times.) The contents of sentences and attitudes are treated as sets of  $\langle w, i \rangle$  pairs, rather than as sets of worlds, as per the invariantist and the contextualist. (I will continue to call such contents ‘propositions’ for the moment, but we will return to this in §3.) The extension of an expression — e.g., the truth value of a sentence — is determined with respect to a context and a  $\langle w, i \rangle$  pair, which is the index. Just as ordinary factual sentences denote truth values with respect to a world of evaluation, so, on this picture, do moral sentences denote truth values with respect to a moral standard (and world). (If one likes, the values “true” and “false” can be understood simply as characterizing sets in the usual way.) We can represent the truth conditions of (1) roughly as follows, in equivalent function- and set-talk (see n. 11).

- (8)     a.   (1) is true in  $c$  at  $\langle w, i \rangle$  iff torturing babies for fun is wrong according to  $i$  in  $w$   
           b.    $\{\langle w, i \rangle$ : torturing babies for fun is wrong according to  $i$  in  $w\}$

Though moral sentences — unlike simple, ordinary factual sentences — have truth conditions that need not place non-trivial conditions on the world parameter, they have truth conditions just the same.

Two points of clarification. First, condition semantics is not simply a kind of “centered worlds” semantics. Centered content is typically understood in terms of self-location. But distinguishing among moral standards does not seem to be reducible to distinguishing among where one might be located in the world. We must distinguish individuals from the moral standards they might accept.<sup>15</sup>

Second, we can now give a more precise analysis of what a “moral sentence” is. Although all propositions may be treated as sets of world-moral standard  $\langle w, i \rangle$  pairs, not all propositions will interestingly depend on the value of moral standard parameter. Morality-*independent* propositions are sets of  $\langle w, i \rangle$  pairs such that if one pair is in the set, so is every pair with the same world coordinate. Ordinary factual sentences like ‘Sara likes brownie sundaes’ express morality-independent propositions. Morality-*dependent* propositions, like (1), are sets of  $\langle w, i \rangle$  pairs that

<sup>14</sup>LEWIS 1980, pp. 84–87 and KAPLAN 1989, pp. 502, 511, 511n.35 are particularly forthright on this point. There is nothing “radical” in a relativist semantics as such (*pace* ?, 272; ?, 399, 408; ?, 274, 278; BJÖRNSSON & FINLAY 2010, 11, 25; and ?, 270, among others).

<sup>15</sup>*Pace* suggestions in EGAN ET AL. 2005, EGAN 2007, STEPHENSON 2010, WEATHERSON 2008. Thanks to Allan Gibbard for pressing me on this point.

are not morality-independent propositions. *Moral sentences*, we might say, are those sentences that express morality-dependent propositions. (So, on this semantic characterization, not all complex sentences containing moral terms will be “moral sentences” in this sense.) Though this is certainly not the only difference between moral and non-moral language, moral sentences, unlike non-moral sentences, place non-trivial conditions on the moral standard parameter.

Condition semantics is an improvement. First, it improves on MORAL INVARIANTISM and captures the intuitions from the beginning of §1. Moral sentences won’t be necessary in the sense of being “true at all indices.” (1) has non-trivial content: it is false with respect to moral standards which permit torturing babies for fun; it rules out those moral standards which do not prohibit torturing babies for fun. We can now see that there is something problematic in the contextualist’s response to our puzzle in §1. Our goal, in part, was to show that moral sentences non-trivially distinguish among possibilities. Since our contextualist treats contents as ordinary sets of worlds, she has to find a non-empty set of worlds to rule out. She does so via the assumption that the contextually relevant group might have different moral views in different worlds. (1), for example, would rule out those worlds in which the contextually relevant moral standard — say, Sara’s — permits torturing babies for fun. Though this indeed distinguishes the contents of (1) and (2), it does so in the wrong sort of way. (1) is not a claim about Sara or her moral views. It is a claim about the moral status of torturing babies. This is what distinguishes (1) from a sentence like (6c), reproduced below, or (9), said in a context in which we are cataloguing what Sara’s moral views are.<sup>16</sup>

- (6c) According to Sara’s moral views, torturing babies for fun is wrong.
- (9) Torturing babies for fun is “wrong.”

The possibilities (1) rules out are possible moral standards, possible ways things might have been morally. Condition semantics captures this intuition.

Second, condition semantics has a natural explanation of SEMANTIC COMPLETENESS. Since the relevant moral standard does not figure into the content of moral claims — it is used in determining extension *from* content — we do not need a suitably specified context to determine what a given moral claim says (abstracting away from other context-sensitive items).

Finally, let’s return to mental content. Condition semantics offers a natural explanation of FIRST-ORDER STATES OF MIND. To see this, start with a conserva-

---

<sup>16</sup>On this distinction, see, e.g., HARE 1952, VON WRIGHT 1963, HANSSON 1969. For related discussion in linguistics on “objective” and “subjective” modalities, see LYONS 1977.

tively enriched standard modal semantics for attitude verbs like ‘believe’ (HINTIKKA 1969). As usual, attitude verbs can be treated as shifting the index. For condition semantics this means that attitude verbs shift not only the world parameter but also the moral standard parameter. So, what the attitude verb does is quantify over the set of possibilities, now represented at the relevant level of abstraction as a set of world-moral standard pairs, that are compatible with the attitude state. ‘S believes  $\phi$ ’ says that S’s belief state entails the content of  $\phi$ , relative to the given context (cf. n. 22). As in unembedded cases, complements that contain moral terms may have extensions that non-trivially depend on the value of the moral standard parameter. ‘Sara believes that torturing babies for fun is wrong’ says, roughly, that for all  $\langle w', i' \rangle$  pairs compatible with Sara’s beliefs (in the world of evaluation), torturing babies for fun is wrong according to  $i'$  in  $w'$ . The belief ascription is true iff Sara’s moral views are a certain way, i.e., iff all live moral standards, given her moral views, prohibit torturing babies for fun. In this way, we can capture FIRST-ORDER STATES OF MIND and the close semantic connection between sentences like (6a) and (6c).

We are now in a position to compare our approach to a Stalnakerian diagonalization strategy and an impossible worlds approach, briefly mentioned in §1 (cf. n. 13). Diagonalization is a means of *reinterpretation*; it is triggered by the assertion of a sentence that would be uninformative or trivial if interpreted with its literal meaning (STALNAKER 1978). First, diagonalization, in this sense, is in tension with MORAL NEUTRALITY: If moral sentences semantically express ordinary possible worlds propositions and are interpreted with respect to an invariant moral standard, then their semantic contents presuppose some particular view on how to live. Irrespective of its story about how we interpret moral language in conversation, a reinterpretation strategy faces the same problems as MORAL INVARIANTISM: it fails to account for what it is to be *semantically* competent with moral language. Second, contingently true moral sentences will not trigger diagonalization. So, for any contingently true moral sentence ‘ $\phi$  is  $\mu$ ’ (pick your favorite), for some moral predicate  $\mu$ , the diagonalizer will have no explanation for how its semantic content differs from that of ‘ $\phi$  is  $\delta$ ’ (for some associated descriptive, non-moral predicate  $\delta$ ) that would express the same possible worlds proposition. The main difference between condition semantics and Stalnakerian diagonalization is that, for condition semantics, *reinterpretation* is not necessary to capture the communicative content of moral sentences. Semantic competence itself requires moral expressions to be interpreted with respect to a given moral standard.

As for the impossible worlds approach, we can now see that it will end up looking quite similar to condition semantics. Whereas condition semantics separates the possible world and moral standard parameters, the impossible worlds theorist

builds the moral standard into the (im)possible world. (For example, the truth conditions of (1) would be the set of (im)possible worlds  $w'$  such that torturing babies for fun is wrong according to the correct moral standard in  $w'$ .) For any  $\langle w, i \rangle$  pair there is a unique (im)possible world identical to  $w$  in all non-normative respects in which  $i$  is the correct moral view. And if we restrict ourselves to impossible worlds whose impossibility is a result of what moral views are correct in them, the mapping from  $\langle w, i \rangle$  pairs to (im)possible worlds is one-to-one. So, which analysis one adopts may depend on one's commitments regarding the broader utility and theoretical role of impossible worlds (see n. 1). But even if one wished to go in for impossible worlds on independent grounds, separating in our theories the world and moral standard parameters of the index can illuminate more perspicuously the structure of our semantic competence with moral terms and what is special about their meanings as compared with the meanings of ordinary non-normative descriptive terms. And depending on one's views about representation and intentional content more generally, one might think that representing moral judgments in terms of  $\langle w, i \rangle$  pairs rather than impossible worlds will lend itself to a preferred account of how particular morality-dependent contents count as being the contents of agents' beliefs and other attitudes. (We will return to this issue in §3.)<sup>17</sup>

### 3 A way forward in metaethics

In §1 I argued that we should not build substantive normative assumptions into the meanings of moral terms. This motivated contextualism and condition semantics over invariantism. In §2 I argued that a variety of further phenomena motivate condition semantics over contextualism: Moral sentences place conditions on moral standards.

Now for a hitch — okay, two hitches. First, condition semantics may seem similar to certain versions of relativism in the sense familiar from recent contextualism-relativism debates. Is condition semantics committed to being “relativist”?

Second, we articulated condition semantics within a familiar truth-conditional framework. However, condition semantics seems to look a lot like the version of expressivism developed in GIBBARD 1990, 2003, which, at least *prima facie*, purports to be *non-truth-conditional*.<sup>18</sup> In Gibbard's (2003) response to the Frege-Geach

---

<sup>17</sup>Thanks to Andy Egan and Eric Swanson for helpful discussion.

<sup>18</sup>However, condition semantics is clearly unlike old-school versions of emotivism and expressivism according to which utterances of moral sentences conventionally express non-contentful states of mind — e.g., states expressed by some variant on ‘Boo!’ or ‘Hurray!’ Thanks to Matthew

problem he develops a normative logic based, roughly, on assigning truth values to sentences relative to triples of context, world, and “hyperplan.” (A hyperplan is a maximal contingency plan, or a plan that, for any occasion or decision one might conceivably be in, either forbids or permits an alternative open on such an occasion. The semantics in his 1990 is couched in terms of “systems of norms” rather than hyperplans.) The content of an attitude or judgment is given in terms of the world-hyperplan pairs it rules out. This formalism should sound pretty familiar. What gives?

In reply, I understand condition semantics as offering an attractive *framework* for thinking about broadly metaethical questions. Much of the formal apparatus is, on its own, neutral on various broader philosophical issues. Certain of these issues can then be understood as characterizing the debates over relativism and expressivism. (This is not to say that the condition semantics framework is neutral on all philosophical issues. For example, it presupposes the substantive picture of content described in §1, and rules out invariantism and contextualism about moral terms.) Contemporary metaethicists can — and, if the arguments in §§1–2 are right, *should* — all accept condition semantics for giving their compositional descriptive semantics for moral terms.<sup>19</sup> Where they will differ is on certain further issues in, e.g., philosophy of language, philosophy of mind, and psychology. We can begin to delineate the landscape of resulting metaethical positions in terms of the following sorts of questions, among others.

#### PROPOSITION

What are propositions? Are they identical to — or, at least, do they determine — sets of possible worlds, or sets of  $\langle w, i \rangle$  pairs? Are the abstract objects that a semantics assigns as the contents of declarative sentences and the complements of attitude verbs the same sorts of entities that are the contents of human states of mind?

#### SEMANTICS

What is explanatorily fundamental in a theory of meaning? Knowledge of truth conditions? Or satisfaction of norms regarding how expressions are conventionally used?

#### MONADIC TRUTH

How should we define a monadic truth predicate? Is there a single moral

---

Chrisman for pressing me on relevant distinctions between expressivist theories.

<sup>19</sup>I mean ‘descriptive semantics’ in the sense of STALNAKER 1997, p. 535; cf. KAPLAN 1989, pp. 573–576 on ‘semantics’.

standard that is correct for all contexts? Or can the moral standard relevant in evaluating whether a moral sentence is true simpliciter vary across contexts?

#### METASEMANTICS

What makes a term *normative*? What is it about a linguistic community that makes it the case that a certain term has a morality-dependent content, and that it has the particular morality-dependent content that it does? Do facts about the term's regulative, directive role in planning and practical reasoning play an essential role in this sort of content-determination?

#### MIND

What makes a particular abstract object count as the content of an agent's attitude? Do facts about the agent's motivational state or emotional capacities play an essential role in this sort of content-determination?

#### PSYCHOLOGY

Are the attitudes agents bear toward morality-dependent contents essentially practical and action-guiding? Is there a special tie between moral judgment and action?

No doubt these questions are at least dialectically, if not logically, related. Though defenses of answers to these questions must be reserved for elsewhere, the following clarificatory remarks may be helpful.

Though all metaethicists can accept that moral sentences have truth conditions, in the sense of sets of indices, whether we say they express "propositions" will depend on whether that label is better reserved for ordinary sets of worlds or for sets of our enriched indices. Given the theoretical roles propositions are usually thought to play — namely, of being the contents of declarative clauses and propositional attitudes — there is little reason to prejudice the issue in favor of the former. However, if one has a more metaphysically robust job description for propositions, calling perhaps for something that carves "fundamental reality" (whatever that is) at its joints, one might have qualms about calling sets of  $\langle w, i \rangle$  pairs 'propositions'.

Though condition semantics makes use of the notion of "truth at an index," as is standard, it can still capture our intuitive notion of truth simpliciter (in a context). According to KAPLAN (1989, pp. 522, 547), an ordinary non-normative sentence is true in a context  $c$  iff it is true at the point of evaluation  $\langle c, w_c \rangle$ , where  $w_c$  is the world of  $c$  (cf. LEWIS 1980, p. 88). By extension, we can say that a (moral) sentence is true — or, we might say, "correct" — in a context  $c$  iff it is true at the point of evaluation  $\langle c, w_c, i_c \rangle$ , where  $i_c$  is the correct moral standard of the context.

Note that this characterization is neutral on whether there is a single correct moral standard for all contexts. I suggest that it is by this question, the question targeted by MONADIC TRUTH, that we characterize “relativist” positions. To capture common relativist claims that morality is not objective — or at least that moral language does not presuppose that morality is objective — the relativist might say that a moral sentence is correct, or true in a context, iff it is true according to the relevant moral standard of the context, where what standard is relevant can vary across contexts. Conflicting moral judgments about a particular case may thus both be able to be correct. Those who defend the objectivity of morality — or at least the objective purport of moral language — would deny this by defining a monadic truth predicate in terms of some universal moral standard.<sup>20</sup>

I suggest that it is the questions in METASEMANTICS, MIND, and PSYCHOLOGY that locate many important points of contention in the debates over expressivism. Is a morality-dependent content assigned to a term because the term plays an essentially regulative, directive role in planning or practical reasoning? Are facts about an agent’s conative, practical, or motivational state essential in determining that some particular set of  $\langle w, i \rangle$  pairs is the content of her moral belief? If so, what role and which facts? (No doubt GIBBARD 1990 will be relevant here.) Though the non-expressivist would answer these questions in the negative, she might grant that moral beliefs are often intimately connected with an agent’s emotional capacities and practical attitudes, and that utterances of moral sentences can non-conventionally express such attitudes in certain contexts.<sup>21</sup> These questions may also shed light on the debate over judgment internalism. If having certain emotional capacities is necessary for one to count as bearing the acceptance relation to a certain morality-dependent content, this could clarify the sense in which the amoralist or sociopath may not count as genuinely accepting any moral claims.

It is widely assumed that expressivism is committed to explaining the semantic properties of a sentence  $S$  (e.g., inconsistencies, entailments) in terms of properties of the psychological attitudes conventionally expressed by utterances of  $S$ , rather than in terms of properties of the content of  $S$ . The above discussion suggests that there are positions continuous with historical versions of expressivism that reject

---

<sup>20</sup>Thanks to Allan Gibbard and Jason Konek for pressing me on this issue. Depending on how one answers MONADIC TRUTH, one’s definition of a monadic truth predicate may introduce a disanalogy between normative and non-normative sentences: Unlike in the case of non-normative sentences, the question of which normative sentences are true simpliciter may be a normative question. Regardless, as noted in §1, one does not need to be able to determine whether a sentence is true simpliciter in order to be semantically competent with it.

<sup>21</sup>Cf. “hybrid” versions of expressivism (see SCHROEDER 2009 for an overview).

this commitment, positions that accept the common assumption in semantics that, at the relevant level of abstraction, semantic properties of sentences are to be explained in terms of formal properties of their contents. Given the alleged problems incurred by this commitment, as discussed in the literature on the Frege-Geach problem, this may prove a welcome result for theorists attracted to the psychological aspects of expressivism. (Though see SILK 2012a for arguments that expressivists can still solve the negation problem even while taking the above explanatory commitment on board.)<sup>22</sup>

In these ways, a plurality of implementations of condition semantics are possible depending on one's broader philosophical commitments. All metaethicists, I have argued, have reason to accept condition semantics for moral terms in giving their compositional semantics. And they can do so in a way that does not prejudice certain further issues concerning the "meaning" of moral terms, broadly construed. Instead, we can locate many classic metaethical debates in certain empirical and philosophical matters in psychology, philosophy of mind, and philosophy of language more broadly. The framework of condition semantics suggests new ways of clarifying what is at issue in these debates.

## 4 Conclusion

I have argued that reflection on what it is one knows when one knows what moral terms mean motivates a "condition semantics." Contents distinguish among possibilities. Semantic competence with moral language requires, perhaps among other things, having the capacity to correctly divide the space of moral standards, or distinguish among ways things might be morally. Just as ordinary factual sentences

---

<sup>22</sup>For arguments against treating expressivism as committed to explaining semantic properties in terms of properties of psychological attitudes, see CHARLOW 2011, App. A. Though I have characterized expressivism as agreeing with non-expressivism on what sort of descriptive semantics to adopt for normative language, there will of course be alternative characterizations without this feature. For instance, though Gibbard's extension of ordinary possible worlds semantics treats sentential contents, roughly, as sets of pairs of worlds and normative standards, an alternative, perhaps more in the spirit of Gibbard's own non-technical discussions, is Swanson's (2012) constraint semantics (*cf.* MOSS 2012). This does away with worlds in the index and treats the primary semantic values of all sentences as sets of admissible states ("constraints") — credal states, in the case of ordinary factual sentences and epistemically hedged sentences — and treats attitude ascriptions as saying that the subject's state satisfies the constraint expressed by the complement. This is, in effect, a static version of one kind of familiar dynamic semantics which treats the primary semantic values of sentences as relations between information states or contexts (perhaps understood in the enriched way described in LEWIS 1979).



place conditions on possible worlds, moral sentences place conditions on moral standards. The resulting type of analysis improves upon invariantism: it's not the case that a mere world of evaluation determines what moral standard is relevant for the interpretation of moral terms. Condition semantics also improves upon contextualist analyses: knowing what a moral sentence means requires being able to correctly distinguish among moral standards, not among worlds based on whether such-and-such relevant standard *G* of the context is at it is; it requires being able to pick out the set of moral standards according to which the sentence is true, not the set of worlds in which the sentence is true according to *G*. The framework of condition semantics also offers a perspicuous way of posing various classical questions in metaethics and related areas. This, I hope, will motivate clearer, better motivated answers and suggest new ways the dialectic may proceed.

## References

- ANAND, PRANAV & ANDREW NEVINS. 2004. Shifty operators in changing contexts: Indexicals in Zazaki and Slave. In *Proceedings of SALT 14*. Evanston, IL.
- BJÖRNSSON, GUNNAR & STEPHEN FINLAY. 2010. Metaethical contextualism defended. *Ethics* 121, 7–36.
- BLACKBURN, SIMON. 1984. *Spreading the word*. Oxford: Clarendon Press.
- BLACKBURN, SIMON. 1998. *Ruling passions*. Oxford: Clarendon Press.
- CHARLOW, NATHAN A. 2011. *Practical language: Its meaning and use*. Ph.D. thesis, University of Michigan.
- CRESSWELL, M.J. 1978. Semantic competence. In F. GUENTHNER & M. GUENTHNER-REUTTER (Eds.), *Meaning and translation: Philosophical and logical approaches*, pp. 9–27. New York: New York University Press.
- CRESSWELL, M.J. 2002. Why propositions have no structure. *Noûs* 36, 643–662.
- CROSS, CHARLES B. 2001a. The paradox of the knower without epistemic closure. *Mind* 110, 319–333.
- CROSS, CHARLES B. 2001b. A theorem concerning syntactical treatments of non-idealized belief. *Synthese* 129, 335–341.

- DEROSE, KEITH. 1991. Epistemic possibilities. *Philosophical Review* 100, 581–605.
- DEUTSCH, HARRY. 2008. Review of Jeffrey C. King, *The nature and structure of content*. *Notre Dame Philosophical Reviews* URL <http://ndpr.nd.edu/review.cfm?id=13165>.
- DREIER, JAMES. 1990. Internalism and speaker relativism. *Ethics* 101, 6–26.
- DREIER, JAMES. 2009. Relativism (and expressivism) and the problem of disagreement. *Philosophical Perspectives* 23, 79–110.
- EGAN, ANDY. 2007. Epistemic modals, relativism, and assertion. *Philosophical Studies* 133, 1–22.
- EGAN, ANDY, JOHN HAWTHORNE, & BRIAN WEATHERSON. 2005. Epistemic modals in context. In G. PREYER & G. PETER (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth*, pp. 131–168. Oxford: Oxford University Press.
- VON FINTEL, KAI & ANTHONY S. GILLIES. 2008. CIA leaks. *The Philosophical Review* 117, 77–98.
- GIBBARD, ALLAN. 1990. *Wise choices, apt feelings*. Cambridge: Harvard University Press.
- GIBBARD, ALLAN. 2003. *Thinking how to live*. Cambridge: Harvard University Press.
- HANSSON, BENGT. 1969. An analysis of some deontic logics. *Nous* 3, 373–398.
- HARE, R.M. 1952. *The language of morals*. Oxford: Oxford University Press.
- HINTIKKA, JAAKO. 1969. Semantics for propositional attitudes. In J.W. DAVIS, D.J. HOCKNEY, & W.K. WILSON (Eds.), *Philosophical logic*, pp. 21–45. Dordrecht: Reidel.
- KAMP, HANS. 1971. Formal properties of ‘now’. *Theoria* 37, 227–274.
- KAPLAN, DAVID. 1989. Demonstratives. In J. ALMOG, J. PERRY, & H. WETTSTEIN (Eds.), *Themes from Kaplan*, pp. 481–563. Oxford: Oxford University Press.
- KRATZER, ANGELIKA. 1981. The notional category of modality. In HANS JURGEN EIKMEYER & HANNES RIESER (Eds.), *Words, worlds, and contexts. New approaches in word semantics*, pp. 38–74. Berlin: de Gruyter.

- KRATZER, ANGELIKA. 1991. Modality/Conditionals. In A. VON STECHOW & D. WUNDERLICH (Eds.), *Semantics: An international handbook of contemporary research*, pp. 639–656. Berlin: de Gruyter.
- LASERSOHN, PETER. 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28, 643–686.
- LEWIS, DAVID. 1970. General semantics. *Synthese* 22, 18–67.
- LEWIS, DAVID. 1973. *Counterfactuals*. Cambridge: Harvard University Press.
- LEWIS, DAVID. 1978. Reply to McMichael. *Analysis* 38, 85–86.
- LEWIS, DAVID. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8, 339–359.
- LEWIS, DAVID. 1980. Index, context, and content. In STIG KANGER & HELLE OHMAN (Eds.), *Philosophy and grammar*, pp. 79–100. Holland: D. Reidel.
- LEWIS, DAVID. 1981. Ordering semantics and premise semantics for counterfactuals. *Journal of Philosophical Logic* 10, 217–234.
- LYONS, JOHN. 1977. *Semantics*, vol. 2. Cambridge: Cambridge University Press.
- MACFARLANE, JOHN. 2007. Relativism and disagreement. *Philosophical Studies* 132, 17–31.
- MOORE, G.E. 1903. *Principia ethica*. New York: Cambridge University Press.
- MOSS, SARAH. 2012. Epistemology formalized. *Philosophical Review* .
- PARFIT, DEREK. 2011a. *On what matters*, vol. I. Oxford: Oxford University Press.
- PARFIT, DEREK. 2011b. *On what matters*, vol. II. Oxford: Oxford University Press.
- ROSS, W.D. 1930. *The right and the good*. Oxford: Oxford University Press.
- SCHLENKER, PHILIPPE. 2003. A plea for monsters. *Linguistics and Philosophy* 26, 29–120.
- SCHROEDER, MARK. 2009. Hybrid expressivism: Virtues and vices. *Ethics* 119, 257–309.

- SCHROEDER, MARK. 2010. *Noncognitivism in ethics*. New York: Routledge.
- SCHROEDER, MARK. 2011. Philosophy of language for metaethics. In GILLIAN RUSSELL & DELIA GRAFF FARA (Eds.), *Routledge Companion to the Philosophy of Language*. New York: Routledge.
- SHARVIT, YAEL. 2008. The puzzle of free indirect discourse. *Linguistics and Philosophy* 31, 353–395.
- SILK, ALEX. 2012a. Negation and disjunction in expressivist semantics. MS, University of Michigan.
- SILK, ALEX. 2012b. Standards and scoreboards. MS, University of Michigan.
- STALNAKER, ROBERT. 1970. Pragmatics. In STALNAKER (1999), pp. 31–46.
- STALNAKER, ROBERT. 1978. Assertion. In STALNAKER (1999), pp. 78–95.
- STALNAKER, ROBERT. 1984. *Inquiry*. Cambridge: MIT Press.
- STALNAKER, ROBERT. 1988. Belief attribution and context. In STALNAKER (1999), pp. 150–166.
- STALNAKER, ROBERT. 1997. Reference and necessity. In CRISPIN WRIGHT & BOB HALE (Eds.), *Blackwell Companion to the Philosophy of Language*, pp. 534–554. Oxford: Basil Blackwell.
- STALNAKER, ROBERT. 1999. *Context and content: Essays on intentionality in speech and thought*. Oxford: Oxford University Press.
- VON STECHOW, ARNIM. 2003. Feature deletion under semantic binding: Tense, person, and mood under verbal quantifiers. *Proceedings of the North East Linguistics Society* 33, 397–403.
- STEPHENSON, TAMINA. 2007. *Towards a theory of subjective meaning*. Ph.D. thesis, MIT.
- STEPHENSON, TAMINA. 2010. Control in centered worlds. *Journal of Semantics* 27, 409–436.
- SWANSON, ERIC. 2012. The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic* pp. 1–28.

- VELTMAN, FRANK. 1976. Prejudices, presuppositions, and the theory of conditionals. In JEROEN GROENENDIJK & MARTIN STOKHOF (Eds.), *Amsterdam papers in formal grammar*, vol. 1, pp. 248–281. Central Interfaculteit, University of Amsterdam.
- WEATHERSON, BRIAN. 2008. Attitudes and relativism. *Philosophical Perspectives* 22, 527–544.
- VON WRIGHT, GEORG HENRIK. 1963. *Norm and action: A logical inquiry*. London: Routledge and Kegan Paul.
- YALCIN, SETH. 2007. Epistemic modals. *Mind* 116, 983–1026.