

Review of *The Meaning of 'Ought'* by Matthew Chrisman

Billy Dunaway, University of Missouri—St Louis

Forthcoming in *The Journal of Philosophy*

In *The Meaning of 'Ought'*, Matthew Chrisman draws on tools from formal semantics, philosophy of language, and traditional metaethics to arrive at an ambitious and suggestive view about a term that is central to ethics, and normative theorizing more generally. The term in question is 'ought' (as the title of the book makes clear) and Chrisman advances a package of theses designed to challenge several received views on the topic.

One element of the package concerns the truth-conditions of sentences containing 'ought'. Philosophers and linguists share general agreement that the correct picture will incorporate the framework from Angelika Kratzer: 'ought' is a modal operator, which produces truths (or falsities) when attached to a declarative sentence *S*.¹ In its most basic form, the Kratzer framework says that *ought S* is true just in case *S* is true in all of the relevant worlds that rank highest according to a contextually supplied ordering source. This is designed to account for the different "flavors" of 'ought': in some contexts it has an *epistemic* flavor, and 'ought' is evaluated against worlds ranked by an ordering source according to what is known. In others it has a *moral* flavor, and is evaluated against worlds ranked by an ordering source according to what ought to happen, morally speaking. (There are of course many other flavors: prudential, teleological, etc.) Thus we might truly say, in a context where the epistemic ordering is salient,

¹ Angelika Kratzer, 'What 'must' and 'can' must and can mean', *Linguistics and Philosophy* 1.3 (1977), 337-355.

“Trump ought to be lying to a reporter now”. Roughly we mean here is that, given our background knowledge about Trump’s speaking engagements and assertoric propensities, he is lying to a reporter at this moment. But of course if the context is one where the *moral* ranking on worlds is salient, this sentence is false. In the morally best worlds, Trump isn’t lying at all right now.

So much for the received orthodoxy about the truth-conditions for ‘ought’. Chrisman does not wish to depart from it entirely, but makes a substantial foray into the domain of formal semantics as he claims that it needs non-trivial modification. His central motivation stems from the distinction between the so-called “agential” and “non-agential” readings of sentences containing ‘ought’. To use Chrisman’s example, we might truly say “Larry ought to win the lottery”, in a context where Larry is the only holder of a ticket in the relevant lottery who has fallen on financial misfortune owing to no fault of his own.² Of course the nature of a (non-rigged) lottery makes it such that Larry has no means to exercise his own *agency* in a substantial sense to bring it about that he wins, once the ticket has been bought and before the winning numbers drawn. Here ‘ought’ receives a *non-agential* reading. But in other cases ‘ought’ seems to receive a different reading. We might peer in from a well-disguised hiding place on our friend Bill’s date with Lucy, when Bill is a romantic novice and in need of advice (which unfortunately we cannot provide in real-time at the risk of scaring off Lucy). And here we might truthfully utter the sentence “Bill ought to kiss Lucy”. By this we don’t simply mean that what ought to happen is that Bill kisses Lucy, no matter what the causal route that brings about the kissing. Instead we mean that Bill *ought to bring it about* that he kisses Lucy—the romantic gesture of Bill

² Matthew Chrisman, *The Meaning of ‘Ought’*, (New York: Oxford, 2016), 110.

initiating the kissing is part of the most desirable outcome, given the circumstances. This is an *agential* reading of 'ought'.

Chrisman argues that the standard Kratzer semantics cannot accommodate the distinction between agential and non-agential 'ought's. And the alleged culprit here is that the Kratzer framework provides truth-conditions by a ranking on *worlds*. So Chrisman's first positive proposal is that the modal 'ought' ranks inputs that are more fine-grained than possible worlds.

This is questionable. Sets of possible worlds are perfectly capable of representing the kind of agency that Bill exercises in kissing Lucy (if things go as desired), but Larry does not exercise in winning the lottery. In some worlds the relevant kind of agency is exercised, and in others it is not. So an agential 'ought' should rank highest only worlds where Bill does the kissing by exercising his full agential capacities. By contrast, a non-agential should include worlds in the ranking where the kissing happens, but not in virtue of any action on Bill's part—just include some worlds where Lucy kisses Bill as he checks his smartphone in the set of worlds which constitutes the proposition 'ought' attaches to. That is a non-agential reading. Likewise, since the best worlds for Larry include worlds where he wins without exerting any causal pressure on the outcome of the lottery, the natural interpretation of the relevant sentence is a non-agential one.

To be clear, there are many issues with providing a compositional semantic theory that produces the correct truth-conditions for sentences containing agential and non-agential readings of 'ought'. But readers are likely to demur when Chrisman argues that it is the source of the

problem lies in the traditional framework's use of sets of possible worlds, rather than more fine-grained proposition-like devices.

Here is a second component of Chrisman's package of views: regardless of how (or whether) we should revise the Kratzer truth-conditions for normative readings of 'ought', this is not a source of division among traditional meta-ethical views. Take the familiar dispute between *expressivists* and their opponents, the *descriptivists*. The latter hold that the meaning of normative sentences is to be explained in the same way as other ordinary sentences. Roughly, the explanation of the sentence's meaning fundamentally proceeds by specifying how the sentence says reality is.

Expressivists, by contrast, think that normative sentences are special. To explain their meaning, we need to point to the state of mind that they express.³

For Chrisman, the debate between descriptivists and expressivists need *not* be construed as a debate at the level of semantics, surfacing in a debate over whether the Kratzer-style truth-conditions (or anything like them) are correct. Instead, expressivists can *agree* with descriptivists that the truth-conditions are roughly as Kratzer says they are. To make room for the traditional meta-ethical debate, Chrisman deploys a distinction from the philosophy of language: this is the distinction between *semantics* and *meta-semantics*. Roughly, the former is the project of specifying the conditions under which a sentence is true. Kratzer's truth-conditions lie squarely in the domain of semantics. *Meta-semantics* is the project of saying *why* a given sentence has the truth-conditions that it in fact has, rather than something else. (For instance: *why* does 'Bill kisses Lucy' have truth conditions that make it true when Bill presses his lips to hers, but not true

³ Cf. Allan Gibbard, *Thinking How to Live*, (Cambridge: Harvard, 2003).

when he ignores her in favor of updating his social media profile on his smartphone?

Presumably part of the answer will involve facts about how English speakers use the word 'kiss': they don't tend to call solitary acts with a smartphone 'kissing'. Though perhaps some amount of semantic shift is possible.)

Proponents of these meta-ethical views can accept the Kratzer semantics for 'ought'. What they disagree over the following question: is the semantics correct because normative uses of 'ought' aim at describing reality? Or is it correct because normative sentences express motivational states of mind? This is the meta-semantic question.

There is an important ambiguity in the 'can' in the relationship between the semantics and meta-ethics. Adherents to different meta-ethical positions, Chrisman says, can agree at the semantic level about the truth-conditions of normative sentences. What he is surely right about is the fact that there is no straightforward entailment from one of the meta-ethical positions to the *denial* of the claim that normative sentences have such-and-such Kratzer-like truth-conditions. There is a well-developed literature on how expressivism is consistent with normative truth.⁴ So the mere fact that the semantics is stated in terms of truth-conditions doesn't directly entail that the meta-ethics isn't expressivistic. But a meta-semantics is supposed to be more than merely consistent with semantic claims: it is supposed to *explain* why the

⁴ See, e.g., James Dreier "Meta-ethics and the problem of creeping minimalism," *Philosophical Perspectives* 18(1) (2004), 23-44, and Billy Dunaway, "Expressivism and Normative Metaphysics," *Oxford Studies in Metaethics* vol. 11, ed. Russ Shafer-Landau, (New York: Oxford, 2016), 241-264.

semantic facts are such-and-such, rather than another way. Chrisman appears to mostly focus on the consistency issue between Kratzer-style semantics and non-descriptivist meta-metaethics. But if meta-ethical positions are treated as meta-semantical, they should give explanatory insight into the semantics, a question Chrisman devotes significantly less attention to.

Here is an illustration. A philosopher might propose, as a meta-semantical claim, that 'ought' applies to those actions that can be performed at a location whose latitude and longitude are both canonically specified by numbers whose fourth decimal is an odd number. This cannot be rejected on the basis of straightforward inconsistency with the accepted truth-conditions for the term. So there is a narrow sense in which the meta-semantics is *consistent* with the standard truth-conditions. But the explanatory role for the meta-semantics will go unfulfilled: for instance it doesn't explain why someone who consciously used the term according to rules that are incompatible with the meta-semantics would still count as meaning the same thing as a other users of normative terms.

This meta-semantic proposals is, to reiterate, crazy. I introduce it simply to emphasize that we need the explanatory relationship between the meta-semantics and first-order truth-conditions to adequately diagnose the craziness. We can't dismiss it on the grounds that it is logically inconsistent with Kratzer truth-conditions alone. Narrow logical consistency is cheap, and not the full story.

A final contribution from Chrisman comes in the form of a third alternative to the traditional descriptivism/expressivism dichotomy. While meta-ethicists have traditionally treated these as the only relevant alternatives, Chrisman claims that there is another: *inferentialism*. Since descriptivism and expressivism are meta-semantic proposals that are consistent with Kratzer-style truth-conditions, inferentialism should be treated similarly. It is a meta-semantic proposal that is consistent with the orthodox truth-conditions. The difference is that the meta-semantics does not work with fundamental representations of reality (descriptivism) or expression of motivational states of mind (expressivism), but instead endorsement of quasi-logical connections between 'ought' and other linguistic expressions or intentions.

Inferentialism is famously championed in the work of Robert Brandom and Wilfrid Sellars, and Chrisman follows them closely in introducing the view. This does the service of presenting inferentialism as a historically important, if frequently ignored, option in meta-semantic space. But it also does a disservice, since historical introductions of inferentialism frustratingly focus on picturesque presentations of a large philosophical picture, and infrequently attend to the details which provide the resources to derive concrete predictions of the truth-conditions of terms with an alleged inferentialist meta-semantics. This is directly relevant to the explanatory relationship between meta-semantics and truth-conditional semantics. The inferentialist can plausibly claim that her meta-semantics will be consistent in the narrow sense with the standard semantics for the normative 'ought'. But the explanatory question remains a mystery. Without moving substantially beyond the hand-wavy tradition of inferentialism, Chrisman leaves us with little traction to assess the explanatory prospects of inferentialism as a serious contender in the meta-semantic landscape.