

Chrisman, Mathew. *The Meaning of ‘Ought’: Beyond Descriptivism and Expressivism in Metaethics*. Oxford: Oxford University Press, 2015. Pp. 280. \$65.00 (cloth).

Chrisman’s recent book develops an account of the semantics and metasemantics of the modal verb and normative expression *par excellence*, ‘ought’. Arguing for his preferred semantic rule for ‘ought’ occupies the majority of the book, chapters 2-5. Chapter 1 gives a brief tour of the argumentative landscape, including a valuable discussion of the distinction between semantics and metasemantics, chapter 6 sketches Chrisman’s own metasemantic picture, and in chapter 7, we get a brief tour of additional issues such as the connection of his view to issues in moral psychology and epistemology as well as discussions of other normative expressions like ‘wrong’ and ‘good’.

Chrisman emphasizes throughout the difference between giving a metasemantics for normative language and giving a compositional semantics for it. For Chrisman, a (formal) *compositional semantics* is a theory of the content of complex constructions in terms of the content of simpler expressions. A *metasemantics* interprets notions like *content* and mechanisms like *composition* in terms of the function of the relevant bit of language. Model-theoretic possible worlds semantics —“worlds” really being mathematical structures of some type—is a compositional semantics; that possible worlds *represent* ways things could have been and that sets of possible worlds *represent* states of information are part of a representationalist metasemantic interpretation of it. Compositional semantics imposes some constraints on metasemantics—since semantic structure partially explains our understanding of language, a proper semantics will demand of a metasemantics the ability to interpret this structure in a reasonable fashion. Even so, Chrisman argues compellingly that standard compositional semantics takes no stand on which metasemantic picture is correct.

Chrisman claims that some accounts of normative language, like expressivism, are best interpreted as metasemantic views. I am broadly sympathetic, though it is combinations of semantics and metasemantics—*interpreted* compositional semantics—that are, have been, and should be of primary philosophical interest. Emphasizing the distinction between metasemantics and formal semantics is most important as a corrective to the perennial mistake of treating algebraic structures, like bare possible worlds semantics, as accounts of meaning instead of structure on which to hang a proper account of linguistic meaning. Chrisman’s metasemantic aim is to develop an alternative to orthodox representationalist and expressivist options. His own view is that ought-talk makes explicit certain positions in a “space of implications” (187). As I will suggest below, this metasemantic picture faces trouble in providing a natural interpretation of a formal semantics, at least given the picture’s current stage of development.

His preferred *semantics* takes ‘ought’ to be a necessity modal. ‘Barry ought to be pirouetting’ is semantically composed of an operator *ought* applied to the proposition *Barry is pirouetting* (the ‘prejacent’). Following Kratzer, Chrisman presumes that context supplies two aspects of the meaning of ‘ought’. First, we only consider certain salient possibilities when we evaluate ‘ought’-claims. When we’re evaluating whether Barry ought to be pirouetting, we’re usually pointedly ignoring possibilities where his legs will shortly fall off. We thus presume that context specifies a presupposed body of information, holding in all salient worlds, which we can represent formally as a function f from a world t to a set of relevant worlds $f(t)$. Likewise, ‘ought’s come in different *flavors*: moral, prudential, evaluative, teleological, and epistemic. Again, following Kratzer, Chrisman presumes that context specifies a body of information about how to evaluate the relevant flavor of ‘ought’ which in turn induces a partial order g on the salient worlds.

I will use ‘correct’, instead of Chrisman’s ‘1’, as a metasemantically neutral analogue for ‘true’. Correctness, for simple non-normative propositions, is treated in the usual fashion as truth-at-a-world. Chrisman takes as half of his analysis of ‘ought’-claims (represented here as *Op*) that:

(R10d) *Op* is correct at t just in case p is correct at every $g(t)$ -maximal world w in $f(t)$.

That is, when *Barry is pirouetting* in every salient possibility which is maximal (not maximum!) in the $g(t)$ -ranking. Presumably, for ‘Barry ought to be pirouetting’, g ranks worlds according to their desirability. For epistemic ‘ought’-claims, like ‘He ought to be there now’, g ranks worlds according to something like their likelihood.

Chrisman is sensitive to the difference between *agentive* ‘oughts’ like ‘You ought to give more to charity’ and *non-agentive* ‘oughts’ like ‘He ought to be home by now.’ To analyze the former, Chrisman makes use of Castañeda’s imperatival analogue to propositions, *practitions* (Castañeda 1975). These can be thought of as connecting an action to an agent in terms of a practical, not a predicative, copula— $\langle \text{Barry to } \varphi \rangle$ instead of $\langle \text{Barry is } F \rangle$. Using pairs of worlds and sets of norms (indicated $\langle w, n \rangle$), instead of simply worlds, Chrisman explains a *practition* p as being correct at $\langle w, n \rangle$ just in case n demands p at w . For example, if t is a world in which Devrim’s house is on fire and k demands that Devrim care for her property, then $\langle \text{Devrim to go home} \rangle$ is correct at $\langle t, k \rangle$.

Chrisman extends the analysis of ‘ought’ to *practitions* by taking g to rank world-norm pairs in $f(t)$, the set of world-norm pairs specified by background conditions. Then:

(R10p) Op , for p a practition, is correct at t just in case p is required by n at w at every $g(t)$ -maximal $\langle w, n \rangle$ in $f(t)$.¹

So, ‘Devrim (morally) ought to go home’ is correct at t just in case p is required at every morally best pair $\langle w, n \rangle$ in the set of world-norm pairs specified by the background conditions in t . (R10d) and (R10p) together yield:

(R10) Op is correct at a world t just in case p is correct at every $g(t)$ -maximal $\langle w, n \rangle$ in $f(t)$

The discussion of semantics for ‘ought’ which precedes the discussion of R10 is useful, if not entirely surprising to those familiar with recent work on modal verbs. Chrisman’s main objections to alternatives are that they (a) stumble on dilemmas about obligation, (b) lack the ability to capture both agentive and non-agentive senses of ‘ought’ in a natural and non-disjunctive fashion, and (c) have trouble with non-ideal contexts. Chrisman’s analysis, designed explicitly to satisfy (b), has the additional virtue of avoiding both (a) and (c) in a natural fashion. Following Kratzer, the ranking function g might rank highest world-norm pairs which are non-ideal, solving (c). (a) can be solved for *agentive* ‘ought’ claims since g might rank highest world-norm pairs whose norms demand inconsistent actions.²

This pushes the question of whether dilemmas exist out of formal semantics and into the domain of substantive ethics, as we are not *required* by Chrisman’s solution to think that any g *would* so rank world-norm pairs. This is a nice result in keeping with Chrisman’s goal of distinguishing substantive ethical and metasemantical projects from the project of developing a linguistically adequate compositional semantics for ‘ought’. His solution does not generalize to non-agentive ‘oughts’ since their prejacent are typically *propositions* which are evaluated for truth at possible worlds. Some will find this a good result as agentive ‘oughts’ are the natural source of normative dilemmas. Some will want more.

Chrisman (149) claims as a crucial advantage of his semantics that R10, applying to both propositions and practitions, is not an ambiguity or polysemy analysis, unlike accounts which explicitly provide distinct lexical rules for different uses of ‘ought’ (§2.3, §5.2). However, the ability to formulate a compositional rule which applies to both practitions and propositions is not sufficient to avoid polysemy. After all, some

¹ I ignore some complications here about how the ranking function g interacts with n

² There is a slight technical problem here. According to Chrisman, g ranks world-norm pairs in terms of how well the relevant standard is met *when* the relevant agent does what is required of them by n (148). However, if n demands impossible things, there is no case where the agent does all of what is required of them. Presumably, g should evaluate how well the consequences of what the relevant agent does meets the standard, whether or not they satisfy *all* of n ’s demands.

‘ought’ claims are really evaluated by R10 in terms of what is required by a set of norms whereas others are really evaluated in terms of what is true at a world. R10 is thus intuitively, if covertly, polysemous as it is naturally factorable into two distinct, but related, lexical entries.

Variations on R10 are possible as well. It isn’t clear to me, for example, that Chrisman’s detour through Castañeda-style practitioners is required.³ As far as the linguistic argument of the book is concerned, we could posit a semantic distinction between suitably structured *agentive* and *non-agentive* propositions, replacing R10p with something like:

(R10p_a) *Op*, for *p* an agentive proposition, is true at *t* just in case the agent of *p* is required by *n* to bring it about that *p* in *w* at every *g(t)*-maximal $\langle w, n \rangle$ in *f(t)*.

Chrisman might claim that R10p_a and R10d together constitute a less unified analysis of ‘ought’ than R10. But, even if true, this seems outweighed by the theoretical gain in using only propositions in analyzing ‘ought’. ‘Barry ought to be pirouetting’ doesn’t intuitively encode imperatival content of the type practitioners are meant to capture, unlike the statement of intention ‘I will go to Istanbul’, so making use of practitioners to capture the former requires, perhaps unnecessarily, disrespecting or explaining away these intuitions.

Overall, chapters 2 through 5 constitute an extremely useful and slightly idiosyncratic introduction to the semantic analysis of ‘ought’ in both philosophy and linguistic semantics as well as a fairly convincing abductive argument in favor of something like R10 as R10 (or a close cognate) looks significantly less theoretically costly than alternatives. I have not detailed his criticisms of the alternatives, but the reader will find the discussion sensible and the landscape well surveyed. The reader is cautioned to read the footnotes carefully since they contain much substantive material.

Turning now to metasemantics, Chrisman quickly dismisses two alternatives to his view: global representationalism and expressivism. The former interprets R10’s structure as more or less directly representing the world as containing moral properties, possible worlds, and so on. I am broadly sympathetic to his reasons for rejecting this view—it is difficult to accept that R10 commits us to actual possibilities and normative properties—though this is more a matter of philosophical taste than principled argument. Chrisman is more sympathetic to a local version of representationalism which acknowledges that some expressions, such as ‘ought’, encode what Chrisman glosses as ‘conceptual manipulations’ instead of representing the world. He does complain that the notion of a conceptual manipulation is underspecified and

³ Much as I appreciate seeing someone make serious use of Castañeda’s annoyingly neglected work.

vague, though both expressivism and his own preferred metasemantics seem to be local representationalist views where ‘conceptual manipulation’ is further specified. More would have been appreciated here; in particular, exploration of other local representationalist options would have been illuminating.

Chrisman glosses expressivism as an ideationalist metasemantics—semantic rules like R10 are articulations of how we should think when uttering ‘ought’-claims—along with the claim that *what* we ought to think when uttering ‘ought’ claims involves directive instead of descriptive content. In my view, this recent trend of interpreting metaethical expressivism as an ideationalist view is a serious mistake, conflating a program in linguistic (meta-)semantics dating back to Grice with the radical picture of the constitutive function of normative thought and talk expressivists have championed. Expressivism claims that the function of normative discourse is to express our evaluative or directive attitudes, ideationalism claims that the content of normative claims is derived from what we should think when so claiming. Expressivism, by itself, makes no claims about the content of normative claims and ideationalism, by itself, makes no claim about the function of normative discourse. The two should be separated, even if some theorists accept both.

Chrisman briefly raises (but does not discuss in detail) some standard issues for this kind of expressivist metasemantics, such as the Frege-Geach problem. More interesting is his suggestion that expressivists have trouble accommodating the distinction between agentive and non-agentive readings of ‘ought’-claims since not all ‘ought’-thoughts are intuitively directive. Just consider ‘Erdoğan ought to be in jail.’ This is a nice problem, but one affecting primarily the plan-oriented expressivism developed by Gibbard. Those who treat the function of normative thought and talk as the expression of *evaluative* attitudes like approval and disapproval instead of *directive* attitudes like plans or intentions will welcome this problem for their competition.

Chrisman’s own metasemantics takes a middle path between global representationalism and expressivism, preserving the valuable aspects of each. On the one hand, he wants to be able to respect the most sophisticated work in linguistic semantics about the meaning of ‘ought’. This is easiest to do on the local representationalist metasemantic picture since it typically interprets non-normative propositions standardly as truth-apt representations of the world. On the other hand, he wants his metasemantics to be resolutely antidescriptivist about normative claims and respect the motivational and psychological aspects of normative thought taken seriously by expressivists. The result is a version of inferentialism where certain bits of language *make explicit* commitments to implicational connections. ‘ought’ claims are “an articulation of what one is committed to in virtue of using the [‘ought’-sentence] to make an assertion in ordinary discursive practice” (186).

Semantic rules like R10 articulate “a position in the space of implications”—a set of implications—which we are committed to in virtue of our use of ‘ought’ claims. That is, ‘ought’ gets “its content from being usable to acknowledge inferential connections between more basic items rather than to refer to things in the world” (197). Chrisman only give a couple of tentative examples of such connections and neglects the difference between the action of *inference* and the relation of *implication*, so it is difficult to evaluate this part of his metasemantic view. Still, I worry that there is far too little agreement about what we are committed to doing when accepting ‘ought’ claims to support an inferentialist metasemantics for ‘ought’, unlike the logical particles ‘and’, ‘or’, and ‘not’.

In contrast to the global representationalist view, we are never told explicitly how to interpret the semantic machinery involved in R10 in terms of the inferentialist metasemantics. On the representationalist picture, there is a specifiable *real content* of ‘ought’ claims—representations of normative reality—which interpret the machinery in R10. Similarly, for the expressivist, the real content of ‘ought’-claims, so far as there is such a thing, can be described in terms of plans or endorsements. In both cases, it is possible to explain why R10 is the way it is in terms of the underlying metasemantics (though adequately matching an expressivist metasemantics to something like R10 is non-trivial.)

We are not told how the inferentialist metasemantics interprets mechanisms like the ranking function g or ‘being required by n ’. Since it isn’t clear how to interpret R10 in terms of an underlying space of implications, it’s not clear what the inferentialist picture takes the real content of ‘ought’-claims to be. One might worry that the gap between the favored formal semantics and an inferentialist metasemantics is simply too wide. This is the reason that inferentialists in the logical tradition—such as the intuitionists—have often treated formal semantics as a useful tool for modeling real content, but one which plays no substantive role in the explanation of our actual understanding of the meaning of the relevant fragments of language. Chrisman seemingly wants his formal semantics to do more, but it is unclear exactly how to connect it to his favored metasemantic picture. If it is to do more, then the formal machinery needs precise interpretation in terms of the underlying metasemantics.

One potential way of filling this gap without directly interpreting the machinery in R10 treats notions like ‘correct’ as insubstantial placeholders, using a few constitutive rules or entire sets of implications governing the behavior of expressions like ‘ought’ (207) and the requirement that rules or implications be correctness-preserving to *implicitly define* the meaning of ‘ought’ and thereby specify rules like R10. Historically, this approach has been beset with tremendous difficulties, even for logical expressions like ‘or’ and ‘not’. The problem, in brief, is that it is difficult to pin down a *unique* semantic rule like R10 in terms of an underlying set of implications or constitutive rules. Even in the simplest cases, any of many pairwise inconsistent semantic rules validates the relevant rules or implications. For example, without disrespect-

ing the classical consequence relation, we can either treat a disjunction with two false disjuncts as sometimes true or, alternatively, as always false. If Chrisman favors the approach just described, then it would be useful to have some discussion of how he intends to finesse this problem. (For the logical case, see Garson (2013) and the references therein.)

Given this lack of detail, it is hard to evaluate the overall plausibility of Chrisman's metasemantic suggestion. I have focused on places I disagree with Chrisman, but there is also much to agree with. Chapters 2-5 are a useful propaedeutic to modal verbs, especially for metaethicists unfamiliar with work in linguistic semantics, as well as a corrective to the idea that language-oriented metaethics can ignore this work. Chrisman's emphasis on the distinction between semantics and metasemantics is both useful and timely. If his metasemantic picture is sketchy, it is nonetheless intriguing and worthy of further detailed development. I would recommend this book to anyone interested in the debate about the meaning of normative expressions. It may frustrate, but the reader will be well compensated with probative questions and worries for their own favorite account of 'ought'.

References

Castañeda, H. *Thinking and Doing: The Philosophical Foundations of Institutions*, D. Reidel Pub. Co. 1975.

Garson, J. *What Logics Mean: From Proof Theory to Model-Theoretic Semantics*, Cambridge University Press, 2013.

Jack Woods
University of Leeds