13

Metanormative Theory and the Meaning of Deontic Modals

Matthew Chrisman

1 Introduction

Metanormative theory is, in part, about the meaning of sentences such as "One ought always act for reasons for which one could consistently allow everyone else to act as well," and "One should proportion degree of belief solely in accordance with evidence." There are other normative terms, of course, but 'ought' ($\approx$ 'should'), is clearly one of the handful of core normative terms, and I will focus on it here.[1]

Given this focus, here is a way to put an idea that I think many philosophers find attractive: ought-claims prescribe possible action, thought, and feeling, rather than describing how things actually stand in reality. Arguably, this idea underwrites the popularity of the is/ought divide, explains some of the attraction in Moore's Open Question Argument, and motivates some of the interest philosophers have had in noncognitivist, emotivist, prescriptivist, projectivist, and expressivist accounts of normative thought and discourse.

Here is another popular line of thought that has sometimes seemed to be in tension with that initial idea: A central part of the best overall account of the meaning of declarative sentences will explain how their truth conditions can be derived from the semantic values of their basic components and their logical form. The word 'ought' can, of course, figure in nondeclarative sentences, but the ought-sentences primarily at issue in metanormative theory are grammatically declarative. This implies that our overall theory of meaning should provide a compositional assignment of truth conditions to ought-sentences as a part of its treatment of declarative sentences more generally. This seems necessary for accounting for the fact that it is obviously meaningful to embed ought-sentences in propositional contexts, such as under the truth

---

[1] Returning only very briefly at the end to reflect on the broader issue of how my conclusions here might bear on similar issues regarding other normative and evaluative vocabulary.

predicate, under epistemic modals, in the antecedent of conditionals, in disjunctions, or in the context of "S believes/knows/doubts/fears/hopes that . . .," etc.

One reason this line of thought has seemed in tension with the distinction many philosophers have wanted to draw between 'is' and 'ought' is that proponents of the most prominent family of antirealist metanormative views have sometimes suggested that normative sentences including ought-sentences are not truth-apt. Or they have suggested that normative sentences have truth conditions but only in a "deflationary sense" supposedly incompatible without reliance on truth conditions for substantive and informative explanations of the meanings of these sentences.

However, it's important not to underestimate the challenges facing this idea. Ought-sentences embed just like other declaratives in propositional contexts, and our best account of the compositional function of these other declarative sentences is (arguably) one that seeks a perspicuous display of the functional interaction of the nodes of the correct tree-like representation of the logical form of each sentence, where these functions are ultimately functions from the semantic values of the parts of the sentence to the value *true*—that is to say, a theory which assigns truth conditions to these sentences. Any theory that would propose otherwise for ought-sentences is under extremely strong pressure to propose otherwise for *all* declarative sentences. At this level of generality, however, the truth-conditionalist approach to the semantics of declarative sentences is widely (though not universally[2]) assumed to be the best game in town for articulating the compositional structure of declarative sentences. This is why many philosophers are strongly inclined to think we should assign ought-sentences truth conditions, and not as a mere deflationary afterthought but as a central part of the explanation of how their meaning is a function of the meanings of their parts.

Does that mean we have to choose: *either* accept the Herculean task of developing a new approach to compositional semantics that eschews the successes of truth-conditionalist approaches even for nonnormative language *or* treat ought-claims as describing reality after all (and "locate" the normative properties that would make sense of that)? Here, I want to argue against thinking that this is an important choice point at the intersection of metanormative theory and compositional semantics. To that end, I pursue three main tasks in this chapter. First, I explain a specific version of the semantic rule for 'ought' that I favor. This is a refinement of the standard treatment of 'ought' as a context-sensitive necessity modal. Secondly, I explain two ways we might interpret the results of truth-conditional semantics: a more familiar "representationalism" and a less familiar "inferentialism." I view these as different

---

[2] Nothing I say in this paper will engage substantively with the question of whether a truth-conditionalist approach to compositional semantics is preferable *in general* to some dynamic alternative such as one based on Discourse Representation Theory (Kamp, 1981), File Change Semantics (Heim, 1982), or Update Semantics (Veltman, 1985, 1996). Here I shall simply assume (along with most metanormative theorists) that a truth-conditionalist approach is the default position for nonnormative declaratives, and so the interesting question is what to make of the fact that it is also appropriate for normative declaratives.

kinds of *metasemantic* accounts of what it is in virtue of which declarative sentences have the contents that they do.[3, 4] I want to explain these views in order to argue that, with my semantic rule for 'ought' in hand, both metasemantic accounts can respect the apparent prescriptivity of (some) ought-claims and both can accept relatively standard compositional truth conditions for ought-claims without committing to the idea that these claims describe how things actually are in reality. What distinguishes them is *how* they respect these ideas, not *whether* they can do so. Finally, I discuss the "ideationalist" thesis that sentences are primarily vehicles for expressing our ideas and so get their truth conditions from their conventional expressive potentials. This may be an attractive home for some sophisticated contemporary expressivist views in metanormative theory. I compare them with inferentialism because I think it is not completely clear how such sophisticated versions of expressivism fit in. Depending on how they are developed, they might be viewed as compatible with and so perhaps usefully combined with inferentialism; or they might constitute a third competing metasemantic view. I hope that comparison of aspects of this somewhat more familiar view to the inferentialist view I develop earlier in the chapter will help to clarify what is at stake.

2 'Ought' and Truth Conditions

As with many other modal auxiliary verbs in English, ought-sentences come in diverse flavors, e.g.

---

[3] By using language of "in virtue of" here, I mean to indicate that they are very general *grounding* explanations that seek to identify the kinds of nonsemantic facts that undergird the semantic facts postulated by empirical semantics. Compare Charlow (2014a, p. 662) and Yalcin (2014, p. 18). Some (e.g. Speaks, 2010) may prefer the term "foundational theory of meaning" or "theory of the foundations of meaning" for what I am calling *metasemantics*. Yalcin (2014) surveys arguments for a distinction between *content* and *semantic value*, where the former is (at least) what is said by uses of a sentence and what grounds the intentionality and causal efficacy of mental states expressed by such uses, and the latter is what is attributed to sentences to explain various semantic facts, e.g. about productivity of language, entailment relations between sentences, the un/acceptability or un/interpretability of particular sentences, etc. Given this distinction, there are two domains about which we might ask second-order questions: we can ask about what it is in virtue of which speech-acts and the mental states they express have the *contents* that they do, and we can ask about what it is in virtue of which sentences have the *semantic values* that they do. While I am highly sympathetic to this distinction, observing it in what follows would have unnecessarily complicated an already complicated discussion. It may be the case that representationalism and inferentialism are best viewed as second-order accounts of content rather than—or at least in addition to—semantic value. However, I write here under the simplifying assumption that "metasemantics" is about what I sometimes call "semantic content," which is meant to cover both.

[4] Perez Carballo (2014) usefully distinguishes a "hermeneutic" question about how to interpret the formal specifications of truth conditions in truth-conditional semantics from an "explanatory" question about what it is that explains the (approximate) correctness of the formal specifications of truth conditions for a sentence. I agree that these are distinct, but I also think the metasemantic views I will discuss here are all at a high enough level of generality that they can be seen as offering package answers to these questions. I'm mainly interested in the explanatory question, but sometimes I will ask about how we are to *interpret* the results of truth-conditional semantics. By this, I don't mean just what hermeneutic gloss of the formalism is appropriate but also what underlying facts about language and language users explain what it is in virtue of which sentences have the truth conditions that they have.

(1)   We ought to do more to relieve great suffering.

(2)   Bob ought to give up smoking.

(3)   To support a loft-bed in plaster walls, anchor bolts rather than screws ought to be used.

(4)   Larry, who has had the most unfortunate time of things recently, ought to win the lottery.

(5)   The storm ought to hit shore before midnight.

In light of this diversity, semanticists typically treat 'ought' as a flexible necessity modal and give its semantics in terms of a universal quantification over possible worlds. For example, they might use a rule something like:

Modal Rule: $[[\text{ought } p]]^{fg} = \text{T}$ iff for all $w \in W_{fg}$, $w \in p$

which says that an ought-sentence is true just in case its prejacent p is true at each possible world in some contextually determined set W (usually, the set of worlds consistent with some set of background conditions f and highly ranked by some ordering function $>_g$).[5]

This rule is attractive for the way it generates decent interpretations of all of the examples (1)–(5) above without positing lexical ambiguity (which is good since the diversity of flavors of modal verbs is cross-linguistically robust, suggesting that 'ought' has univocal meaning). In each case, the difference in flavor is explained as a difference in the contextual determination of the set of worlds quantified over. As long as context can determine different values for f and g to input into the modal rule, we can provide a very systematic interpretation of diverse ought-sentences.

Some metaethicists may be tempted to suggest an old-school reductive analysis of 'ought' in terms of obligations, fittingness, or reasons. However, without prejudicing the separate issue of whether our ought-concept is constitutively connected to these

---

[5]   This is a simplified version of the rule predicted by Kratzer (1981) for necessity modals. In addition to the fact that it suppresses world-relativity of semantic values and elides the distinction between sentences of the object language and metalinguistic variables, it also prescinds from the interesting issue of how to account for 'ought's apparent weakness compared to 'must'. Kratzerian proposals for handling this are explored in Fintel and Iatridou (2008), Swanson (2011), and Silk (2012). Wertheimer (1972), Finlay (2009), and Portner and Rubinstein (2016) make somewhat different proposals, which are nonetheless consistent with the basic idea that 'ought' is a necessity modal and to be modeled in terms of a universal quantification over a set of worlds. Here I will also suppress the fascinating question of how (if at all) this rule should be modified to handle the distinction between the objective 'ought' of what's actually best and the so-called subjective 'ought' of what best optimizes ones choices given imperfect information and varying standards for practical decisions. For discussion, see Kolodny and MacFarlane (2010), Björnsson and Finlay (2010), Dowell (2013), Cariani (2013), Charlow (2013), Carr (2012), Silk (2014), Bronfman and Dowell (2016), Wedgwood (2016). Furthermore, there are other metaethically important challenges to this semantics that will remain off stage here. For instance, the standard account also has difficulty with sentences of the form "If p, then it ought to be the case that p" (see Zvolenszky, 2002; Carr, 2014). It seems to rule out strong ought-dilemmas by semantic fiat rather than philosophical argument. See Swanson (2011) and Fintel (2012) for discussion. I think each of these problems can be addressed within a broadly intensional semantics for 'ought', which is all I really need for my argumentative purposes below; but I won't discuss this further here.

other concepts in important ways, we should recognize that the modal rule's ability to provide such a systematic treatment of sentences as diverse as (1)–(5) without positing lexical ambiguity is an impressive feat—arguably one that any semantic analysis of 'ought' in terms of obligations, fittingness, reasons, or what's best will fail to do.[6]

Nevertheless, I still believe there is a lesson proponents of the modal rule might learn from the urge amongst metanormative theorists to connect 'ought' to notions such as obligations, fittingness, or reasons. Ought-sentences seem, when used in the sense that is of most interest in metanormative theory, to call on *agents* to *do* something; and arguably this is a more fine-grained notion than a proposition's being true at possible worlds. Here is Geach in a classical statement of the objection:

If the deontic operator applied to whole propositions, then the results of modifying the active and the equivalent passive form with an 'ought' must likewise be equivalent. But such equivalences often appear not to hold. 'John beats up Tom' and 'Tom is beaten up by John' are equivalent; but it looks as though 'John ought to beat up Tom' and 'Tom ought to be beaten up by John' are not necessarily equivalent.   (1982, p. 3)

In light of this good point, rather than abandoning the modal rule altogether,[7] I think we should enhance it by integrating it with a semantics for imperatives. After all, saying that S ought to ϕ—in the sense of most interest in metanormative theory—seems closely related to issuing the imperative, "$S \phi$!"

How might we enhance the modal rule in a way that respects this connection? Here's a start: Formally, semantic models typically assign declarative contents (propositions) truth values relative to points of evaluation, which are conceived as (at least) possible worlds (which in turn might be thought of in a variety of ways, e.g. as concrete possible universes, abstract ways reality could be, sets of propositions, or something else). As far as formal semantics goes, we might also want to assign *imperative* contents (what I'll call "prescriptions") a kind of semantic value relative to a point of evaluation—after all, imperatives are also logically-composed sentences of language standing in semantic relations such as entailment and equivalence to other sentences—but the kind of semantic value and the conception of the point of evaluation has to be different (since imperatives do not seem to have truth values).[8] One proposal that I like assigns imperatives the semantic values legitimate/illegitimate relative to possible norms (which in turn might be thought of in a variety of ways, e.g. as concrete laws of

---

[6] I consider various attempts to reductively analyze 'ought' in these terms in chapter 2 of Chrisman (2015); the main stumbling block is plausibly capturing the meaning of epistemic ought-sentences such as (5) while retaining a plausible treatment of paradigmatic normative ought-sentences such as (1) and (2).

[7] Schroeder (2011) argues that it motivates treating 'ought' as ambiguous between raising-verb and control-verb readings. I discuss this proposal and show why it isn't linguistically motivated in Chrisman (2012a). See also Finlay and Snedegar (2013) for pragmatic explanations of many of the phenomena that lead Schroeder to posit ambiguity.

[8] A point which is argued for in greater detail in Charlow (2014b, section 4), Starr (2014, section 2), and Chrisman and Hubbs (unpublished, section 2).

possible practices, abstract ways agents might behave, sets of general prescriptions, or something else).[9]

It is not important here whether this is the right semantics for imperatives, but if we posit the imagined prescriptional contents with legitimacy values rather than truth values, then we can embed this in our semantics for 'ought' by adopting something like the following as a new general rule articulating the truth-conditional contribution of 'ought':

Enhanced Modal Rule: $[[\text{ought } p]]^{fg} = \text{T}$ iff for all $<w, n> \in P_{fg}$, $<w, n> \in p$

Here we no longer restrict p to propositions that might be expressed by declarative sentences and modeled as a set of possible worlds; instead we allow that it could also be the sorts of prescriptions that might be expressed by imperative sentences, and we model it in general as a set of world-norm pairs (the worlds where the prescription is legitimate relative to the norm). This rule says that an ought-sentence is true[10] just in case its prejacent is correct relative to each member of some contextually determined set of world-norm pairs. This means, if we allow that the prejacent p could be either propositional or prescriptional, we get two species of ought-claims out of a univocal semantic rule. When the prejacent is propositional, the rule works similarly to the one above: the ought-sentence is true just in case the prejacent is true relative to each of some contextually determined set of world-norm pairs, where the norm parameter is usually idle. By contrast, when the prejacent is practical, the rule evaluates its legitimacy relative to each of some contextually determined set of world-norm pairs, where the norm parameter is usually crucial.[11]

---

[9] Compare Castañeda (1975, chapters 3–5) for the most detailed way of working out this basic idea about imperatives (but note that he uses the term "prescription" in a different way to how I am here). Charlow (2014b) defends a related idea in an expressivistic framework which treats imperatives and declaratives as having properties of representations of states of mind as their semantic values, where ordinary declaratives denote properties of representations of beliefs and imperatives denote properties of representations of preferences. Hauser (1980), Segerberg (1990), Portner (2004), and Hanks (2015, chapter 9) defend competing views about the semantic value of imperatives that can still be used to generate sentence-level contents that are not propositions, which in turn can be thought to be embedded under 'ought's. In Chrisman and Hubbs (unpublished), we show how treating such nonpropositional contents as the semantic values of imperatives can be motivated by action-theoretic considerations ancillary to the project of compositional semantics. As indicated above, what is crucial here is not the correctness of this approach to the semantics of imperatives but rather that it can handle the content embedded under prescriptive-seeming 'ought's. A perhaps more orthodox view is that imperatives have action-related *properties* as their contents and their semantic relations to other sentences are derivative of the semantic relations between the propositions that result from ascribing these properties to individuals. I think this view has trouble explaining the semantic difference between necessarily coextensive but nonsynonymous imperatives, and I doubt that it gets the logic of imperatives right. But that doesn't matter here since I am not relying on the view of prescriptions above as the right view of the content of imperatives.

[10] Ought-sentences are clearly declaratives and so they get *truth* values and have *truth* conditions on this view. Although this treatment of their truth conditions allows that they can embed prejacents which are not truth-apt, this is not a version of the failed noncognitivist idea that ought-sentences are neither true nor false.

[11] This bears some similarity to the proposal developed in Gibbard (1990, chapter 5) and Gibbard (2003, chapters. 3–6) to extend the standard possible-world semantics, adding an extra norm-index relative to

With a rule such as this one, I think we overcome the worry from Geach about applying truth-conditional semantics to 'ought.' The prescriptions *John, beat up Tom!* and *Tom, be beaten up by John!* are different from each other (and from the proposition *that John beats up Tom*); they could easily differ in whether they are legitimate relative to a norm. The enhanced modal rule lets us understand the truth conditions of ought-sentences in terms of a universal quantification over a contextually determined set of possibilia.[12] Even within a univocal treatment, we can recognize an important semantic difference between obviously *propositional* ought-sentences, such as (4) and (5), and obviously *prescriptional* ought-sentences, such as (1) and (2). It may be less clear where other ought-sentences fall in respect of this divide, but that is plausibly because determining this requires more information about the context in which they are used.

In any case, the enhanced modal rule is the rule articulating 'ought's contribution to truth conditions that I will use in the rest of this paper. As I indicated above, I will not provide any more argument for this semantics of 'ought' than is already in the barebones explanation of the rule, and I fully grant that further refinements of it are needed for various purposes in semantics. My purpose in explaining it is mainly to frame the question I want to pursue in the rest of this paper: how should we interpret this rule? Recall at the outset I suggested that a central debate in metanormative theory seems to be pulled in two directions based on whether one takes more seriously the idea that 'ought's (at least the ones relevant to metanormative theory) are distinctively prescriptive, or the idea that all declarative sentences should be given a truth-conditionalist semantics. With something like the enhanced modal rule in hand, I think we can see this to be a false choice by exploring two different metasemantic accounts of what this rule tells us about the facts in virtue of which ought-sentences have the contents that they do. Both of these accounts are consistent with the idea that some ought-claims function not as descriptions of what is actually the case but rather prescriptions for what someone is to do in various circumstances. The interesting question is how they differ in capturing this idea.

3   Truth Conditions as Ways Reality Could Be

The first account stems from what I take to be a pervasive view in contemporary philosophy. This is that declarative sentences have the contents that they do in virtue

which some sentences could be evaluated. However, Gibbard uses this extension to provide semantic values for *declarative* sentences with normative elements, in essence committing to either a kind of noncognitivism or a kind of norm-relativism about the possibility of true normative declaratives. By contrast, I'm applying this idea only to nonpropositional contents, like those we might assign to imperatives, which I deny are truth-apt and for which I think it is much more plausible to assign some kind of semantic value other than truth.

[12]   An alternative attempt to do this posits elided "stits" ("sees-to-it-that") in the prejacent proposition of 'ought's intuitively thought to be practical. As a semantic proposal, I think this is ad hoc and non-explanatory. But arguing for that is not necessary to the points I want to make here.

of how they represent reality as being. To see how this works, consider first a simple sentence, such as

(6)   Grass is green.

It is very tempting to think that sentences such as this one are true just in case reality is the way they represent it as being. That is to say, e.g., that (6) represents grass as having the property of being green. Generalizing from this example, it is tempting to think that pursuit of a truth-conditional semantics boils down to the development of a recursive method for spelling out the way reality would have to be in order for each declarative sentence of the language to be true. Using the terminology I prefer, we can call this a *representationalist* explanation of truth conditions. Arguably, it is a core assumption of all traditional realist and error-theoretic views in metaethics. They would say that

(7)   Stealing is wrong.

is true just in case reality is the way (7) represents it as being. Then there is room for debate among these philosophers about the nature of this putative piece of reality (naturalism vs nonnaturalism) and whether it actually obtains (success-theory vs. error-theory).

The pervasiveness of this representationalist interpretation of truth-conditional semantics can help to make some sense of why anyone would have ever denied, as early noncognitivists did, that normative sentences have truth conditions. For, on the face of things, it seems obvious that some normative sentences are true and others are false; and from the point of view of ordinary language, embedding ethical sentences such as (7) under the truth predicate and other propositional contexts such as belief/knowledge attributions and epistemic modals seems beyond reproach. However, if your view is that normative language at its base is prescriptive rather than descriptive, and you accept (at least tacitly) the representationalist conception of what truth-conditional semantics is telling us, then you might be tempted to argue that normative sentences do not have truth conditions (at least not in any theoretically interesting sense).

Below I'll suggest that there's a way that even those metanormative theorists inspired by noncognitivism can resist this temptation—as they should, since it is difficult to see any hope of realizing the aspirations of compositional semantics if we pursue a bifurcated approach to semantics, using truth conditions as our basic semantic framework for nonnormative sentences such as (6) but something else for normative sentences such as (7).[13] But before we get there, I want to explain why I think the metanormative situation with respect to 'ought' complicates things in a way that allows

---

[13] This is the main lesson of Hale (1993), Dreier (1996), Unwin (1999). Schroeder (2008) shows how difficult it is for an expressivist to avoid this even by advancing a psychologistic semantics across the board, i.e. applied also to paradigmatically descriptive sentences.

even the hardcore representationalist to accept, indeed to champion the idea that normative language (at least insofar as it is "fraught with ought") is prescriptive rather than descriptive, undercutting some of the traditional noncognitivist motivation for denying that normative sentences have truth conditions.

Here's why things are more complicated. Even if you accept the representationalist conception of truth conditions in the basic case, you will want to allow there are some declarative sentences that do something other than describe what is actually the case. This is because some sentences are formed with intensional operators, e.g.

(8)   Jack and Mary might be in the house.

And intensional operators do not seem to function like normal predicates. Instead, they are generally thought by semanticists to serve the purpose of shifting some parameter of the circumstance relative to which we evaluate the content that they embed. That is, for instance, they tell us to evaluate some embedded proposition not at the actual world but at some possible world(s). For example, the idea is to think of (8) as true just in case the embedded proposition *that Jack and Mary are in the house* is true in some world regarded as possible given a particular body of evidence. This is why its truth conditions are usually given with something like the following:

$[[(8)]]^c$ = T iff $[[Jack\ and\ Mary\ are\ in\ the\ house]]^c$ = T in some possible world consistent with our evidence in this world

As a working hypothesis, this is all relatively familiar from textbook treatments of intensional operators in semantics.[14] It may need to be refined, but the important question at this stage is not semantic but metasemantic: What is the representationalist gloss of these truth conditions? One fairly natural thing for the representationalist to say is that these truth conditions tell us how this sentence represents reality as being, but the piece of reality it represents is more complicated than anything having to do merely with what is actually the case. The thought is that (8) represents how things are not just in the actual world but also in other possible worlds; this involves representing some kind of relation between possible worlds. So, if we assume an ontology of possible worlds,[15] we can continue to think of all declarative sentences, even those formed by intensional operators, as representing ways reality could be. But (here's the payoff) we can simultaneously deny that all declarative sentences describe how things actually are. For on this view some sentences represent a piece of reality that involves

---

[14]   Not that it is uncontroversial. Even Kratzer (1991) recognizes that its application to epistemic modals faces challenges making sense of apparent disagreement, which has led some to follow MacFarlane (2011) in complicating the definition of truth for such claims, relativizing to a context of assessment; but see von Fintel and Gillies (2008, 2011) for criticism and a more conservative response.

[15]   What about linguistic *Ersatz*-theorists who reject the reality of possible worlds but contend that some other part of our ontology can play the role of possible worlds in our interpretation of the truth conditions of modal sentences? Whether they count as representationalists in the sense I'm trying to capture here is going to depend on how exactly they conceive of the other part of our ontology as playing the role of possible worlds.

a relation between the actual world and other merely possible worlds. Hence, on the representationalist interpretation of the standard truth-conditionalist semantics for them, might-sentences are no mere description of how things actually are.

4 Application to 'Ought'

I hope it's already clear that this is relevant to the apparent tension between the idea that 'ought's do not describe what is actually the case but prescribe possible action, thought, and feeling, and the idea that all declarative sentences deserve the same general semantic treatment, e.g. in terms of compositionally specified truth conditions. Even if we just use the simple modal rule for 'ought,' it can be viewed like 'might' as a device for shifting some parameter of the circumstance relative to which we evaluate the propositions it embeds rather than for describing some feature of what is actually the case. Then, representationalists will interpret the truth conditions this rule predicts for ought-sentences as specifying how the sentences represent a complex modal feature of reality as being—not as merely describing how things actually are.[16] But more to the point, if we use the enhanced modal rule, I think the representationalist can even capture an important sense in which at least some 'ought's are prescriptive rather than descriptive. For example, we can use this rule to interpret

(9)   Jack and Mary ought to leave.

predicting the following truth conditions:

$[[(9)]]^{fg}$ = T iff [[Jack and Mary, leave!]] = L in for all $< wn > \in P_{fg}$

This means that (9) is true just in case its prejacent prescription is legitimate relative to all of the relevant points of evaluation (in this case world-norm pairs highly ranked by $>_g$ that are consistent with background conditions f). Then, assuming an ontology of possible worlds and possible norms, representationalists can interpret these truth conditions as representing a modally complex way reality could be: a way the actual world is related to possible worlds and possible norms. More specifically, representationalists will think the formula above tells us that the prescription for Jack and Mary to leave is legitimate across a contextually determined range of possible norms in light of the circumstances common in a contextually determined range of possible worlds. In this way, (9) is assigned truth conditions, but since these are truth conditions involving what is prescribed by norms in various possible worlds, this

---

[16] So it's possible to reject my suggestion that 'ought's can embed prescriptive prejacents and still buy the rest of what I say here as long as there's an alternative way to capture the prescriptivity of some 'ought's. Some may be inclined to do this by treating one of the ordering sources evoked by 'ought' (according to the modal rule) as ordering worlds in terms of whether some set of prescriptions is satisfied. I think this is going to be too coarse-grained to respect Geach's point from above, but a lot will depend on how we work out the details of prescriptive ordering sources.

sentence can be recognized as prescriptive rather than descriptive of what is actually the case.

So, if that representationalist story is on the right track, then it's wrong to think we're forced to choose between a truth-conditionalist approach to semantics and the idea that ought-sentences are often prescriptive rather than descriptive of what is actually the case. By conceiving of ought-sentences as representing a complex modal piece of reality, involving possible norms and what they prescribe in various possible worlds, this apparent dilemma is avoided.

It is avoided, we should note, by divorcing the notions of *describing what's actually the case* and *representing reality*, allowing that the truth conditions of some sentences (at least those involving intensional operators) tell us that these sentences have the contents that they do in virtue of representing more than what is actually the case. They represent something like merely possible worlds and merely possible norms and their relations to the actual world. Hence, the representationalist way of avoiding the dilemma comes with ontological (and associated epistemological) costs.

This is a familiar worry in metanormative theory, but I think the representationalist interpretation of a workable truth-conditional semantics for 'might' or 'ought' reveals how it extends into the metaphysics of modality. By treating intensional operators as generating representations of reality, the articulation of their truth conditions gestured at above would force us to understand anyone committed to the truth of sentences such as (8) or (9) as committed to the reality of possible worlds and possible norms (or their *ersatz*-replacements). Are the merely possible worlds, e.g. where there are dragons in the command of a princess, and the merely possible norms, e.g. which require dragons to blow fire on the princess's enemies, just as real as the greenness of grass or the location of people in the house? Once in the representationalist mindset, it may be difficult to see how the answer could be anything other than "yes"; but taking a step back, many metanormative theorists are going to worry that this is an ontologically profligate position affronting to common sense about the difference between what is real and what is imaginary/virtual/fictional/merely-possible (not to mention the more theory-laden distinction between what *is* the case in reality and what reality *might/ought/*etc. to be like).[17]

Perhaps this ontological cost (and any related epistemological costs) should be paid. I'd certainly be willing to pay provided it was required to get the benefits of systematicity and seamlessness provided by truth-conditionalist approaches to compositional semantics. Below, however, I want to suggest that there are alternative metasemantic accounts of the results of compositional semantics—ones not committed to the idea that ought-sentences represent a (complex modal) way reality could be. I think this

---

[17]   Are these things perhaps real but not "just as real" as the greenness of grass or the location of people in the house, because they are *abstract*? I doubt that the distinction between concrete and abstract elements of reality helps here, as many of the things we probably want to say are "more real" than possible worlds and possible norms would, on many ways of drawing this divide, count as abstract.

coheres with the popular stance among many philosophers that possible-worlds talk in semantics is a convenient fiction or a mere notational device, or that possible worlds are "pleonastic entities" rather than real parts of our ontology.

5 Truth Conditions as Positions in a Space of Implications

Next, I want to introduce a different metasemantic view inspired by Sellars's (1953, 1974) inferential-role account of semantic and mental content and developed in more systematic detail in Brandom's (1994, 2000) inferentialist version of pragmatism about intentional content. I believe this is somewhat less familiar, so I'll take a bit more space to introduce it and bring it to bear on the metanormative issues introduced above. The core idea is to reverse the normal order of explanation that moves from a view about the semantic content of bits of language to a view about proprieties of use in a social discursive practice. Instead, these inferentialists seek to explain why sentences have the contents that they do in terms of what they are vehicles for doing in a social discursive practice. However, not everything we can do with language is relevant.[18] When it comes to declarative sentences, these are conceived primarily as vehicles for making assertions, which conception is then understood as undertaking (acknowledging, or embracing) a particular kind of commitment.

More specifically, looking downstream from a sentence, various further sentences would typically follow from its truth—e.g. part of the commitment carried by an assertion of "Grass is green" is to grass's being colored.[19] Then, looking upstream from a declarative sentence, it would typically follow from the truth of various other sentences; and one who uses a declarative sentence to make an assertion is open to a potential challenge to entitle oneself to this commitment by asserting one of these other sentences—e.g. one who asserts "Stealing is wrong" might, if legitimately

---

[18] This is, I believe, what distinguishes the view from Wittgenstein-inspired "use theories of meaning" based on skepticism about drawing any principled line between the uses of language that reveal semantic competence and those that reveal merely pragmatic understanding of the way language can be used to achieve sundry goals. Brandom has recently characterized his linguistic pragmatism as the combination of two principles. *Methodological pragmatism*: "the point of associating meanings, extensions, contents, or other semantic interpretants with linguistic expressions is to codify proprieties of use," and *semantic pragmatism*: "all there is to effect the association of meanings, contents, extensions, rules or other semantic interpretants with linguistic expressions is the way those expressions are used by the linguistic practitioners themselves"(2011, pp. 21–6). As long as we restrict the focus to the uses of language to make assertions, which are distinguished in the way they license inference to other assertions and require defense via other assertions, I think we get close to what most workaday semanticists take as their principal data: ordinary speaker intuitions about the entailment, inconsistency, equivalence properties of sentences.

[19] One may not always acknowledge the inferential consequences of what one asserts. So its being part of what one is committed to is not the same as *believing* that it follows. But the more someone fails to acknowledge what we take to follow from what one asserts, the less we will think she understands what she is saying.

challenged, have to back it up with reasons, for instance arguing that stealing involves a gross violation of the respect we owe to each other as human beings.[20]

Importantly, these upstream and downstream inferential connections are not merely *logical* entailments but rather in the broader family sometimes called semantic implications, which are usually defeasible and context-sensitive.[21] For this reason, competence with the meaning of sentences can be thought of, on the inferentialist picture, as coming in degrees. How much one understands what one is saying depends on how much one knows one's way around a space of semantic implication-relations containing the sentence and the types of things that would defeat these relations.[22]

Before we can bring this idea to bear on the metanormative issues introduced above, there are two questions I want to address in order to flesh out the inferentialist idea.

First question: What does endorsing inferentialism mean for truth-conditional semantics? Brandom sometimes writes as if inferentialism is inconsistent with appealing to truth conditions in one's explanation of the semantic composition of declarative sentences,[23] but I doubt inferentialism and truth-conditional semantics are really in tension, at least not when we understand "truth-conditional semantics" as the approach in compositional semantics where one attempts to recursively derive some specification of when a declarative sentence should be assigned the preferred semantic value *true* relative to some parameters, thereby displaying the semantic function of the sentence's subsentential parts. Because of the pervasiveness of representationalism, the representationalist gloss on the results of this approach is sometimes labelled "truth-conditional semantics." But I think it's helpful to view representationalism

---

[20] Often one may not actually be challenged to defend what one asserts, and, regarding certain kinds of topics, one may be default entitled in most or all of the everyday situations where one would make an assertion; but being the sort of linguistic move whose form makes it challengeable for reasons is essential for being an assertion. (Brandom 2000, pp. 193–4)

[21] We can then seek to recover specifically "logical" entailments by restricting ourselves to consideration of the semantic implications ensured by the "logical" form and words of the sentence.

[22] Sellars and Brandom tend to take a very expansive view of the implication-relations that are relevant to the determination of meaning. However, any plausible version of their views would embrace the idea that *how* central such a relation is to determining the meaning of a particular word/concept is something that comes in degrees, and the lesson we were supposed to have learned from Quine's attack of the analytic/synthetic divide is that it is not possible, in general, to draw a sharp line between the implications that hold in virtue of meaning alone and the implications that hold in virtue of something else (though if we hold enough other stuff fixed, we might be able to draw a practically useful line in particular contexts).

[23] For example, he writes, "…*truth* is not a concept that has an important *explanatory* role to play in philosophy. Appearances to the contrary, are the result of *mis*understanding its distinctive *expressive* role. The word 'true' does indeed let us *say* things that in many cases we could not say without it. But when we understand *what* it lets us say, and *how* it does that, we will see that the very features that make it *expressively* useful make it completely unsuitable to do the sort of theoretical *explanatory* work for which philosophers have typically enlisted its aid." (2009, pp. 158–9) Later, however, he strikes a slightly more conciliatory tone, "What about the role of truth in semantic explanation, via a definition of propositional content in terms of truth conditions? We certainly do use 'true' to *say* what the content of a claim is. … But it would be a mistake to infer from this sort of appeal to truth conditions to express propositional contents that one can explain what propositional contents are by appeal to the conditions under which sentences are true." (ibid., p. 162)

and inferentialism as a metasemantic interpretation of the results of first-order compositional semantics (whatever these happen to be). More specifically, I think an inferentialist should view the specification of the truth conditions of a sentence as an initial spelling out of what one who asserts the sentence is committed to as a premise with further downstream inferential consequences, and what one would have to entitle oneself to by looking to upstream inferential antecedents were one to be legitimately challenged. In this way, we might say that an assignment of truth conditions to a sentence is an articulation of the position in a network of implications occupied by one who uses the sentence to make an assertion in ordinary discursive practice.[24]

In cases such as (6) and (7), this may appear trivial. Of course, one who asserts that grass is green has undertaken a commitment describable as a commitment to the proposition that grass is green; and one who asserts that stealing is wrong has undertaken a commitment describable as a commitment to the proposition that stealing is wrong. But when it comes to more complex cases where it is not entirely clear what truth conditions to assign to a sentence, this way of thinking of things lets us interpret candidate assignments as articulations of different positions in a network of implications. For example, if someone asserts

(10)   Visiting relatives can be boring,

is she committed to the proposition that those relatives who visit are boring, or is she committed to the proposition that going to visit relatives is boring? Disambiguating this sentence by assigning truth conditions in line with one or the other of these interpretations has implications for what further commitments downstream from an assertion of this sentence would carry and what further assertions one would have to make upstream from this assertion in order to entitle oneself to it if legitimately challenged. In this way, assigning truth conditions can be viewed as an attempt to identify what inferential commitments are implicitly endorsed, embraced, or acknowledged by someone who uses the sentence to make an assertion. Rather than being—in the first instance—commitments to reality's being a particular way, they are commitments that carry downstream inferential consequences and that must, in principle, be defensible via upstream inferential antecedents.

Of course, representationalists will typically agree that declarative sentences have upstream and downstream inferential connections to other sentences that are somehow relevant to understanding semantic content. What distinguishes inferentialism is that it appeals directly to the spectrum of these connections (rather than to anything about representing reality) in its metasemantic account of what it is in virtue of which sentences have the truth conditions that they do.

Second question: If declarative sentences are understood first and foremost as vehicles for undertaking inferentially articulable commitments, what does this mean for the

---

[24] Compare Williams (1999, pp. 552–7; 2010), who also argues that inferentialism is compatible with a truth-conditionalist approach to compositional semantics.

intuitive thought that many declarative sentences describe the way things are in the world around us? Many philosophers seem to think that inferentialism undermines the very idea of a sentence's *describing reality*. The thought is something like this: if semantic contents are positions in something we characterize as merely a "game" of giving and asking for reasons *rather than* specifications of how reality has to be in order for the relevant sentence to be true, then surely *no* sentences should be viewed as describing reality. Alternatively, inferentialism is sometimes portrayed as taking what is usually a local denial that some area of discourse is representational (e.g. normative or modal) and generalizing it to all areas of discourse. Here, the thought is something like this: If everything we say is somehow about how we should infer, and you think this "should" is prescriptive rather than descriptive, then surely that means that nothing we say is a description of how things are in the world around us.[25]

While granting that some things some inferentialists have written may suggest that this is the view,[26] I think a wholesale rejection of the idea that some declarative sentences describe the world around us is an extreme version of the view not actually endorsed by some of its central proponents. For example, Sellars was famously an inferentialist about content and a realist about science, writing "in the dimension of describing and explaining the world, science is the measure of all things, of what is that it is, and of what is not that it is not" (1956, §41). And Brandom characterizes his inferentialism as supported by a "sophisticated expressivism" about logical, modal, and normative vocabulary that "…is essentially, and not just accidentally, a *local* expressivism." This is because, in the way he understands the expressive role of these vocabularies, it depends on the existence of other vocabularies that play a descriptive role: "Autonomous discursive practices *must* contain vocabularies playing *other* expressive roles—for instance, observational vocabulary that reports features of the non-linguistic bits of the world" (2013, p. 102). In the context of the arguments here, this is good because any view that implies that ought-sentences aren't descriptive because *no* declarative sentence is descriptive could achieve only a pyrrhic victory.

[25] For example, Price imagines what a quasirealist expressivist like Blackburn might say if he adopted a Brandomian inferentialist conception of assertion. Price suggests the expressivist might say, "I used to think of my Humean expressivism as a local position, applicable to some vocabularies but not others. . . . However, I have now come to realise that for no vocabulary at all is it theoretically interesting to say that its function is to 'represent' particular kinds of states of affairs.…This…does nothing to undermine the interesting observations that got me started, about the distinctive—and different—functions of moral and modal vocabulary. On the contrary, it simply implies that they are exemplars of an approach to language we should be employing everywhere. In other words, what I took to be linguistic islands are simply the most visible extremities of an entire new continent—a universal program for theorising about language in expressivist rather than representationalist terms . . ." (2011, pp. 101–2).

[26] Macarthur and Price (2007) appear very sympathetic to this idea. But, in later work, Price (2013, chapter 3) has developed a more nuanced position, which trades on a distinction between thoughts that carry information (what he terms *i-representations*) and a species of these which empirically track the mind-independent world (what he terms *e-representations*). See also Blackburn (2013) and Williams (2013) for useful discussion and repudiation of global expressivism in the sense at issue here.

So, in my view, inferentialism is not committed to global expressivism or the idea that language use is *merely* a "game" of giving and asking for reasons, rather than a central part of how humans conceptualize the world and reason about what to do in it. To the contrary, I think someone sympathetic with inferentialism might argue that some or even many sentences have direct inferential consequences for our views about the way reality is. Knowing how to draw the line between those that are and those that aren't is difficult, but maybe "Grass is green." is such a sentence, in which case, one who asserts it might be thought to be inferentially committed to the proposition that grass has the (real, mind-independent, describable) property of being green. That is to suggest that one of the direct downstream inferential connections of this sentence is to another sentence deploying explicitly representational notions. Maybe we think this is the best way to make sense of the idea that the use of such sentences is causally regulated by tightly corresponding features of the world, such as the color of grass, and that such use is plausibly treated on a tracking model of the relation between language use and extralinguistic reality.[27] If this were right about this sentence or any other class of sentences, inferentialism might be seen as overlapping with the core representationalist idea outlined above but as also offering a deeper explanation of the relevant phenomena.

Of course, if we maintained that *all* declarative sentences had such direct inferential consequences for our views about the way reality is, then we would have a metasemantic view that was supposed to be an alternative to representationalism but which turned out to be a sophisticated version of it.[28] As we'll see below, the textbook expressivist move in this context is to insist on a distinction between mental states with descriptive and directive roles in the psychology of motivation ("belief-like" and "desire-like" attitudes as they are often called, with different "directions of fit") and argue that normative claims express mental states with a desire-like direction of fit with the world. Whether or not this is consistent with truth-conditional approaches to compositional semantics, especially as applied to 'ought,' is a complicated and

---

[27] Williams (2010, pp. 323–8) explores a more sophisticated way to cash out this distinction in terms of various clauses in an explanation of meaning in terms of use. As Williams argues, some such explanations may be ontologically conservative in that they do not mention referents (properties when it comes to predicates) of the terms whose use they are appealing to in explaining meaning, but others won't be. My suggestion here (following Williams) is that the explanations of the meaning of color vocabulary in terms of its use might not be ontologically conservative, while the explanation of deontic modal vocabulary is a plausible candidate for an explanation that can be ontologically conservative.

[28] Drawing on Peacocke (1986; 1987) and Block (1987), Wedgwood (2001, 2007, chapter 4) articulates a theory of conceptual roles as determining the representational function of the words that express concepts with these roles. On this view, each word refers to something, and what it refers to is determined in part by the conceptual role of the concept it expresses; and it is the logical composition of these concepts that determines how each declarative sentence represents reality as being. Because he uses a truth-conditional semantic rule similar to the modal rule stated above, he ends up defending a conceptual-role interpretation of this rule that bears important similarities to the one I develop below. However, because he treats all words as referring to something in reality, his inferentialism makes no room for distinguishing between representational and nonrepresentational vocabulary. By contrast, I see in the inferentialist metasemantics resources for motivating a distinction at least when it comes to 'ought.'

controversial matter that I'll touch on below. But as I understand inferentialism, at least of the sort developed by Sellars and Brandom, it offers a different set of resources for marking out ought-claims from those we typically think of as representing reality. In short, this is to classify 'ought' with other modal terms and then to treat modal terms along the lines of sentential logical operators. Typically, inferentialists explain why logical operators have the content they do, not in terms of what in reality they represent but instead in terms of something they allow us to do: state and endorse/reject claims about the inferential commitments carried by the use of other sorts of terms. How exactly does this work for 'ought'?

6 Application to Necessity Modals Including 'Ought'

There is another idea championed by Sellars and Brandom that I think is not consti-tutive of an inferentialist metasemantics but a natural extension of it. The idea is that some terms do not earn their keep from the way they let us represent pieces of reality but rather by serving as what we might call *metaconceptual devices* for embracing commitment to features of the conceptual (inferential) framework in which ordinary representational terms have determinate application. Sellars (1956, 1958) defended an explicitly metalinguistic and inferentialist version of this claim, arguing that in addition to descriptive words applying to empirically knowable pieces of reality, our language contains "categorical" words whose core function is not to describe reality but to embrace commitment to the inferential relations which must hold between descriptive words for them to count as descriptive in the first place. In this way, Sellars viewed philosophically puzzling words as many philosophers view logical words: not as adding to the descriptive content of the sentences in which they figure but rather as devices for affirming commitment (at least implicitly) to the inferential connections between other words (often words which *do* add to the descriptive content of the sentences in which they figure). On Sellars's way of working out this line of thought, implicit commitment to these kinds of inferential connections is seen as a crucial element of ordinary competence with other descriptive words; and the metaconceptual words provide, as Brandom puts it, means for making explicit those implicit commitments.

There may be other philosophically interesting words that are metaconceptual rather than representational in this sense.[29] However, I think one of the most per-suasive uses of this idea is for conceptual necessity claims. For example, consider:

(11)   Necessarily, if x is east of y, then y is west of x.

---

[29] Kant (1787, A80/B106), for example, identified twelve "pure concepts of the understanding" interpret-ing them in something like this way. In a similar vein, Sellars (1956) argued that terms such as "substance" and "quality" are metalinguistic devices that get their content from the way they can be used to endorse the types of inferential connections between other empirical terms, connections which must be in place, at least implicitly, for these terms to count as applying to empirically knowable pieces of reality in the first place.

Sellars and Brandom would say that this sentence does not represent a piece of reality involving the modal relation between locations which are situated east/west of each other. Rather it is a device for acknowledging the conceptual validity of inferences of the form

x is east of $y \rightarrow y$ is west of x

thereby (partially) spelling out and explicitly committing to the inferential connections between more basic terms such as 'east' and 'west'.[30]

Does this mean that, according to Sellars and Brandom, conceptual necessity claims such as (11) are *mere* "inference tickets" in a "game" of giving and asking for reasons, rather than sentences with truth-conditional content? I think it would be foolish to make that claim.[31] Instances of (11) can clearly be embedded in "It is true that…" and other truth-functional contexts, and they express something that can be doubted, believed, known, etc. So, it would be very strange to say that they do not have truth-conditional content. Indeed, something akin to the standard possible-worlds semantics for alethic modality provides a pretty good articulation of these truth conditions.

Assuming this is (roughly) right, the crucial metasemantic question is what it is in virtue of which this truth-conditional specification of the content of some instance of (11) is correct (insofar as it is)? As we have already seen, the representationalist answer is that it is correct in virtue of correctly identifying how that instance of (11) represents reality as being. By contrast, an inferentialist following Sellars and Brandom would say these truth conditions are correct in virtue of correctly modeling the position in space of implications taken up by someone who asserts that instance of (11) in normal discursive practice. One who accepts this sentence is committed to y's being to the west of x, if she is committed to x's being to the east of y. In most contexts, were it to come up, in order to entitle oneself to an instance of (11), one might need only to affirm one's competence with the words 'east' and 'west'.[32]

---

[30] Compare Thomasson (2007, 2013) who argues that this is the appropriate account of metaphysical necessity claims. Here my suggestion on behalf of Sellars and Brandom is the weaker idea that at least one species of necessity claim plays this inference-rule endorsing role. If Thomasson is right, however, that's a congenial conclusion.

[31] Kant and Frege sometimes seem to be making something like this claim, e.g. "…the modality of judgments is a very special function thereof, which has the distinguishing feature that it does not contribute to the content of the judgment" (Kant, 1787, A74/B100), and "By saying that a proposition is necessary I give a hint about the grounds for my judgment. But, since this does not affect the conceptual content of the judgment, the form of the apodictic judgment has no significance for us" (Frege, 1879, p. 5). However it is probably anachronistic to view them as speaking of the compositional semantic content of modal sentences rather than something like the empirical or descriptive content of the thoughts expressed by canonical uses of these sentences.

[32] What about conceptual necessity claims whose prejacents are not explicitly conditional in form? As with all modal claims, we'd still articulate their truth conditions relative to background conditions in the context of use, and these truth conditions could be viewed as depicting the idea that one who takes those background conditions to hold is committed to the truth of the explicit prejacent.

This is how I see an inferentialist metasemantics, when combined with the further Sellarsian/Brandomian idea that some terms are metaconceptual devices, as providing resources for interpreting the truth conditions of some sentences along non-representationalist lines. So, now we should turn back to the main topic of this paper: 'ought.'

Above I claimed that the standard treatment of this word in compositional semantics is as a necessity modal rather than an ordinary predicate-forming term. That already suggests interpreting the truth-conditional contribution of 'ought' in terms of its being a metaconceptual device. The main objection to this basic idea, I think, is that the standard possible-worlds semantics for necessity modals faces difficulty in making sense of the distinctive prescriptive uses of 'ought' common in normative discourse. That is why I favor the enhancement of the standard account sketched in §2. This continues to treat 'ought's truth-conditional contribution in terms of a universal quantification over possibilia, but it refines the kinds of prejacent content 'ought' is thought to embed, and it makes correlative enhancements to the conception of the possibilia at which these prejacents are to be evaluated.

I think this can now be leveraged into a very plausible inferentialist treatment of 'ought' as a metaconceptual device, which then opens up space for distinguishing all ought-claims from descriptive claims. The core idea is to interpret 'ought' as getting its content from its role in embracing commitment to inferential connections between more basic items, rather than from its representational purport. This is analogous to the metasemantic story about conceptual necessities provided by Sellars and Brandom. Then we deploy the enhanced modal rule from before to generate specific explanations of the types of more basic claims whose inferential connections are embraced via an ought-claim.

For instance,

(12)   We ought to do more to relieve great suffering.

would be said to have truth conditions something like the following:

$[[(12)]]^{fg}$ = T iff [[Let's do more to relieve to relieve great suffering!]] = L in all $<w, n> \in P_{fg}$

where this is conceived as a modal operation on a prescription, evaluating its legitimacy across a contextually determined set P_{fg} of world-norm pairs.

However, inferentialists won't gloss this as an articulation of how (12) *represents* modal reality as being (e.g. as containing various possible worlds and possible norms, according to which the relevant prescription is legitimate). Rather they could interpret the truth-conditional specification of the content of (12) as a depiction of a point in a network of inferentially specifiable commitments one takes on when one endorses (12). Downstream, one is committed to prescriptions to act (in particular contexts) in ways that comply with the prejacent prescription; and upstream one might defend (12) by appealing to more general normative claims that support the norms

relative to which the prescriptional content is legitimate at the contextually relevant worlds.

Although inspired by Brandom, it is worth noting that my proposal here is considerably more conservative than his suggestion (in Brandom, 2000, chapter 2) that *all* normative and evaluative terms should be understood as metaconceptual devices governing specifically *practical* inferences. He considers such sentences as:

(13)   You're obliged, if you're a bank employee, to wear a necktie to work.

(14)   It is wrong, if it's done to no purpose, to harm someone.

and claims a sentence like these "... makes explicit the endorsement ... of material proprieties of practical reasoning" (2000, p. 89). His idea is that, just as conceptual necessity claims such as (11) could be seen as devices for embracing inferential connections between the applications of more basic concepts such as "east" and "west," sentences such as (13) and (14) should be seen as devices for embracing inferential connections between more basic concepts such as "bank employee" and "harm to no purpose" and practical thoughts about what to do, e.g. in the inferences:

I am a bank employee → so I shall wear a necktie to work

Repeating the gossip would harm someone to no purpose → I shall not repeat the gossip

where these are conceived as "language-exit" transitions because the termini are the expressions of intentions that—when the time comes, assuming things go normally—will result in the agent's doing something other than making further linguistic moves.

Leaving aside what to say about other normative terms, this surely cannot be the right metasemantic story about why 'ought' makes the semantic contribution that it does. For, as we saw above, not all ought-sentences are about what someone is to do. For instance,

(15)   Larry, who has had the most unfortunate time of things recently, ought to win the lottery.

does not seem to prescribe an action to Larry (or anyone else). So, although such sentences may be "normative" or "evaluative" in some broader sense, they are not plausibly construed as endorsing patterns of narrowly practical inference, i.e. inferences that end with a language-exit expressible with an "I shall . . . ".

Nevertheless, I think we can and should make sense of these sentences as nondescriptive. If we use my enhanced modal rule for 'ought' and think of it as taking propositional rather than prescriptional prejacents for these sentences, then (simplifying slightly) we will get truth conditions for these sentences that say that they are true just in case some proposition is true at all of some contextually determined set of possible worlds. When it next comes to a metasemantic interpretation of this result, we could view the relevant specifications of truth conditions as an initial spelling out

of a position in a space of implications, which determines an inference rule; and these sentences could then be seen as devices for embracing (acknowledging, endorsing) this rule. For example, (15) might be thought of as a device for embracing (and making explicit one's commitment to) the defeasible semantic validity of inferences of the form:

> Larry has had the most unfortunate time of things recently → Larry's winning the lottery would be most preferable (compared to the salient alternatives).

where the conclusion of this inference isn't itself a language-exit but a commitment to a certain constraint on one's preferences.

That's admittedly a very rough idea. The purpose of mentioning it here is just to indicate that full consideration of the types of sentences in which 'ought' figures suggests that it cannot be solely a device for endorsing inferential connections between propositional commitments and narrowly practical commitments in the way that Brandom suggests. That does not mean, however, that he's wrong about its being a metaconceptual device for endorsing inferential connections more generally, on the model of conceptual necessity. This is precisely how I think an inferentialist should interpret the enhanced modal rule for 'ought' spelled out above, which always involves determining the truth conditions of an ought-sentence in terms of a universal quantification over possibilia. This allows for an evaluation of the legitimacy of an embedded prescription across a range of norms as a special case but also the truth of an embedded proposition across a range of worlds as another special case.

Recall the apparent tension between the ideas that (i) 'ought's do not describe what is actually the case but prescribe possible action, thought, and feeling, and (ii) all declarative sentences deserve the same general semantic treatment, e.g. in terms of compositionally specified truth conditions. In §3, I argued that there's a representationalist resolution of this tension, which treats all declarative sentences as having the truth conditions that they have in virtue of their representing a way reality could be, but which also allows that ought-sentences don't describe what is actually the case because, as their truth conditions suggest, they represent how things stand with respect to possible worlds and norms. In this section, I have developed an alternative resolution of this tension. The inferentialist metasemantics says that all declarative sentences have the truth conditions that they have in virtue of the position in a space of implications that accepting the sentence commits one to, but it also allows us to draw a distinction between sentences playing a first-order descriptive role and those playing a higher-order metaconceptual role. It's not obvious where to draw the line between these, but empirical claims we think of as tracking features of our perceptual environment are a pretty good candidate for the former, and claims containing logical terms we think of as providing the conceptual scaffolding within which complex logical interrelations can be specified are a pretty good candidate for the latter. The truth-conditionalist semantic rule for 'ought' encourages us to think of it along the lines of a necessity modal, and I have suggested that this means we should

think of it as a metaconceptual device rather than a term for representing pieces of reality.

This carries an ontological benefit over the representationalist way of resolving the tension between (i) and (ii): we don't have to commit ourselves to the reality of possible worlds and possible norms in order to recognize some ought-claims as true. But, for this, it incurs the cost of explaining in more detail why language should have terms that function as metaconceptual devices and how to tell the difference between those terms and representational terms. Although I am not going to pay it here, it's a debt I view as worth paying. (However, I also admit to difficulty in knowing how to think about the exchange rate between such high-level theoretical costs and benefits.) In any case, my main claim is that the deciding factor between representationalism and inferentialism should not be which of ideas (i) and (ii) we are more impressed by but rather what we think of the relative merits of how these two views respect these ideas.

7 Comparison to Expressivism

I want to conclude by discussing a more familiar constellation of ideas that might also be thought to provide an alternative to representationalist treatments of normative sentences. This is born out of the noncognitivist thesis that normative claims express conative attitudes rather representations of reality, the sorts of attitudes that are supposed to class with desires rather than beliefs in their motivational potentials (or "directions of fit") in a Humean psychology of motivation. This noncognitivist thesis has received considerable attention and refinement over the past half century, resulting in a number of different positions now defended under the banner of "expressivism." Moreover, expressivism has been applied in several domains besides normative discourse (modal, epistemic, aesthetic, alethic, etc.) with importantly differing theoretical constraints. So, I find it far from clear what the core of expressivism is and how it relates to the inferentialist ideas explored above. As we already saw, some think of inferentialism as a kind of "global expressivism"—but that's not the view one finds in Sellars and Brandom. Moreover, many philosophers conceive of expressivism as premised on rejecting truth-conditionalist semantics in favor of a "psychologistic" alternative,[33] at least in its application to normative discourse—but that's not how I think we should think of inferentialism.

I'm not going to be able to discuss all possible (or even many) versions of expressivism here.[34] But I think it will prove useful to discuss one expressivist line of thought about 'ought' that I think comes closest to the inferentialist line of thought developed above. This will put me in a position to explain how some ways of developing the view might result in an alternative to inferentialism that has various problems while others

---

[33] This is, for instance, the way expressivism is conceived by critics such as Rosen (1998, pp. 387–88), Wedgwood (2001, p. 26), Kalderon (2005, pp. 66–82), and Schroeder (2010, p. 30).

[34] See Chrisman (2011) for my take on some historical and contemporary versions of expressivism.

might result in friendly amendments to inferentialism. Hopefully, this will also clarify the inferentialist view outlined above.

As long as we continue to observe the distinction between semantics and metasemantics, I think it should be clear that expressivism needn't be understood as a view about compositional semantics.[35] We can find a version of expressivism as part of a larger *metasemantic* view that has, until now in this paper, remained conspicuously off stage. This is the view that the truth conditions assigned by compositional semantics are articulations of what one who asserts one of these sentences ought to think. For instance, in the most basic case, it's plausible to say that one who asserts

(16)   Grass is green.

ought to think (judge, affirm) that grass is green. The idea is that someone who asserts this sentence but does not think this thought has violated the semantic rule associated with (16) in virtue of its core communicative role. So, rather than starting with the view that truth conditions articulate the way reality must be in order for a sentence to be true, an "ideationalist" metasemantics would encourage us to start with the view that semantic assignments to sentences tell us what idea one ought to have when one uses the sentence canonically, in order to conform with the core communicative norms associated with the sentence in the language.[36] This is one plausible way to flesh out the vague and popular suggestion that different sentences are canonical and conventional means for expressing different mental states.[37] Then, on this ideationalist view, the truth-conditionalist articulation of the subsentential elements of language could be viewed as an explanation of how an indefinite number of such expressive commitments are generated by a finite number of terms and logical forms.

Although ideationalism starts in a different place from representationalism, it's wholly compatible with the representationalist view that sentences such as (16) represent reality. For an ideationalist can say that the truth conditions of (16) specify the way this sentence represents reality as being by virtue of articulating the content of a particular kind of thought, a representational thought about the way reality is,

[35] Compare Suikkanen (2009), Chrisman (2012b), Yalcin (2012), Silk (2013), Charlow (2014a, p. 661), and Ridge (2014, chapter 1).

[36] Precursors and versions of this idea can be found in Grice (1989), Blackburn (1984), Davis (2003). Although defending a somewhat similar idea, I think Richard (2008) would say that it's wrong to think of these as *truth* conditions but that compositional semantic values can be interpreted as modeling a space of commitments to various thoughts.

[37] See Schroeder (2008, chapter 2) and Ridge (2014, §4.3) for further discussion of the expression relation appropriate for developing a form of metanormative expressivism. Both end up developing forms of expressivism with an account of expression very similar to the account I am using here. The key commonality is in construing the relation as a hypothetical linguistic norm, something like: *if* one uses the sentence to make an assertion, one ought to think the relevant thought *in order to conform with the core communicative rules of the language*. In Bar-On and Chrisman (2009), we argue that distinguishing the way in which *users* of a sentence express mental states from the way in which *sentence-types* express their semantic contents provides for a much simpler and more plausible explanation of the apparently distinctive connection between moral claims and motivational attitudes than is on offer in standard forms of metaethical expressivism.

which one who asserts the sentence ought to have in order to conform with the core communicative rules of the language.

Then, a "metasemantic" form of expressivism begins to come into view when we observe that, although the ideationalist can agree with the representationalist about sentences like (16), she needn't agree that all declarative sentences represent reality. For she can argue that some sentences are conventional means for expressing nonrepresentational thoughts, the content of which is specified by the truth conditions the correct semantic theory assigns to the sentence.[38] For example, an expressivist might argue that one who uses a sentence such as:

(17)   Stealing is wrong.

to make an assertion ought to think a particular thought, i.e. *that stealing is wrong*. Then, however, she will argue that the functional role of this thought in the psychology of motivation is directive of action rather than representative of reality.

This move obviously requires giving up on the idea that truth is something like correct representation, which is why most contemporary expressivists embrace some sort of minimalism about truth (though I think some kind of pluralism about truth would also serve this purpose).

But that raises an important question: how could the expression of a directive thought be something that one could doubt or know to be true, even in a deflationary sense? After all, paradigmatic directive mental states are desires, preferences, and plans, and these are not the proper objects of doubt and knowledge. In the present context, I think expressivists might insist that the characterization of directive thoughts as "desire-like" rather than "belief-like" is simply misleading and distracting. The important claim is that some thoughts are thoughts about reality, they represent reality as being a particular way; other thoughts are about what to do, they direct us to make reality become a particular way. The standard Humean view of the psychology motivation is that action always results from the cooperation of these two kinds of thoughts. We should allow expressivists to call both "beliefs" in many core cases; the important suggestion is that one of them has a directive rather than descriptive functional role in the psychology of human motivation. It is because of the way in which this role stands in contrast with the descriptive functional role of other thoughts that this kind of expressivism promises to carry whatever attractions were in earlier views

---

[38]   For similar ideas, compare Blackburn (1984, pp. 11–17), Silk (2013, p. 218), Ridge (2014, chapter 4), Perez Carballo (2014). Charlow (forthcoming) develops a metasemantic form of expressivism in the context of a dynamic "test" semantics for deontic and epistemic modals. This is importantly different in its rejection of the semantic explanations of truth-conditional semantics, but it still shares in conceiving of expressivism as a metasemantic thesis. Yalcin (2012, p. 156) characterizes his expressivism as a pragmatic rather than a semantic thesis. However, if that interpretation is going to carry the ontological benefits expressivism is supposed to carry, then I think the pragmatics needs to be part of a broader metasemantic theory—perhaps one that explains why sentences have the truth conditions they do in terms of what conversational moves they are canonical means for making. I believe this is consonant with the characterization of metasemantics in Yalcin (2014).

in the expressivist tradition. Moreover, if one can make sense of evaluating directive thoughts as being more and less reasonable, justified, or warranted, a proponent of this view would be on the way to making sense of the way that normative beliefs seem to be the objects of epistemic evaluations.[39]

So, at least one way to conceive of metanormative expressivism is as the combination of the Humean psychology of motivation with an ideationalist metasemantic interpretation of the truth conditions of normative sentences. More specifically, we will need to think of the truth conditions of normative sentences such as (17) as the articulation of directive rather than descriptive thoughts. We can give them truth conditions such as

[[(17)]] = T iff stealing is wrong

and think of this as an articulation of what one who asserts this sentence ought to think, in order to conform to the core communicative rules associated with this sentence. However, in their philosophy of mind, expressivists will argue that this thought plays a directive rather than a descriptive role in the psychology of motivation. By understanding expressivism as part of an overall metasemantic package, rather than a controversial thesis in compositional semantics, we could then grant expressivists access to all of the successes of standard truth-conditional approaches to compositional semantics, while leaving room in the philosophy of mind for them to make the Humean argument that normative thoughts are nondescriptive because they are directive.

There is a lot more we might want to know about metasemantic forms of expressivism, but the crucial question for me at this stage is whether this generates a competitor to the inferentialist metasemantic view about 'ought' outlined above or a possible partner for it in a coalition against representationalism about 'ought.' I think that's going to depend on what expressivists say about several further questions.

First Question: What sort of thoughts do metasemantic expressivists say are expressed by logically complex sentences? Those versed in debate about the Frege-Geach problem might grant that cleaving to a truth-conditionalist semantics and now attempting to develop expressivism in one's metasemantics provides for a better answer to this question than old-school versions of expressivism did. For the expressivist can now begin by saying that a sentence such as

(18)   Either stealing is common or stealing is wrong.

has the truth conditions of a disjunction, i.e. a logical form ensuring that it is true if either of its disjuncts is true. Still, if our metasemantic expressivist says that it

---

[39] Though we will have to reject Cuneo's suggestion that it is platitudinous that "... propositional attitudes display one or another epistemic merit (or positive epistemic status) such as being a case of knowledge, being warranted, being an instance of understanding, insight or wisdom and the like, only insofar as they are representative in some respect." (2007, p. 57) See Chrisman (2009, 2010, 2012c) for discussion of the prospects of nonrepresentational attitudes' achieving positive epistemic status.

has these truth conditions in virtue of expressing the disjunctive thought *that either stealing is common or stealing is wrong*, we can ask: what kind of thought is that? In the simple Humean psychology of motivation we were working with above, there are only descriptive and directive thoughts. And while we may want to say that this thought has descriptive and directive elements, the expressivist metasemantics seems to require us to be able to say what kind of mental state the whole thought represents. After all, as I glossed it above, the expressivist's way of securing the claim that normative sentences don't describe reality is by claiming that these sentences express directive thought rather than descriptive thought.

In response, expressivists might try to argue that logically complex sentences such as (18) are directive, but in an attenuated sense. The idea would be to treat them as the expression of conditional dispositions to infer, where inferring is conceived as a mental action. So, in this case, we might think of (18) as expressing a thought about what to do: if I come to reject one of the disjuncts I shall embrace the other.[40]

As someone sympathetic with inferentialism, naturally I think there's something importantly right about understanding the meaning of logically complex sentences in terms of their inferential interrelations with other sentences. However, I doubt using this attenuated sense of "directive" is going to let us draw the line between representational and nonrepresentational thoughts anywhere near where we want to draw it for the purposes of metanormative theory. For notice that even a paradigmatically descriptive sentence such as (16) might be thought to direct inference (e.g. if you accept this, do not infer any conclusion that entails that grass is not green!). Moreover, it would be highly odd to place

   (19)   Grass is green or roses are red.

   (20)   Grass is green and roses are red.

on the directive side of the ledger, given that we place each of the atomic parts on the descriptive side of the ledger.[41]

Because of this, the answer I think expressivists should give to the first question is that the Humean psychology of motivation is too stark; there are more than just descriptive and directive thoughts, and the thoughts expressed by logically complex sentences are neither purely descriptive nor purely directive. But, if expressivists gives that answer, then they owe us a new account of what distinguishes the class of sentences that have their truth conditions in virtue of expressing a representational thought and

[40]   This is one way to make sense of Blackburn's (1998, pp. 71–74; 2006) suggestion that logically complex thoughts "tie one to a tree."

[41]   I think a similar problem arises for Ridge's (2006, 2007, 2014) suggestion that logically complex sentences express hybrid states. The idea is that sentences such as (18) express beliefs with a descriptive content that is logically isomorphic to the normative sentence but they also express interlocking directive attitude (e.g. the endorsement of an ideal adviser or a normative perspective). The problem I see for this idea is that it ends up committing one to treating *all* logically complex sentences as expressive of a (partially) desire-like state of mind, or else it involves an ad hoc treatment of the role of logical particles in various logically complex sentences.

all of the other sentences that supposedly express nonrepresentational thoughts. If being directive isn't necessary for a thought to be nonrepresentational, what is?

I don't think it's obvious how expressivists should answer this question, but one way to look for an answer is to think about the downstream inferential consequences of the thoughts in question. For example, expressivists might insist that the thought *that grass is green* directly implies something about how reality is (i.e. without the aid of substantive auxiliary premises) but the thought *that stealing is wrong* does not directly imply something about how reality is. If that's right, there might be diverse reasons why some thoughts do not directly imply something about how reality is, but it would be the absence of this downstream inferential connection that marked out the class of nonrepresentational thoughts. Obviously, if expressivists go in this direction, their position begins to look like a coalition partner for rather than a competitor with the inferentialist view sketched above.

Second Question: What do metasemantic expressivists say about 'ought'? By asking this question, I don't mean to reopen the issue of whether a modal analysis of the truth conditions of ought-sentences is right (say, in comparison with attempts to reductively analyze 'ought' in terms of values or reasons). Indeed, since the sort of expressivists I have in mind are now on board with truth-conditional semantics, I shall assume they accept something like the enhanced modal rule outlined above. But what do they say it is in virtue of which the truth conditions predicted by this rule are (approximately) correct?

The initial response is clear: it's in virtue of the fact that ought-sentences are conventional and canonical means for expressing ought-thoughts. But what are ought-thoughts, descriptive or directive or do they fall in some yet-to-be-specified third category? Again, if we work with the stark Humean division between descriptive and directive thoughts, it may seem as if expressivists are going to say that ought-thoughts are uniformly directive. However, when we consider the data, that doesn't seem plausible. At least it seems pretty implausible to say that all ought-sentences have the truth conditions that they do in virtue of expressing a thought whose primary functional role is to provide the goal-setting rather than means-providing element in a Humean explanation of action. Consider:

(21)   I ought to have begun my higher education at a different college.

(22)   NATO ought not to expand any further.

I won't ever be in a position to begin my higher education again; I won't ever be in the position of NATO; so why think that the thoughts I would express by asserting these sentences have the motivational efficacy characteristic of directional thoughts in a Humean explanation of action? To be sure, an expressivist might insist that these sentences express something like plans or preferences *conditioned* on being in exactly the situation of my former self or of NATO. But that looks like a patch: something expressivists have to say rather than an intuitive gloss of the core communicative

function of sentences. So, this is a cost for the overall plausibility of the theory. Moreover, as we have already observed, there are non-agentive but at least weakly normative ought-claims, such as the teleological and evaluative ought-sentences (3)–(4) from above. Whatever connection there is between having the thought expressed by these sentences and being motivated to perform some specific act is going to be rather indirect and highly defeasible. We are right, I think, to question whether someone understands what she is saying if she says something of the form "All things considered, I ought to ϕ," and yet has no motivation to ϕ. However, it is much less plausible to think that competence with 'ought' more generally—even in many paradigmatic normative uses—requires any particular motivational propensities.

This is why I think expressivists should instead answer the question by arguing that many ought-thoughts fit into some third category of thoughts, neither purely descriptive nor purely directive.[42] Again, however, we should ask: what are they like and what about them makes them nonrepresentational? And again I think the inferentialist story sketched above might provide some helpful resources for expressivists. That is to say that expressivists could treat some thoughts (paradigmatically ones expressed with necessity modals) as not themselves representational but rather ways of encoding metaconceptual manipulations on more basic pieces of semantic content, resulting in thoughts explicating the structure of a space of inferential relations. Maybe there are other way expressivists could answer this question,[43] but obviously if they answer it this way, the view is again a plausible partner to rather than a competitor with the inferentialist view sketched above.

Third Question: If semantic content of sentences is to be explained in terms of the content of the mental states (or "thoughts") they express, as expressivists maintain, then what explains the content of these mental states? The ideationalist conception of truth conditions sketched above takes for granted mental content in order to explain semantic content, which means that anyone endorsing it owes a further account of mental content.

---

[42] Although Blackburn (1993, 1998) is deeply influenced by Hume, he finds in Hume a much richer diversity of types of mental states besides the descriptive and the directive, including all sorts of different projectable stances and mental dispositions.

[43] Another way, which I do not quite know how to categorize, is to follow Yalcin's treatment of claims of probability and epistemic modality (see especially Yalcin, 2007, 2011). He replaces the traditional Humean distinction between descriptive and directive thoughts with a distinction between mental states representing the world as being one way or another (modelable as conditions on worlds) and second-order properties of one's first-order representational states (not modelable conditions on worlds but rather as a function from sets of worlds). In a similar vein but within a dynamic framework, Charlow (forthcoming) distinguishes between information-carrying thoughts and thoughts that are instructions to "test" one's overall belief and preference state for certain structural features. The more general idea behind Yalcin's and Charlow's versions of expressivism is that there may be declarative sentences whose function is not to express representational beliefs, nor to express desire-like states, but rather to coordinate on a higher-order property of one's global mental state, e.g. credence distributions and preference orderings. Depending on how these are related to inferential commitments, this too might constitute a useful ally with rather than competitor with the inferentialist view sketched above.

Early expressivists were, I think, tempted to say that the mental states expressed by normative sentences are not thoughts with normative content; rather, all content is representational but we can have nonrepresentational attitudes towards such content (e.g. we can desire that stealing be less common). Then normative sentences were treated as conventional and canonical means for expressing such attitudes. If we proceed this way, expressivists could insist on a purely representationalist conception of mental content, but argue that it's not the content of but the type of attitude that makes a mental state nonrepresentational. However, I think it is wildly implausible to say that there are no thoughts with normative content (what then is one doubting, assuming, hoping, etc., when one doubts, assumes, hopes, etc. that stealing is wrong?).[44]

Another way expressivists could go at this point is to apply something akin to the sort of inferentialist account of content I sketched above to the case of mental content.[45] If this is right, the ideationalist conception of truth conditions might be correct as far as it goes, but the inferentialist view would be providing a deeper explanation of what content is in virtue of which sentences have the meaning that they do.

8 Conclusion

As I see things, the key issue between representationalists, inferentialists, and ideationalists is one about order of theoretical priority. Everyone should agree that some sentences describe things in our environment, that most sentences have inferential connections to other sentences, and that all sentences are vehicles by which we express our minds. The interesting question is which of these is most theoretically basic for understanding that in virtue of which sentences have the contents that they have. Here I have argued that each of these views can accept the predictions that truth-conditionalist approaches to compositional semantics give for 'ought'; in effect that it is a modal operator that quantifies over possibilia at which it evaluates embedded contents. Given this, I think we can take representation as the master concept and argue that these truth conditions reveal that ought-claims don't describe how things actually are but rather how things are with respect to possible worlds and possible norms. But we might also take inference as the master concept and argue that these truth conditions reveal that ought-claims are not ways of encoding a representation of modal space but rather ways of explicating inferential relations standing between other things we can say. It's tempting to see the resulting inferentialist view in the lineage of expressivism. And I've suggested that there may be sophisticated forms

---

[44] Because of this, I think appealing to off-the-shelf Fodorian (1987, 1990) or Millikanian (1984) or Dretskean (1981, 1988) accounts of mental content is not going to work for expressivists, since these all proceed in terms of representation relations between concepts and reality, whereas expressivists need something different for normative concepts. Moreover, I suspect going in this direction opens expressivists up to Dorr's (2002) wishful-thinking challenge.

[45] See Blackburn (2006) for an expressivist who ends up moving in the direction of inferentialism.

of expressivism that are naturally allied with inferentialism about 'ought.' However, it's important to the way I am thinking of inferentialism that inferential connections rather than conventional expressive potentials are most fundamental. This is the idea that I want to recommend as an interesting metasemantic account of that in virtue of which ought-sentences have the truth-conditional content that they do, which I see as providing new foundations for a form of antidescriptivism that has often been ignored in metanormative theory.[46]

Acknowledgments

References

Bar-On, Dorit and Matthew Chrisman (2009) Ethical Neo-Expressivism. In *Oxford Studies in Metaethics*. Volume 4, edited by Russ Shafer-Landau. pp. 133–65. Oxford; New York, NY: Oxford University Press.

Björnsson, Gunnar, and Stephen Finlay (2010) Metaethical Contextualism Defended. *Ethics*. 121. 1. pp. 7–36.

Blackburn, Simon (1984) *Spreading the Word: Groundings in the Philosophy of Language*. New York, NY: Oxford University Press.

Blackburn, Simon (1993) *Essays in Quasi-Realism*. New York, NY: Oxford University Press.

Blackburn, Simon (1998) *Ruling Passions: A Theory of Practical Reasoning*. New York, NY: Oxford University Press.

---

[46] I say "foundations" because I have not tried to motivate this idea for *all* claims one might want to classify as nondescriptive—not even for all normative claims, let alone evaluative claims. I see three general ways the foundations I have laid here with respect to 'ought' could be developed to address other normative and evaluative terms. First, one might pursue a priority view, according to which 'ought' is *the* fundamental normative/evaluative term, perhaps because when one is deciding what to think, feel, or do, all other considerations should be thought to bear on a question most perspicuously posed with an 'ought.' Compare Broome (2013, pp. 9–11 and chapter 4). Along these lines, one might then try to analyze other terms, insofar as they have normative uses, in terms of more fundamental normative 'ought's which one treats as metaconceptual devices along the lines I have sketched above. Secondly, one might pursue a no-priority view, according to which other terms with normative or evaluative uses deserve a different treatment, which is not somehow derivative of the treatment of 'ought' presented here, but which also does not provide a reduction going the other way for 'ought.' Along these lines, maybe the sort of treatment of 'ought' as a metaconceptual device presented here works for what is clearly one of our core normative terms; and we simply have to dig deeper to untangle the bilateral relationship between ought-claims and other normative and evaluative claims. Finally, one might want to maintain that other terms with normative import (e.g. 'right' and 'good') are themselves intensional operators. Along these lines maybe they are metaconceptual devices for endorsing different sorts of inferential commitments.

Blackburn, Simon (2006) The Semantics of Non-Factualism. In *The Blackwell Guide to the Philosophy of Language*, edited by Michael Devitt and Richard Hanley. pp. 244–52. Malden, MA, Oxford: Blackwell.

Blackburn, Simon (2013) Pragmatism: All or Some? In *Expressivism, Pragmatism, and Representationalism*, edited by Huw Price. Cambridge: Cambridge University Press.

Block, Ned (1987) Functional Role and Truth Conditions. In *Proceedings of the Aristoteian Society*. Supplementary Volume LXI. pp. 157–83.

Brandom, Robert (1994) *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.

Brandom, Robert (2000) *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press.

Brandom, Robert (2009) *Reason in Philosophy: Animating Ideas*. Cambridge, MA: Harvard University Press.

Brandom, Robert (2011) *Perspectives on Pragmatism*. Cambridge, MA: Harvard University Press.

Brandom, Robert (2013) Global Anti-Representationalism? In *Expressivism, Pragmatism and Representationalism*, edited by Huw Price. pp. 85–111. Cambridge: Cambridge University Press.

Bronfman, Aaron, and Janice Dowell (2016) Contextualism About Deontic Conditionals. In *Deontic Modality*, edited by Nate Charlow and Matthew Chrisman. Oxford: Oxford University Press.

Broome, John (2013) *Rationality Through Reasoning*. London: Wiley-Blackwell.

Cariani, Fabrizio (2013) 'Ought' and Resolution Semantics. *Noûs*. 47. 3. pp. 535–58.

Carr, Jennifer (2012) Deontic Modals Without Decision Theory. *Proceedings of Sinn Und Bedeutung 17*: Chemla, E., Homer, V., and Winterstein, G. (eds.) École Normal Supérieure, Paris, 8–10 September. pp. 167–82.

Carr, Jennifer (2014) The If P, Ought P Problem. *Pacific Philosophical Quarterly*. 95. 4. pp. 585–3.

Castañeda, Hector-Neri (1975) *Thinking and Doing: The Philosophical Foundations of Institutions*. Berlin: Springer.

Charlow, Nate (2013) What We Know and What to Do. *Synthese*. 190. 12. pp. 2291–323.

Charlow, Nate (2014a) The Problem with the Frege-Geach Problem. *Philosophical Studies*. 167. pp. 635–65.

Charlow, Nate (2014b) Logic and Semantics for Imperatives. *Journal of Philosophical Logic*. 43. 4. pp. 617–64.

Charlow, Nate (forthcoming) Prospects for an Expressivist Theory of Meaning. *Philosophers' Imprint*.

Chrisman, Matthew (2009) Expressivism, Truth, and (Self-)Knowledge. *Philosopher's Imprint*. 9. 3. pp. 1–26.

Chrisman, Matthew (2010) Constructivism, Expressivism and Ethical Knowledge. *International Journal of Philosophical Studies*. 18. pp. 331–53.

Chrisman, Matthew (2011) Ethical Expressivism. In *The Continuum Companion to Ethics*, edited by Christian Miller. Chapter 2. London: Continuum.

Chrisman, Matthew (2012a) 'Ought' and Control. *Australasian Journal of Philosophy*. 90. 3. pp. 433–51.

Chrisman, Matthew (2012b) On the Meaning of 'Ought'. In *Oxford Studies in Metaethics*. Volume 7, edited by Russ Shafer-Landau. pp. 304–32. New York, NY: Oxford University Press.

Chrisman, Matthew (2012c) Epistemic Expressivism. *Philosophy Compass*. 7. 2. pp. 118–26.

Chrisman, Matthew (2015) *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. New York, NY: Oxford University Press.

Chrisman, Matthew and Graham Hubbs. unpublished. The Pratitional Approach to Imperatives.

Cuneo, Terence (2007) *The Normative Web*. New York, NY: Oxford University Press.

Davis, Wayne (2003) *Meaning, Expression, and Thought*. Cambridge: Cambridge University Press.

Dorr, Cian (2002) Non-Cognitivism and Wishful Thinking. *Noûs*. 36. pp. 97–103.

Dowell, Janice (2013) Flexible Contextualism About Deontic Modals: A Puzzle About Information-Sensitivity. *Inquiry*. 56. 2–3. pp. 149–78.

Dreier, James (1996) Expressivist Embeddings and Minimalist Truth. *Philosophical Studies*. 83. 1. pp. 29–51.

Dretske, Fred (1981) *Knowledge and the Flow of Information*. Bradford, MA: MIT Press.

Dretske, Fred (1988) *Explaining Behavior: Reasons in a World of Causes*. Bradford, MA: MIT Press.

Finlay, Stephen (2009) Oughts and Ends. *Philosophical Studies*. 143. 3. pp. 315–40.

Finlay, Stephen and Justin Snedegar (2013) One Ought Too Many. *Philosophy and Phenomenological Research*. 86. 1.

Fintel, Kai von (2012) The Best We Can (Expect) to Get? Challenges to the Classic Semantics for Deontic Modals. Prsented at the Central APA, 17 February.

Fintel, Kai von and Anthony S Gillies (2011) Might Made Right. In *Epistemic Modality*, edited by Andy Egan and Brian Weatherson. pp. 108–30. Oxford: Oxford University Press.

Fintel, Kai von and Sabine Iatridou (2008) How to Say Ought in Foreign: The Composition of Weak Necessity Modals. In *Time and Modality*, edited by Jacqueline Gueron and Jacqueline Lecarme. pp. 115–41. Berlin: Springer.

Fodor, Jerry (1987) *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.

Fodor, Jerry (1990) *A Theory of Content*. Bradford, MA: MIT Press.

Frege, Gottlob (1879) *Begriffsschrift, Eine Der Arithmetischen Nachgebildete Formelsprache Des Reinen Denkens*. Halle a/S.: L. Nebert.

Geach, Peter (1982) Whatever Happened to Deontic Logic? *Philosophia*. 11. 1–2. pp. 1–12.

Gibbard, Allan (1990) *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.

Gibbard, Allan (2003) *Thinking How to Live*. Cambridge, MA: Harvard University Press.

Grice, H. P. (1989) *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Hale, Bob (1993) Can There Be a Logic of Attitudes. In *Reality, Representation, and Projection*, edited by Haldane John and Crispin Wright. Oxford: Oxford University Press.

Hanks, Peter (2015) *Propositional Content*. Oxford: Oxford University Press.

Hauser, Roland (1980) Surface Compositionality and the Semantics of Mood. In *Speech Act Theory and Pragmatics*, edited by Searle, J., Kiefer, F. and Bierwisch, M. pp. 71–96. Dordrecht: Reidel.

Heim, Irene (1982) The Semantics of Definite and Indefinite Noun Phrases. Ph.D. thesis. University of Massachusetts, Amherst.

Kalderon, Mark (2005) *Moral Fictionalism*. Oxford: Oxford University Press.

Kamp, Hans (1981) A Theory of Truth and Semantic Representation. In *Formal Methods in the Study of Language, Part I*, edited by Jeroen Groenendijk, Theo Janssen, and Martin Stokhof. pp. 189–222. Amsterdam: Mathematisch Centrum.

Kant, Immanuel (1787) *Kritik Der Reinen Vernunft*. Riga: Johann Friedrich Hartknoch.

Kolodny, Niko and John MacFarlane (2010) Ifs and Oughts. *Journal of Philosophy*. 107. pp. 115–30.

Kratzer, Angelika (1981) The Notional Category of Modality. In *Words, Worlds, and Contexts*, edited by Eikmeyer, H. J. and Reiser, J. Berlin: de Gruyter.

Kratzer, Angelika (1991) Conditionals. In *Handbuch Semantik/Handbook of Semantics*, edited by von Stechow Arnim and Dieter Wunderlich. Berlin, New York, NY: de Gruyter.

Macarthur, David, and Huw Price (2007) Pragmatism, Quasi-Realism and the Global Challenge. In *New Pragmatists*, edited by Cheryl Misak. 91. Oxford: Oxford University Press.

MacFarlane, John (2011) Epistemic Modals Are Assessment-Sensitive. In *Epistemic Modality*, edited by Andy Egan and Brian Weatherson. pp. 144–78. Oxford: Oxford University Press.

Millikan, Ruth G. (1984) *Language, Thought and Other Biological Categories*. Bradford, MA: MIT Press.

Peacocke, Christopher (1986) *Thoughts: An Essay on Content*. Oxford: Blackwell.

Peacocke, Christopher (1987) Understanding Logical Constants: A Realist's Account. *Proceedings of the British Academy*. 73. pp. 153–99.

Perez Carballo, Alejandro (2014) Hermeneutic Expressivism. In *New Essays in Metasemantics*, edited by Burgess, A. and Sherman B. New York, NY: Oxford University Press.

Portner, P. (2004) The Semantics of Imperatives Within a Theory of Clause Types. In *Proceedings of SALT* editor: Young, R. B. Northwestern University, 14–16 May. 14: pp. 235–52.

Portner, Paul and Aynat Rubinstein (2015) Extreme and Non-Extreme Deontic Modals. In *Deontic Modality*, edited by Nate Charlow and Matthew Chrisman. Oxford: Oxford University Press.

Price, Huw (2011) Expressivism for Two Voices. In *Pragmatism, Science and Naturalism*, edited by Knowles, J. and Rydenfelt, H. pp. 87–113. Frankfurt am Main: Peter Lang.

Price, Huw (2013) *Expressivism, Pragmatism and Representationalism*. Cambridge: Cambridge University Press.

Richard, Mark (2008) *When Truth Gives Out*. New York, NY: Oxford University Press.

Ridge, Michael (2006) Ecumenical Expressivism: Finessing Frege. *Ethics*. 116. 2. pp. 302–36.

Ridge, Michael (2007) Anti-Reductionism and Supervenience. *Journal of Moral Philosophy*. 4. 3. pp. 330–48.

Ridge, Michael (2014) *Impassioned Belief*. Oxford: Oxford University Press.

Rosen, Gideon (1998) Blackburn's Essays in Quasi-Realism. *Noûs*. 32. 3. pp. 386–405.

Schroeder, Mark (2008) *Being for*. New York, NY: Oxford University Press.

Schroeder, Mark (2010) *Noncognitivism in Ethics*. New York, NY: Routledge.

Schroeder, Mark (2011) Oughts, Agents, and Actions. *Philosophical Review*. 120. pp. 1–41.

Segerberg, Krister (1990) Validity and Satisfaction in Imperative Logic. *Notre Dame Journal of Formal Logic*. 31. pp. 203–31.

Sellars, Wilfrid (1953) Inference and Meaning. *Mind*. 62. 247. pp. 313–38.

Sellars, Wilfrid (1956) Empiricism and the Philosophy of Mind. In *Minnesota Studies in the Philosophy of Science*. Volume 1, edited by Feigl, H. and Scriven, M. pp. 253–329. Minneapolis, MN: University of Minnesota Press.

Sellars, Wilfrid (1958) Counterfactuals, Dispositions, and the Causal Modalities. In *Minnesota Studies in the Philosophy of Science*. Volume II, edited by Herbert Feigl, Michael Scriven, and Grover Maxwell. pp. 225–308. Minneapolis, MN: University of Minnesota Press.

Sellars, Wilfrid (1974) Meaning as Functional Classification. *Synthese*. 27. pp. 417–37.

Silk, Alex (2012) Modality, Weights, and Inconsistent Premise Sets. *Proceedings of SALT* Chereches, Anya (ed.) University of Chicago, Chicago, IL, 18–20 May 22. pp. 43–64.

Silk, Alex (2013) Truth Conditions and the Meaning of Moral Terms. In *Oxford Studies in Metaethics*. Volume 8, edited by Russ Shafer-Landau. pp. 195–222. New York, NY: Oxford University Press.

Silk, Alex (2014) Evidence Sensitivity in Weak Necessity Deontic Modals. *Journal of Philosophical Logic*. 43. 4. pp. 691–723.

Speaks, Jeff (2010) Theories of Meaning. In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter, 2010 edn.

Starr, William (2014) Mood, Force and Truth. *Protosociology*. 31. pp. 160–81.

Suikkanen, Jussi (2009) Metaethics, Semantics, and Metasemantics. Pea Soup Blog. [Online] Available from: http://peasoup.typepad.com/peasoup/2009/07/metaethics-semantics-and-metasemantics.html. [Accessed 10 November 2015.]

Swanson, Eric (2011) On the Treatment of Incomparability in Ordering Semantics and Premise Semantics. *Journal of Philosophical Logic*. 40. 6. pp. 693–713.

Thomasson, Amie (2007) Modal Normativism and the Methods of Metaphysics. *Philosophical Topics*. 35. 1 and 2. pp. 135–60.

Thomasson, Amie (2013) The Nancy D. Simco Lecture: Norms and Necessity. *Southern Journal of Philosophy*. 51. 2. pp. 143–60.

Unwin, Nicholas (1999) Quasi-Realism, Negation and the Frege-Geach Problem. *Philosophical Quarterly*. 49. pp. 337–62.

Veltman, Frank (1985) Logics for Conditionals. Ph.D. thesis. University of Amsterdam.

Veltman, Frank (1996) Defaults in Update Semantics. *Journal of Philosophical Logic*. 25. 3. pp. 221–61.

Wedgwood, Ralph (2001) Conceptual Role Semantics for Moral Terms. *The Philosophical Review*. 110. 1. pp. 1–30.

Wedgwood, Ralph (2007) *The Nature of Normativity*. Oxford, New York, NY: Oxford University Press.

Wedgwood, Ralph (2016) Objective and Subjective 'Ought'. In *Deontic Modality*, edited by Nate Charlow and Matthew Chrisman. Oxford: Oxford University Press.

Wertheimer, Roger (1972) *The Significance of Sense*. Ithaca, NY: Cornell University Press.

Williams, Michael (1999) Meaning and Deflationary Truth. *The Journal of Philosophy*. 96. 11. pp. 545–64.

Williams, Michael (2010) Pragmatism, Minimalism, Expressivism. *International Journal of Philosophical Studies*. 18. 3. pp. 317–30.

Williams, Michael (2013) How Pragmatists Can Be Local Expressivists. In *Expressivism, Pragmatism and Representationalism*, edited by Huw Price. Cambridge: Cambridge University Press.

Yalcin, Seth (2007) Epistemic Modals. *Mind*. 116. 464. pp. 983–1026.

Yalcin, Seth (2011) Nonfactualism About Epistemic Modality. In *Epistemic Modality*, edited by Andy Egan and Brian Weatherson. New York, NY: Oxford University Press.

Yalcin, Seth (2012) Bayesian Expressivism. *Aristotelian Society*. 112. pp. 123–60.

Yalcin, Seth (2014) Semantics and Metasemantics in the Context of Generative Grammar. In *Metasemantics: New Essays on the Foundations of Meaning*, edited by Alexis Burgess and Brett Sherman. Oxford: Oxford University Press.

Zvolenszky, Zsófia (2002) Is a Possible-Worlds Semantics of Modality Possible? *Proceedings of SALT* Jackson, B. (ed.) University of California, San Diego, CA, 8–10 March. 12. pp. 339–58.