Edinburgh Research Explorer

Knowing what one ought to do

Knowing What One Ought to Do [1]

Matthew Chrisman

University of Edinburgh
3 Charles St., Edinburgh EH8 9AD
matthew.chrisman@ed.ac.uk

6 April 2016

This paper considers two competing pictures of knowledge of what one ought to do—one which assimilates this to other propositional knowledge conceived as partial 'locational' knowledge of where one is in a space of possibilities, the other which distinguishes this from other propositional knowledge by construing it as partial 'directional' knowledge of what to do in particular circumstances. I argue that the apparent tension can be lessened by better understanding the contextualized modal-*cum*-prescriptive nature of 'ought' and enriching our conception of the kinds of possibilities within which we can locate ourselves.

I   *Introduction*

We often want to acquire more knowledge about what we ought to do. For example, we may know that we ought generally to look out for our friends, but should we sign a false affidavit to protect our friends provided we're confident it won't harm anyone? Or should we, prudentially, not risk getting caught? Or should we, morally, never lie under oath? When asking ourselves questions like these, we might weigh reasons for and against various answers in the attempt to determine which one is correct. In doing so, we want to come to _know_ what one ought to do rather than come to an answer, even the correct answer, in a haphazard or illogical way.

My main question in this paper is this: When we acquire such knowledge, what is it like? In very rough outline, the answer I want to encourage is that this knowledge is a special sort of modal-*cum*-prescriptive knowledge; it's knowing what is prescribed not just in one's actual situation given the normative demands actually one actually faces but rather across a range of possible situations given a range of normative demands that could be in place.

Recent discussions of moral epistemology has tended to focus on the mechanism or method that would lead to correct (because) warranted belief in a normative proposition. Do these beliefs have to be formed by a reliable process, do we have

---

[1] I appreciate helpful feedback on this material from audiences at the University of Edinburgh, the Aristotelian Society, Graham Hubbs and Michael Ridge.

a special faculty of intuition for knowing such normative truths, how does coherence with other beliefs fit in with warrant for normative beliefs? These are important and difficult questions that I won't endeavour to answer here. This is because I think their difficulty stems, in part, from a deeper unresolved tension between two programmatic pictures of what this knowledge could be like. Here I address this deeper tension. More precisely, I want to suggest that it can be lessened by making a few careful refinements of these pictures and situating them in their appropriate theoretical context. In my view, the key to doing so is recognizing the peculiar modal-*cum*-prescriptive character of knowledge of what one ought to do (which is a consequence of the peculiar modal-*cum*-prescriptive character of 'ought').

II   *Two Pictures and Some Stage Setting*

The first picture stems from the dominant *truth-conditionalist* framework for compositional semantics. The core idea is to explain the contribution that each part of a sentence makes to the meaning of the whole sentences in which it figures by determining what it contributes to a determination of the conditions under which these sentences would be true. These conditions are typically modeled as a restriction on a set of *possibilia*: the possible worlds that the sentence correctly describes. From this, we can generate attractive explanations of semantic relations between sentences (such as entailment and inconsistency) in terms of set-theoretic relations between the worlds modelling their contents. Moreover, an intuitive picture of propositional knowledge falls out: knowing that *S* is knowing whether the actual world is in the set of possible worlds correctly described by *S*. Propositional knowledge is a kind of descriptive-locational knowledge. It's like knowing which maps in an indefinite library of maps correctly represent the relevant part of reality.

It's important to recognize in this context that ought-sentences embed in similar ways to other declarative sentences (for example under sentential connectives, propositional attitude verbs, epistemic modals, etc.). Hence, if we are attracted to the semantic explanations of these (and other) phenomena offered by this framework, it seems that we should apply it to ought-sentences as well. (Otherwise we'll be committed to *ad hoc* explanations of the semantic operation of sentential connectives, propositional attitude verbs, epistemic modals, etc. when they embed 'ought's). But, if we do that, shouldn't we extend the descriptive-locational picture of knowledge to this case as well? Knowing that one ought to ϕ would then be construed as something like knowing that the actual world is amongst the worlds correctly described by the sentence 'One ought to ϕ'.

But some philosophers will demur: Isn't normative knowledge fundamentally different from factual knowledge? Didn't Hume teach us to distinguish thoughts about what *is* the case from thoughts about what someone *ought* to do; and didn't Kant teach us that the mode of understanding involved in knowing a theoretical fact about empirical reality is fundamentally different from that involved in knowing a practical fact about what one ought to do?

The intuition of difference evinced in these questions often seems to be about what metaethicists would now call 'all-things-considered normativity'. Roughly, the idea is that knowledge of what one ought all-things-considered to do is fundamentally different from factual knowledge of conditions of the world; the former may often be based on the latter, but because of the *practical* reasons-giving character of normative[2] knowledge, it cannot be knowledge of reality. That may be true, but the second programmatic picture considered here stems from a more basic intuition: Statements about what one ought to do[3]—even many of which aren't intended as or taken to be 'all-things-considered normative' are *prescriptive* rather than *descriptive*; they're for telling people what to do rather than what reality is like. This suggests that knowing that one ought to ɸ is practical-directional rather than descriptive-locational. It's like having someone draw starting and ending points on the map, telling one where to go from where one is.

To fix ideas, let's call these the 'locational' and 'directional' pictures of ought-knowledge stemming, respectively, from the (semantic) 'Framework' and the 'Intuition' (of a difference between description and prescription). The traditional cognitivist positions in metaethics cohere best with the locational picture and find support in the popularity of the Framework amongst philosophers of language. The traditional noncognitivist positions in metaethics are often motivated by the Intuition, and, insofar as the countenance normative knowledge at all, they seek it in something like the directional picture. That is to say that this traditional debate encourages us to choose which of the two pictures of ought-knowledge we find more compelling (in light of *inter alia* varying strengths of background commitments to the Framework and the Intuition). But ideally we wouldn't have to choose. The Framework is very powerful and the picture of knowledge that *p* as locating ourselves amongst possibilities where *p* holds is attractive, but the intuition that 'ought-to-do's are prescriptive rather than descriptive is very strong and the picture of ought-knowledge as providing at least *prima facie* practical direction is also attractive. Isn't there a way for some locational knowledge to be directional—for the Framework to be squared with the Intuition?

---

[2]What about other kinds of 'normative' knowledge, such as knowing what's best or when there are reasons or what is wrong? I think it's reasonable to think such knowledge is 'fraught with ought' at least in the weak sense that it is often partially prescriptive rather than entirely descriptive. So, it should be integrable into the picture I develop; however, I shall not argue for this point here and am happy to restrict my focus for the purposes of this paper to ought-knowledge.

[3]This is not meant to discount the possibility that some ought-statements are merely about what ought to be the case instead of about what someone ought to do. Plausible candidates includes so-called evaluative (or bouletic) 'ought's such as 'There ought to be no childhood death and disease' and predictive (or epistemic) 'ought's such as 'The storm ought to hit before midnight.' The semantics for 'ought' sketched below is intended to handle both species of 'ought' without positing ambiguity (see Chrisman (2012a) for arguments against ought-to-do/ought-to-be ambiguity) while making sense of the prescriptive character of 'ought-to-do's.

In what follows, I want to take three steps towards a positive answer to this question and, hopefully, a better picture of knowledge of what one ought to do. First, I will discuss the contextualist view that 'ought' is semantically sensitive to the context in which it is used. Although it has sometimes been obscured in metaethics, this is surely right.[4] The interesting issue here will be how much that helps with harmonizing the two pictures of ought-knowledge. Although the account of the meaning of 'ought' I favour is resolutely contextualist, I'll argue that contextualism alone doesn't do enough. Second, I will discuss the norm-relativist view that normative sentences are semantically special in that they should be assigned semantic values not at various possible worlds but rather relative to various possible norms. I think this isn't quite right, but I'll suggest that it helps us to break the connection between location and description: location amongst possible worlds _and_ possible norms might provide direction rather than mere description. Finally, I will argue for a kind of synthesis of these two ideas, suggesting that ought-sentences sometimes embed prescriptions (rather than propositions) and function semantically to evaluate the correctness of these relative to a range of possible norms (and worlds). The resulting view is one according to which 'ought' is a special kind of necessity modal, sometimes saying what prescriptions are correct across a range of possible norms and worlds. Knowing what one ought to do (morally, prudentially, …) is thereby portrayed as locating oneself in a space of possible world-norm pairs, partially understanding what reality is like _and_ what the relevant norms prescribe in those circumstances.[5]

III   Semantic Contextualism

The received view in linguistic semantics is that 'ought' is a modal verb, and like other modal verbs it's meaning is sensitive to context of use in several distinctive ways. One can get a rough and ready appreciation of the motivation for this view by just considering all of the things that a statement made with the sentence 'Jerry ought to win his race' could mean given varying contexts of use: (moral) _that he is morally obliged to keep his promise to do so,_ (prudential) _that his life would go best by winning,_ (teleological) _that it would promote his political career to stand in an election he's sure to win,_ (epistemic) _that it would now be highly unexpected if his boat didn't drift across the finish line before the competitors,_ (bouletic/evaluative) _that it'd be generally best if he happened to win this time…_

---

[4]I argue for this in more detail in Chrisman (2012b), and I critically evaluate several proposals for reducing the meaning of 'ought' to other normative or evaluative terms (such as 'reason' or 'what's best') in Chrisman (2015 ch. 2).

[5]This is not, I should note, to promise a reductive account of what it is for norms to prescribe or any sort of explanation of why some prescriptions have all-things-considered normative force for us. I think it is incumbent on a semantics for 'ought' to explain its canonical use in prescribing rather than describing, but semantics is not the place to proffer accounts of the sources of normativity any more than it is the place to proffer explanations of the sources of necessity. At the very end of the paper, I'll briefly mention some ways I see the semantics for 'ought' sketched here as fitting with broader issues in metanormative theory.

Clearly this abstract sentence-type must interact with the context in which it is used in order to determine a specific meaning, such as one of these.

This is relevant here for two reasons. First, it suggests that not all ought-claims are prescriptive. For not all of the possible meanings of an ought-sentence bear on what one is to do (for instance, the 'epistemic' and 'bouletic' readings of the sentence above don't seem to have anything special to do with what Jerry is to do.) Second, the fact that 'ought' is a context-sensitive modal means that applying the truth conditionalist semantic framework outlined above to ought-sentences is more complicated than applying it to ordinary subject-predicate sentences. So, perhaps there is room in this complication for alleviating some of the tension between our two pictures of normative knowledge.

The standard way to extend the Framework to ought-sentences is to treat 'ought' as a necessity-like modal operator, taking an embedded 'prejacent' proposition and evaluating whether it holds true across a contextually determined set of possible worlds, where the contextually determined 'flavour' (moral, prudential, teleological, epistemic, bouletic, etc.) of this 'ought' partially determines the set of worlds. The basic idea is to say that sentences of the form 'ought[$\underline{p}$]' are true just in case $\underline{p}$ is true at all of the relevantly accessible possible worlds. And what counts as 'relevantly accessible' is fixed by context of use in a complex way usually depending on what background conditions and standards are assumed.[6] Moreover, context is also thought to be relevant for recovering the precise value of $\underline{p}$ from the surface sentence.

Given this, if we follow standard practice and model $\underline{p}$ as the set of possible worlds where $\underline{p}$ is true, we can extend the Framework to model the truth condition of a sentence of the form 'ought[$\underline{p}$]' as the set of possible, worlds relative to which all relevantly accessible worlds are worlds where $\underline{p}$ is true. Although this surely needs refinement as a thesis of natural language semantics[7], it seems to be a pretty good starting point for explaining the diversity of things an ought-sentence can mean, and it more or less retains the Framework's prospects for displaying the compositional structure of a language and explaining semantic relations such as entailment and inconsistency in terms of set-theoretic relations.

How might that help with the tension we observed at the beginning? One suggestion worth considering is that, because 'ought' is context sensitive in the way it selects a set of possible worlds where it evaluates the prejacent proposition, perhaps context can mark out some 'ought's as special in being prescriptive for the agent in question. For example, perhaps one of the things that 'Jerry ought to win his race' could mean is that, in all possible worlds where the relevant moral demands, prudential advice, … are followed, the following proposition is true: *Jerry's counterpart wins his race*. If that's right, then although there are other

---

[6]Cf. Kratzer (1981).

[7]For example, it doesn't respect the apparent weakness of 'ought' compared to 'must', and it faces a number of tricky problems having to do with evidence-sensitivity, embedding in conditionals, and dilemmas. There is a vast literature on these topics; some of the most recent contributions are in Charlow and Chrisman (forthcoming).

things this sentence could mean in different contexts, in some contexts, 'ought' has a distinctively prescriptive meaning. That may be the case even while we continue to use the Framework for articulating the meaning of this and all other ought-sentences.

Does that resolve the tension between the locational and directional pictures of ought-knowledge? I think the contextualist suggestion represents a step in the right direction, but I still worry that it doesn't properly respect the Intuition that ought-knowledge can be distinctively directional in a way that ordinary descriptive knowledge is not.

For, the simple contextualism above suggests that, once context has fixed the flavour of an 'ought', the sentence in which it is used does specify a condition of reality under which it would be true, albeit a modal condition—roughly, being such that the prejacent proposition is true at all of the worlds that are 'accessible' in the contextually relevant way. Knowing what one ought to do is still a matter of understanding something about where the actual world is in a space of possible worlds, albeit in a somewhat more complex way that involves knowing which further worlds are 'accessible' from the actual world. And sometimes what makes a world 'accessible' is whether the relevant normative demands in the actual world are satisfied by one's counterpart in that possible world.

But there's something odd about this idea. The Intuition of difference described above implies that knowledge that I ought to ϕ can provide direction, telling me, that ϕ-ing is the thing for me _to do_, at least in some respect, which seems different from knowing that, in all possible worlds where the relevant norms aren't violated, it is true that my counterpart ϕs. Why would the shenanigans of my counterpart in other possible worlds be relevant to my practical deliberations about what to do in the actual world?[8] In terms of the previous metaphor of maps, on this contextualist view, it is as if some of the maps that correctly represent a corner of reality also have flags on them indicating that a goal is accessible (that is some possible person like me gets there). But these flags do not by themselves capture the way in which ought-knowledge can actually direct _my_ action.[9]

IV   Norm Relativism

Let's turn the crank a bit further. Rather than using a semantic model that treats declarative sentences (in context) as having truth values relative to possible worlds, some philosophers of language think we should use an enriched version of the possible worlds model which evaluates declaratives relative to other kinds of parameters. There are a variety of ways this might go, but the core idea is to

---

[8]Kripke (1980, 45).

[9]Compare Korsgaard (2008: 315): "If to have knowledge is to have a map of the world, then … the ability to act [well] cannot be given … by having little normative flags added to the map of nature that mark out certain spots or certain routes as good. You still have to know how to use the map before the little normative flags can be any good to you.)"

say instead that declarative sentences hold (or don't) at things like world-evidence pairs (to handle probabilistic sentences), world-time pairs (to handle future contingent sentences), at world-standards pairs (to handle sentences of personal taste), or world-individual pairs (to handle distinctively _de se_ and _de te_ sentences).[10]

An extension of this idea is available for normative sentences.[11] Suppressing any other semantic relativity, we might say that sentences hold (or don't) relative to world-norm pairs, depending on what propositions are true at the world parameter and what possible states of affairs are 'correct' relative to the norm parameter (where often these two things interact). The content of any declarative sentence _S_ can then be modeled as a set of world-norm pairs: the set of pairs of possible worlds and possible norms relative to which _S_ holds. When it comes to purely descriptive sentences, the norm-relativist's idea is to think of them as norm-invariant, such that they hold at a world-norm pair just in case they are true relative to the world parameter. Then traditional truth conditions, modeled as a set of possible worlds, are predicted as a special case of the more general enriched semantic model. This allows the norm-relativist to say, in contrast, that purely normative sentences are world-invariant, such that they hold at a world-norm pair just in case what they demand is correct relative to the norm parameter. (Out of this falls an explanation of why high-level normative propositions often seem to be necessary: whether they hold doesn't depend on any factual conditions.) In this enriched framework, all other declarative sentences will be treated as mixed. This means that they hold depending on both the way things are factually and the way things are normatively.

That's rather abstract. So let's consider an example: _S_= 'Sally ought to devote her life to philosophy,' and let's assume context renders this 'ought' prescriptive (in whatever sense is relevant to the tension between the locational and directional pictures of ought-knowledge above) rather than bouletic, epistemic, etc. On the norm-relativist view, _S_ holds at a world-norm pair <_w_,_n_> just in case the factual conditions of _w_ are such that the state of affairs where Sally devotes her life to philosophy is correct relative to _n_. That means the content of _S_ can be modeled as the set of all such world-norm pairs. This is supposed to be a relatively conservative extension to the way we previously modeled the content of a sentence _S_ as the set of worlds where _S_ is true. However, now we use a broader notion of 'holding' relative to more complex parameters. This allows for truth-relative-to-world and correctness-relative-to-norm as special cases.[12]

---

[10]Often this is developed as a form of 'truth-relativism' in that these enriched notions of 'holding' are still treated as 'holding true'. Cf. MacFarlane (2014) amongst others. The relationship between the theoretical notion of holding true at a series of parameters deployed in a semantic model and _truth_ is not at all clear, and nothing in the norm-relativist idea I go onto describe requires thinking of the view as a view about truth. For example in Gibbard (1990) something like this idea was used for normative sentences but it was also denied that normative sentences have truth values.

[11]Cf. Gibbard (1990, ch. 5), Silk (2013), MacFarlane (2014, ch. 11).

[12]It's worth noting that this approach is neutral on 'relativism' in the traditional metaethical sense, that is whether there is one (perhaps complex) universally correct

I think the relativist idea is intriguing. Semantically, it's more or less isomorphic to the standard possible worlds model for truth-conditional semantics, so it promises similar resources for explaining semantic relations between sentences and amongst their parts in set-theoretic terms, but it also provides a way to respect the Intuition that ought-sentences are not descriptive (even of complex modal conditions of reality).[13] For, on this view, these sentences don't have traditional truth-conditions; whether they hold at a world-norm pair depends on the norms rather than (only) how things are descriptively.

Nevertheless, I think we shouldn't be completely satisfied with norm-relativism. For one thing, although there are answers to this question worth exploring, It remains highly unclear what it means on this view when a normative sentence is embedded under a propositional attitude verb or epistemic modal. For instance, if someone says 'It might be true that we ought sign a false affidavit to protect our friends,' the norm-relativist cannot treat 'might' as it is typically treated in terms of a proposition being true at some possible worlds in the set of accessible worlds. The obvious alternative is to treat 'might' in terms of a proposition 'holding' at some world-norm in the set of accessible world-norm pairs. But how do we determine which are the accessible world-norm pairs when the sentence is embedded under 'might'?

More importantly, I suspect relativism still obscures the way ought-knowledge can be directional. For it seems to make understanding ought-sentences into something like understanding what possible states of affairs count as correct relative to some norms, rather than understanding what the sentence tells the relevant agent to do in various possible conditions. It's as if some of the maps have regions coloured red and green, not to represent a factual condition of reality but to mark out which parts of reality accord with the relevant norms and which don't. While telling someone what to do is related to dividing states of affairs into correct/incorrect, it's not obvious these are equivalent.

V   Reapplication and Extension

In this section I want to reapply the core idea of norm-relativism and extend the core idea of semantic contextualism. I think this provides for an attractive synthesis of these ideas and a better way to resolve the tension between the two programmatic pictures of ought-knowledge from the beginning.

As I have already mentioned, declarative ought-sentences have all of the linguistic hallmarks of being truth-apt in exactly the same way as any other

---

normative standard. Whether a normative sentence holds relative to a world-norm pair is parallel to the way we previously treated sentences as true relative to a world. That does not imply that there are multiple incompatible norms in play in any given context any more than the standard possible worlds model implies that there are multiple incompatible worlds considered actual in any given context.

[13]As Dreier (2009) shows, there are formally isomorphic 'genuinely' relativist and expressivist implementations of this semantics, and the difference between them comes in what they say about indirect contexts and disagreement.

declarative sentences, which makes it prima facie odd to assign them complex holding-values at world-norm pairs rather than ordinary truth-values at worlds. In contrast, however, *imperatives* don't have all of the linguistic hallmarks of being truth-apt. Indeed it's highly weird to think about imperatives as true or false at possible worlds; and they don't embed under propositional attitude verbs or epistemic modals. Nonetheless, imperatives do seem to stand in semantic relations (something) like entailment and inconsistency, and they seem to be semantically composed out of more basic parts. Moreover, imperatives can combine with declaratives under some sentential connectives to form logically complex mixed sentences.[14] So, any full semantic theory will need a way to model the contents of imperatives and the semantic relations in which they stand.[15]

Developing and defending such a view will be a highly complex affair, but one way we might begin the story is by applying something like the norm-relativist approach sketched above to imperatives in the first instance rather than normative declaratives. As we standardly think of declaratives (or their contents: propositions) as true (or not) relative to possible worlds, we might think of imperatives (or their contents: prescriptions?) as correct (or not) relative to possible norms.[16] More precisely, since the status of an imperative often depends on the factual conditions of the situation to which it is supposed to apply, we find some traction for a semantics and logic of imperatives in a semantic model that assigns imperatives (in context) correctness values relative to world-norm pairs. Then we could try to explain semantic relations such as entailment and inconsistency as they hold between imperatives in terms of set-theoretic relations amongst sets of world-norm pairs. At base, this is just the norm-relativist idea reapplied to imperatives.

Generalizing, the idea would be to enhance the semantic Framework with which we began so that the semantic content of a declarative *or an imperative* sentence *S* determines a set of world-norm pairs relative to which *S* holds. But now we

---

[14]This is part of the difficulty a satisfaction-condition semantics for imperatives; it's unclear what results when satisfaction conditions are combined with truth conditions under a sentential connective such as disjunction.

[15]See Charlow (2014) for an excellent discussion of this in the context of compositional semantics. For further conceptual arguments against accounts of the meaning of imperatives that give them propositional contents or meanings derivative of propositional contents, see Chrisman and Hubbs (unpublished).

[16]What is a norm and how does it determine the correctness of a prescription? There are several ways we might answer this question, many of which are going to be formally adequate for the purposes of compositional semantics. What distinguishes them is how much they illuminate linguistic competence and fit with one's ancillary metanormative commitments. (More on this towards the end of the paper.) But a simple answer we can start with is that norms are higher-level prescriptions that entail lower-level prescriptions given in various factual circumstances. (Note that this question is exactly parallel to the question of what are *worlds* and how do they determine the truth of a proposition. Again there are several ways we might answer this question, many of which are going to be formally adequate for the purposes of compositional semantics but distinguishable in light of how they illuminate linguistic competence and fit with other philosophical commitments.)

recognize two species of 'holding' depending on whether the sentence is declarative or imperative. In case it is declarative, it holds relative to a world-norm pair iff the proposition it expresses is contained (i.e., true at) in the world parameter. In case it is imperative, it holds relative to a world-norm pair iff the prescription it expresses is contained (i.e., correct relative to) in the norm parameter (given whatever is true at the world parameter). The upshot is an enhanced semantic model where the holding-conditions of declarative and imperative (and mixed) sentences can be modelled as a set of world-norm pairs: the set of pairs of possible worlds and possible norms where the sentence holds.

So far that has nothing specifically to do with the semantics of 'ought', but it will do if we refine the contextualist view about 'ought' sketched above in the right way. Above I mentioned that, on the standard contextualist view about modals, context is relevant for recovering the precise prejacent proposition from the surface sentence. The prejacent proposition is the embedded content a modal like 'ought' is standardly thought to operate on, evaluating whether it is true across a the range of 'accessible' worlds. The suggestion I now want to make is that we extend this idea and allow that 'ought' can operate on a more complex range of prejacents—both propositions and prescriptions.

When 'ought' operates on a prescriptive prejacent, (as a first pass) we'll say that it evaluates whether this prescription is correct across a range of 'accessible' world-norm pairs. This is like we previously said (also as a first pass) that, as a necessity modal, 'ought' evaluates whether its propositional prejacent is true across a range of 'accessible' worlds. Indeed, the new model continues to predict the old result as a special case. My idea is that sometimes we should continue to understand 'ought' as operating on propositional prejacents (this is how it will work for epistemic and bouletic 'ought's); it's just that I now want to allow that there are other times where we should understand 'ought' as operating on prescriptive prejacents. The surface sentence alone won't determine which it is; for this we also need context of use. Prescriptive 'ought's are ones which context determines to embed a prescription as their prejacent. This is the sense in which this is an refinement of the standard contextualist view about the semantics of 'ought'.

As a semantics for 'ought', the preceding is still rough in several important respects, but I hope the basic idea is clear enough: we interpret 'ought' as a special sort of necessity modal, one which doesn't just have the function of evaluating the truth of prejacent propositions at a set of possible worlds but which also has the function of evaluating the correctness of prescriptions at a set of possible norms (given what is true at possible worlds). This is what I had in mind above when I said that we need to recognize the peculiar modal-*cum*-prescriptive character of 'ought'.

I suggested that recognizing this would help to alleviate the tension between the locational and directional pictures of ought-knowledge stemming from the Framework and the Intuition. Here's how: On this view, understanding an ought-claim can be thought of as having an ability to correctly divide a space of possibilities into those relative to which its prejacent holds and those relative to which it doesn't. However, since we have enhanced our conception of prejacents and correlatively what it is for them to 'hold' relative to some possibility, we can

now recognize two special cases depending on whether the prejacent is a proposition or a prescription.[17] That means that, although we'll continue to think of knowledge of what one ought to do as a sort of modal knowledge, it will *not* be construed as the ability to locate (only) the actual world in a space of possible worlds. We can still use the idea of locational knowledge, but now this knowledge involves locating agents in a space of possible world-norm pairs. Knowing that one ought to φ is knowing what one is to do across a range of possible worlds given a range of possible norms that apply.[18]

The key idea here is that by allowing prescriptions into the analysis of ought-claims, we generate resources for capturing the distinctive directional aspect of a particular species of modal locational knowledge. It's like having a temporally extended map with a dynamic navigational system capable of calculating where you are and telling you what to do to get where you're supposed, in some respect, to be.

VI   Conclusion

Modelling the contents of declarative sentences as sets of possible worlds does a lot of important work in compositional semantics, and it encourages us to conceive of propositional knowledge as partial knowledge of where the actual world is in a space of possible worlds. However, there seems to be a very important distinction between knowledge of describable conditions of reality, which is like having an accurate map, and knowledge of what one ought to do,

---

[17]I discuss the challenge of capturing the distinction between ought-to-do and ought-to-be senses of 'ought' in much more detail in Chrisman (2012a) and Chrisman (2015 ch. 5). There are several interesting alternative proposals worth considering, including the possibility of stit-propositions and even that 'ought' is syntactically ambiguous. Here I am simply gesturing at my preferred way through this debate, in order to show how it can be deployed in a refined picture of ought-knowledge.

[18] In fn. 16, I suggested there are many ways to develop the idea of a prescription's being correct relative to a norm that are compatible with using this idea in our semantics for 'ought'. One might conceive of norms as high-level general prescriptions, such as we might express with 'Maximise happiness!' or 'Act for genuine reasons!'; but one might instead conceive of them as high-level general ought-propositions of the right flavour, such as we might express with 'One ought always to maximize happiness' or 'One ought always act for genuine reasons.' Here, for the purpose of urging the integration of the locational and directional pictures of ought-knowledge, I have wanted to remain as neutral as possible about this issue (though I have been working mostly tacitly with the former conception). However, we should recognise that, even if both of these conceptions are usable the purposes of semantics, neither is an ideal place for metanormative theory to terminate. For they immediately invite the questions: what makes it the case that a high-level general norm 'holds'; and how do we know which of these norms 'hold'? Treating norms as high-level prescriptions threatens to trivialise the holding of the highest-level norms (for almost any prescription will hold relative to itself); but treating norms as high-level ought-propositions threatens to render the treatment of 'ought' suggested here uninformatively circular. In the following section, I say a bit more about the issue while remaining unsatisfied with my ability to answer these questions. (Thanks here to Kathryn Lindeman and Adrian Moore for helpful discussion.)

which is like having directions for how to navigate within the space that is mapped. I have argued that this tension can be alleviated by enhancing the possible worlds framework in a way that is independently motivated by the need to account for the semantics of imperatives. If we do this, then I think we can embed the results in a more sophisticated contextualist semantics for 'ought'— one that allows us to make sense of the idea that some locational knowledge is directional by recognizing our ability to locate ourselves amongst a space of possible norms and know which prescriptions they legitimate (given our factual circumstances).

I want to close by responding to two related questions so far left open by my discussion. I doubt my responses are adequate, but I hope they give some sense of how I am thinking about the theoretical landscape at the intersection of semantic theory, normative theory, and metanormative theory.

First Question: Given that there are indefinite many possible norms, what determines which norms are such that the prescriptions they render 'correct' have reason-giving or 'normative' force for particular agents in particular circumstances?

Second Question: If, in the end, there is a formula determining the _truth_ conditions of ought-sentences, doesn't that mean that statements deploying these sentences are descriptive after all?

My answer to the first question is mainly negative: Nothing about the meaning of 'ought' can tell us what one ought all-things-considered to do. A general contextualist semantics for 'ought' might help us to understand how context of use determines some ought-statements to be distinctively relevant to practical deliberation (or even to be practically deliberative in a distinctively all-things-considered way). And my particular version of the contextualist semantics might help us to understand that such deliberative contexts make the prescriptive character of various norms most salient. However, determining which norms are 'normative' for various agents requires first-order normative reflection about what kinds of considerations do and don't have reason-giving force.

By contrast, I think the second question encourages further positive reflection on what we're up to when we seek to assign truth conditions to declarative sentences as part of the project of compositional semantics. In my view, this is a _metasemantic_ question rather than a question in first-order semantics. It's not a question about what the correct semantics is, for example, for 'ought' but rather about what it is in virtue of which a semantics like the one I sketched above might be (approximately) correct. In making a specific semantic proposal, what kind of nonsemantic facts is one trying to identify, characterize, or model that would ground the semantic facts we observe (for example about entailment, inconsistency, and composition)?[19]

---

[19]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. By using

Taking recent metaethics as a guide, I think we can discern two very programmatic answers to this question to which I would like to add a third for consideration.

First, we can follow a _descriptivist_ route and, as the question suggests, understand the truth conditions assigned to declarative sentences by our best compositional semantics as an attempt to spell out how that sentence describes reality as being. This means that one plausible way to interpret the account above is as a determination of how ought-sentences describe reality as being. How they describe reality being depends on which 'flavor' of 'ought' is relevant in the context, which in turn determines (assuming the semantics above) which set of possible world-norm pairs count as accessible. But, in the end of the day, the idea would be that ought-claims describe reality as being such that the embedded prejacent holds at a variety of world-norm pairs. On a realist manifestation of this view, possible worlds and possible norms are parts of reality, ought-claims are attempting to describe them, and our semantics for these sentences is part of an attempt to spell out what in reality would make these claims true. The direction provided by distinctively prescriptive 'ought's would need to somehow derive from the complex modal pieces of reality described by those sentences.

Second, we might follow an _ideationalist_ route and argue that truth-conditional semantics is an attempt to spell out what one ought, by virtue of the core communicative rules of the language, to think when one uses a sentence to make an assertion.[20] This is one way to make sense of the idea that various sentences are canonical vehicles for expressing thoughts with particular contents. For example, we might say (assuming the semantics above) that one who asserts "Bill ought to kiss Lucy" has expressed the thought (ought to think that) the prescription _Bill, kiss Lucy_ is correct across the contextually relevant possibilities (worlds and norms). In general, this is an alternative way to understand why sentences have the truth conditions our best semantic theory says that they have. And it makes room for a non-standard form of expressivism[21] because, although it assigns truth conditions to ought-sentences, it leaves open whether the thoughts they express are descriptive thoughts about how reality is or conative

---

language of 'in virtue of' here, I mean to indicate that we're not looking for mere interpretations of formal machinery (though that may come with particular metasemantic theories), rather we're looking for very general _grounding_ explanations that seek to identify the kinds of nonsemantic facts that undergird the semantic facts postulated by empirical semantics (as opposed, for example to causal or genealogical explanations or mere supervenience bases). Compare Yalcin (2014, 18). Some (for example Speaks (2010)) may prefer the term 'foundational theory of meaning' or 'theory of the foundations of meaning' for what I am calling _metasemantics_.

[20]Precursors and versions of this idea can be found in Grice (1989), Blackburn (1984), Davis (2003).

[21]It's non-standard because it conceives of expressivism as a thesis in metasemantics rather than compositional semantics. As such it doesn't face the sorts of objections to expressivism's adequacy as a semantic thesis pressed by _inter alia_ Rosen (1998) and Schroeder (2008). Expressivism as a metasemantic thesis is suggested by Suikkanen (2009) and Chrisman (2012a); a particular 'hybrid' version is worked out and defended by Ridge (2014).

thoughts capable of motivating action. Expressivist might then argue on independent grounds having to do with the psychology of motivation and the motivational character of normative thinking that it is the latter (at least in cases where the 'ought' is moral or all-things-considered).

Finally, we could start instead with the _inferentialist_ idea that a specification of the truth conditions of a sentence is 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.[22] In this way, we might see an assignment of truth conditions to a sentence as 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. Then, if I may cryptically register an idea I try to make better sense of elsewhere[23], the metaethically and ontologically interesting distinction wouldn't be between descriptions of reality and expressions of conative thoughts motivating action but rather between descriptions of reality and endorsements of inference rules characterizing the logical structure of the concepts we use to think both about the actual world and what to do in it. If we go on to treat modal claims as a species of broadly logical claims then we could hold that, although they can be thought of as characterizing a set of _possibilia_ that are 'accessible' from the actual world, that characterization is not an ontologically committing _description_ of modal space, nor an _expression_ of a conative attitude, but rather an inferentially committing _endorsement_ of the logic of particular concepts. The rough idea is that, as modal claims, 'ought's sometimes function to endorse inferential relations to prescriptions implicitly wrapped up in some of our other concepts. When we think that knowledge that one ought to ф in circumstances _C_ settles the question of whether to ф in _C_, that ought-knowledge is all-things-considered normative for the agent.

For my part, I suspect that the notions of _describing reality_, _expressing our minds_, and _committing to inferences_ are all crucial to developing a full theory of meaningfulness. But because I think the latter notion is most fundamental in the order of explanation, I favor the third metasemantic account of what it is in virtue of which ought-sentences have the truth conditions that they do. And this is why I think assigning truth conditions to them, as I suggested we do above, doesn't by itself commit me to descriptivism (or expressivism) about them.

---

[22]This is inspired by Sellars (1953; 1974) and Brandom (1994, 2008), but their views on truth-conditional semantics are less than clear (often they seem to reject it, though it is not clear what they propose as an alternative approach to compositional semantics, as opposed to metasemantics).

[23]In my 2015, ch. 5, and 2016, I explain how modal operators can, on the inferentialist view, be treated as second-order (or 'metaconceptual') devices for making explicit and endorsing the logico-semantic relations between first-order terms and concepts.

References

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

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

——— 2008: *Between Saying and Doing*. New York: Oxford University Press.

Charlow, Nate 2014: 'Logic and Semantics for Imperatives'. *Journal of Philosophical Logic*, 43 (4), pp. 617-664.

Charlow, Nate, and Matthew Chrisman, eds. 2016: *Deontic Modals*. Oxford University Press.

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

——— 2012b: 'On the Meaning of "Ought"'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics, Vol. 7*. New York: Oxford University Press, pp. 304–32.

——— 2015: *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. Oxford University Press.

——— 2016: 'Metanormative Theory and the Meaning of Deontic Modals'. In *Deontic Modality*, edited by Nate Charlow and Matthew Chrisman. Oxford University Press.

Chrisman, Matthew, and Graham Hubbs unpublished: 'The Prescriptivity of the Imperitival Form'.

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

Dreier, James (2009): 'Relativism (and Expressivism) and the Problem of Disagreement'. *Philosophical Perspectives* 23(1), pp. 79-110.

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

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

Korsgaard, Christine (2008): *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford: Oxford University Press.

Kratzer, Angelika 1981: 'The Notional Category of Modality'. In Eikmeyer & Reiser (eds.),*Words, Worlds, and Contexts*, Berlin: de Gruyter.

Kripke, Saul 1980: *Naming and Necessity*. Cambridge, Mass.: Harvard University Press.

MacFarlane, John 2014: *Assessment Sensitivity: Relative Truth and Its Applications*, Oxford: Oxford University Press.

Perez Carballo, Alejandro 2014: 'Hermeneutic Expressivism'. In Burgess, Alexis and Brett Sherman (eds.), *Metasemantics: New Essays on the Foundations of Meaning*, Oxford: Oxford University Press.

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

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

Schroeder, Mark 2008: *Being For*. New York: Oxford University Press.

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

———1974: 'Meaning as Functional Classification'. *Synthese* 27, pp. 417–37.

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

Speaks, Jeff 2010: 'Theories of Meaning.' In Zalta, E. (ed.), *Stanford Encyclopedia of Philosophy Winter 2010 edition*.

Suikkanen, Jussi 2009: 'Metaethics, Semantics, and Metasemantics'. On *Pea Soup: A blog dedicated to philosophy, ethics, and academia*, <URL=http://peasoup.typepad.com/peasoup/2009/07/metaethics-semantics-and-metasemantics.html> last accessed February 5, 2015.

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