THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

What is a theory of normative concepts for?

Citation for published version:

Chrisman, M 2019, 'What is a theory of normative concepts for?', *Philosophy*, vol. Royal Institute of Philosophy Supplement 86 , pp. 63-85. <https://doi.org/10.1017/S1358246119000079>

Digital Object Identifier (DOI):

[10.1017/S1358246119000079](https://doi.org/10.1017/S1358246119000079)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Philosophy

Publisher Rights Statement:

This article has been published in a revised form in Royal Institute of Philosophy Supplement 86 (Supplement to "Philosophy"). <https://doi.org/10.1017/S1358246119000079>. This version is free to view and download for private research and study only. Not for re-distribution, re-sale or use in derivative works. © The Royal Institute of Philosophy and the contributors 2019.

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

What Is a Theory of Normative Concepts For?

Matthew Chrisman

*Abstract: This paper compares and contrasts two recent approaches to the theory of normative concepts with each other and with more traditional theories in metaethics, in order to highlight several different projects one could be engaged in when developing a theory of normative concepts. The two accounts derive from Millgram, *The Great Endarkenment* (Oxford University Press, 2015) and Chrisman *The Meaning of 'Ought'* (Oxford University Press, 2016). These accounts share in rejecting traditional attempts to explain what 'ought' is about or expresses. Instead these accounts treat 'ought' as a quantificational word. However, the nature and range of the quantification are importantly different in the two accounts, which impacts on the ways in which the accounts integrate with the various projects one could be engaged in when developing a theory of normative concepts.*

1. Introduction

Much of 20th century metaethics was focused on one or both of the projects of:

1. Explaining what ethical and then eventually all of normative thought is about, in order to frame investigation into the nature and source of normativity (and ultimately the possibility of knowing it and being rationally moved by it)
2. Explaining the meaning of ethical and eventually all normative terms, ideally in nonnormative terms so as to 'reduce' the normative to the nonnormative or at least show how normative thought and talk fits into the (natural) world.

Nonnaturalism and expressivism form two centres of gravity in the relevant literature. Suppressing nuance, we might say that nonnaturalists (or: 'nonreductive realists') start with 1 and assume that normative concepts are – like most other concepts – about some (putatively real) properties or relations. Then they argue the properties or relations that ethical/normative concepts are about are not the sort of things whose nature can be discovered empirically or investigated scientifically. In light of this, they urge us to conclude that ethical/normative concepts (and so also terms) must be about some part of reality that is fundamentally different from (though possibly metaphysically intertwined with) the natural world. Then it is natural for nonnaturalists to go on to claim that we know what people ought to do in a *sui generis* way, and this is in part because being rationally moved by our own thoughts about what we ourselves ought to do is fundamentally different from other forms of reasoning.

By way of contrast, again suppressing nuance, we might say that expressivists (or: 'quasirealists') start with 2 and reject the assumption that ethical/normative terms are about (robustly real) properties or relations; instead they attempt to analyse the meaning of ethical/normative sentences by describing the systematic potential of these sentences to express noncognitive attitudes of the sort that have the motivational potential of preferences or plans. Then, they attempt to integrate that naturalism-friendly view of the meaning of ethical/normative terms into a broader use-theoretic account of meaning (for all terms), where semantic content is basically conceived as a function of usability to express attitudes (both cognitive and noncognitive). In light of

this, they urge us to conclude that the way ethical/normative thought rationally motivates action is much more like the way an intention or plan rationally motivates action; and having normative knowledge is more like knowing what to do than knowing a piece of reality.

There are other views and projects, of course, but I think it's fair to say that a lot of the literature in 20th century metaethics has treated projects 1 and/or 2 as foundational, with nonnaturalism and expressivism forming the main loci or at least foils of the discussion. Moreover, I think this focus has resulted in important advances in our understanding of normative concepts that should be preserved in metaethics and insisted on in other areas of philosophy. We've seen sophistication of our theories of supervenience; development of new kinds of conceptual analysis (e.g., network and a posteriori); increased care about the psychology of motivation and its connection to normative thought; better appreciation of semantic embedding constraints; the mixed nature of most language expressing some kind of concept; and the fact that specific words don't come marked as 'normative' (or even 'ethical'); enriched treatments of the categories normative, evaluative, ethical, and moral; improved theories of expression and navigation of the semantics-pragmatics divide; and better appreciation of the possibility of separating truth/belief-aptness from word-world representationality.

Nevertheless, sometimes I hear philosophers express dissatisfaction with the research program of 20th century metaethics, suggesting that it is degenerate or at least stagnate. And contemporary metaethics does sometimes seem to me to have increasingly entrenched intellectual camps (gravitating even more strongly towards nonnaturalism or expressivism) and mainly focused on developing more and more sophisticated ways to accommodate the reasons their opponents give for favouring the opposite view, or even to endorse naïve statements of the opposite view in letter rather than spirit.

This stagnation and dialectical polarization might suggest we need to come up with a radically new account of what normative thought is about or a new way of explaining the meaning of normative terms without using normative terms. However, I'm more inclined to think of 1 and 2 as two approaches to developing a theory of normative concepts and to start to consider other approaches to the theory of normative concepts. Here are two more things one might be up to in developing a theory of normative concepts:

3. Explaining what differentiates normative concepts from other concepts, while capturing what they have in common qua concepts
4. Explaining the distinctive role normative concepts have or, possibly, *could and should* have in thinking and reasoning about what to do, think, or feel.

Of course, nonnaturalism, expressivism, and their competitors have implications for projects 3 and 4, but I think it's fair to say that the core research program of 20th century metaethics was implicitly committed to the assumption that we should first sort out projects 1 and 2, and then approach projects 3 and 4 on that basis. By contrast, I want to suggest here that there's progress to be made by bracketing (and indeed possibly even abandoning) projects 1 and 2 and focusing directly on projects 3 and 4 in the pursuit of a theory of normative concepts. I also want to suggest this shift in focus might bring other related projects more clearly into view.

To be more specific, here I consider two accounts of the concept expressed by the word 'ought' which take this concept to be foundational for any full theory of normative concepts. The first account is developed by Millgram in his book *The Great Endarkenment*¹. The second account is one that can be gleaned from my book *The Meaning of 'Ought'*². So you could read this paper as a sort of critical compare and contrast. I'm going to argue the latter account is more promising (naturally!), though I'm even more interested to stress the similarities between the views, since they were developed independently from very different starting points and in quite different theoretical contexts, suggesting mild confirmation for the similarities. I am also keen to highlight how both accounts can be viewed as emerging from a relatively distinctive foregrounding of projects 3 and 4 in the theory of normative concepts, at the expense of projects 1 and 2. In this context, Millgram stresses the cognitive function of ought-thoughts, while I stress the inference-governing function of ought-claims conceived as a class of modal claims. But we both view our focus as an important break with the sort of metaphysical and semantical analyses that characterized much of 20th century metaethics.

2. Why Millgram Thinks We Need a New Account of the Ought-concept

For the purposes of this paper, I'm just taking it for granted that normative concepts are distinctive in how they are 'fraught with ought'. I know that idea means different things to different people, but what I have in mind is the idea that normative concepts are centrally caught up in our normative thinking and talking about what people ought to do, think, and feel. That idea is not, of course, an *account* of normative concepts or normativity, it's just a vague framework that encourages focusing in the first instance on OUGHT.³ In this paper, I examine Millgram's and my competing accounts of this concept as an illustration of how one might begin to develop a theory of normative concepts focused on projects 3 and 4, while bracketing and maybe even abandoning projects 1 and 2.

Millgram claims that nearly a century of semantic analysis of normative terms has failed to generate a satisfactory account of what ought-thoughts are about or an explanation of what 'ought' means in nonnormative terms – i.e. failed at projects 1 and 2. This

¹ Elijah Millgram, *The Great Endarkenment: Philosophy for an Age of Hyperspecialization* (New York: Oxford University Press, 2015).

² Matthew Chrisman, *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics* (New York: Oxford University Press).

³ Our focus is OUGHT and not GOOD, which has been the more traditional focus in metaethics. This may make it less surprising that we reject attempts to explain what normative terms are about or what noncognitive state they express, since after all 'ought' is a modal auxiliary verb rather than a predicable adjective. However, it raises the question: Is GOOD also a normative concept; and, if so, how do Millgram and my respective treatments of OUGHT extend to it? I don't know what Millgram thinks, but my view is that it is quasi-definitional of 'normative concept' that it is 'fraught with ought' and so either GOOD is not fraught with ought and so not a normative concept or its normativity can be captured in its being fraught with ought. I tend to think 'good' is usually used in a way that is merely evaluative and not normative, and a kind of contextualised relativism to a standard for evaluating is one attractive theoretical option. But I'm setting all of that off stage here in order to focus on two similar but contrasting accounts of ought, and to highlight broader meta-philosophical questions about the project of developing a theory of normative concepts.

impugns what he sees as the standard methodology of metaethics: semantic analysis. His main argument for this claim, as I read him, appeals to three key desiderata on an account of the ought-concept, conceived as foundational for any broader account of normative concepts *writ large*. He thinks none of the main accounts of normative concepts from 20th century metaethics meets these satisfactorily or even looks to be on the right track.

Millgram's first desideratum is to explain why it often seems in moral philosophy that we reach a point in our thinking where the best we can say is that some ought-claim just seems true. He writes, 'Again and again over the history of moral theory, we find [moral philosophers] adopting a way of thinking on which there are actions you *just* ought to do...not because there are arguments in favour of doing them...but as a matter of brute metaphysical fact'.⁴ Because any further justification we might attempt manifestly rests on less secure premises, we're inclined to think we know these ought-claims to be true without argument and by intuition.

Millgram's second desideratum is to account for the way 'ought' has its central use in prescribing action but also has many other uses. For example, if we're discussing Charlie's recent slight towards Sarah, we might say 'Charlie ought to apologize to Sarah,' as a way of prescribing action. But Millgram points out quite rightly that we also use 'ought' in several other ways. He mentions three, but there are others. First, gesturing at what decision would be supported by informed preference (e.g. 'Supermarket sushi ought to be avoided if possible unless it is made fresh on the premises,' is more naturally heard as a heuristic rather than a prescription). Second, articulating sovereign dictates (e.g. 'One ought not drive over 70mph in New Jersey,' can be said as a way of reporting the traffic law in a particular jurisdiction rather than prescribing action). Third, summarizing domain-specific expert know-how (e.g. 'The chain ought to be replaced when the front sprocket is replaced' could occur in a bicycle manual rather than as someone telling someone else what to do). Millgram thinks a view of ought-claims which said that such claims are always prescriptive would be wrong; but equally a view which couldn't explain how these other sorts of ought-claims are related to each other and related to prescriptive ought-claims would be inadequate.

Finally, Millgram suggests that an account of OUGHT must make sense of the usability of ought-claims as premises in practical arguments, while respecting the non-monotonicity of practical reasoning. He has in mind the apparent defensibility of reasoning from a normative premise to an action. If, for instance, I think that I ought to leave work now, and I know that taking the elevator is usually the fastest way out of the building, and on this basis I decide to push the button for the elevator, then clearly my decision to push the button can be represented as the culmination of a practical argument beginning with the premise that I ought to leave work now. Millgram stresses, however, that in this and similar cases some further premise could always come in to undercut the rational support the ought-claim gives to the practical decision. For example, if I come to realize that the elevator is out of order, then the fact that I ought to leave work now no longer supports the decision to push the button. He thinks this is generally the case: practical reasoning is non-monotonic and our theory of normative concepts should make sense of that fact.

⁴ *Op. cit.* 128.

I don't want to endorse Millgram's desiderata as the main thing we should be focused on in explaining what distinguishes normative concepts from other concepts, and I'm not sure Millgram would either. But articulating these as desiderata on a theory of OUGHT helps to explain why Millgram is so sceptical of the standard views in 20th century metaethics. In his view, the first desideratum and the general methodology of semantic analysis in 20th century analytic philosophy has motivated many smart metaethicists (via arguments such as Moore's Open Question Argument and its descendants) to suggest that normative terms refer to some irreducible normative property or relation, something we can know brutally or by intuition rather than argument. The core thought is that, if we cannot semantically decompose a normative term such as 'ought' into more primitive nonnormative terms, then we must assume that it refers to an irreducibly normative property or relation, such as the correct-to-do relation, something we can 'just know' to hold in certain intuitive cases. However, like many other metaethicists before him, Millgram regards this answer as one that peddles metaphysical and epistemological mystery.

Moreover, and more originally, he also thinks this answer obscures any path towards satisfying his second and third desiderata. If 'ought' refers to the correct-to-do relation, then it's clear why ought-claims are prescriptive, but it's not clear then why claims about this irreducibly normative relation are usable to gesture at what decision would be supported by informed preference, articulate sovereign dictates, summarize domain-specific expert know-how, and any myriad of other functions. One might suggest that 'ought' ambiguously refers to several relations. But the different ought-claims seem to have important similarities, which would not be explained by the ambiguity view. Moreover, if 'ought' refers to the correct-to-do relation, then it's also not clear why reasoning from the attribution of this nonnatural relation is generally non-monotonic. In Millgram's view, a metaphysical view about what 'ought' refers to is just not the right place to find an explanation of the related but not obviously unified set of functions the ought-concept appears to play in our thought and discourse, or the defeasible way its application functions in practical reasoning.

If one has a distaste for the metaphysical and epistemological challenges generated by nonnaturalism, a tempting response has been to argue that normative terms do not describe reality but rather express a primitive noncognitive stance towards reality. This is the expressivist response, tracing all the way back to Ayer's emotivism, which was motivated by intuitions such as those behind Moore's Open Question Argument but committed to naturalism in its ontology and epistemology. And still today, expressivists say that, in order to explain the meaning of normative terms such as 'ought', we shouldn't look for what in reality it is about but look at what one is doing when one uses it. And they claim one is expressing a distinctive kind of attitude, one with the motivational character of a preference or plan, not a representation of reality.

In Millgram's view this expressivist response might go some way towards meeting the first desideratum, but he thinks it offers limited progress on the second and third desiderata. There is no reason to think the same noncognitive stance is involved in all of the different but related uses of 'ought'; and although a noncognitive stance might, on some views, be a possible starting point of a piece of practical reasoning, it's not clear why its rational support for practical decision be non-monotonic as ought-thoughts

(and other normative thoughts) seem to be? Moreover, he doesn't discuss this, but every metaethicist also knows how expressivism as a semantics for normative terms has raised a whole host of tricky issues around the embeddability of normative sentences in unasserted contexts.

I don't want to suggest there's nothing nonnaturalists and expressivists can respond to these criticisms, but if we accept Millgram's claim that the main traditional accounts of normative concepts in 20th century metaethics have failed to meet these three key desiderata, what should we metaethicists do? Should we dig deeper into the possibility of nonnatural normative relations or toil harder on the possibility of logically embeddable noncognitive stances? There have certainly been many attempts in these directions, but Millgram thinks the root of the problem is the methodology: semantic analysis. By attempting to explain what 'ought' is about or what noncognitive stance it expresses, as part of the foundation of a theory of the meaning of normative terms, he thinks we are misled in our theoretical focus and so inevitably going to fail to meet the second and third desiderata on an account of normative concepts.

3. Millgram's Approach and the Beginnings of an Account of Normative Concepts

Instead of further semantic investigation into what in reality normative terms are about or the noncognitive stance normative terms serve (in some meaning-constituting way) to express, Millgram proposes to replace the methodology of semantic analysis that metaethics has inherited from 20th century analytic philosophy with a methodology based on a 'design characterization of an intellectual device.' His idea is to explain first what the central ought-concept does for us, what its cognitive function is, i.e. what it lets us think and say, at least in central cases. Then, he thinks we need an account of how this cognitive device could have been 'exapted' to do other things in more complex environmental, technological, and/or social circumstances – circumstances like the ones we have come to face as human cognition and reasoning has evolved and developed. Moreover, he thinks this characterization is potentially open-ended in that it should also include some indication of how the ought-concept, conceived as a malleable intellectual device could be exapted to do other things in even more complex or at least different environmental, technological, and/or social circumstances that we may come to face in the future.

Millgram thinks that, in addition to being better positioned to meet the three desiderata above, such a 'cognitive function analysis' would helpfully exhibit the work an ought-concept does and could do within a larger cognitive system and intellectual environment. Moreover, it could provide such theoretical illumination without incurring obligations to explain what 'ought' means in some reductive way that is barred from using normative concepts at the base of our account of normative concepts.

In the book cited above Millgram only barely begins to execute this project, but I think we can see enough of the beginning to understand the basic direction it might go and what is distinctive about it. In his view, 'the central (though not the only) use of a practical ought is as an existential quantifier over arguments for doing something...to say that you ought to do something is, again in this central use, to say that there is a

good argument for doing it'.⁵ What does he mean by an 'argument for doing something'? I think he just means a set of considerations that would rationalize doing it; the sorts of things one might cite in everyday life (not the philosophy classroom) as reasons which support doing it, the sorts of thoughts one might think when considering what to do and then deciding on one of one's options for reasons.

Moreover, in Millgram's view, 'Just about any practical inference form that you would have real occasion to deploy is defeasible,' meaning that those existential quantifiers expressed by 'ought' 'range over defeasible arguments'⁶. And because he thinks that the central ought-concept adverts to the existence of an argument for doing something, he thinks ought-claims can be useful for committing to a practical argument while screening off the content of that argument much like non-disquotational truth-claims (e.g. 'What James said is true') commit to the truth of a proposition while screening off the content of what is claimed to be true. One might say, e.g., 'You ought to apologize to Sarah' thereby affirming the existence of an argument for apologizing to Sarah without revealing the specific reason that supports this action. That is, Millgram claims, 'just like truth lets you affirm claims, while screening off their content, oughts (in their central use) let you recommend choice and mark them as supported, while screening off the arguments that support them'⁷.

Moreover, because Millgram thinks 'ought' in its central use adverts to the existence of a defeasible argument for action, this word can be used to prescribe action in a distinctive way. Telling someone they ought to do something is subtly different, according to Millgram, from simply commanding them to do it. As he puts this, 'oughts allow you to tell someone to do something, to tell him that you're not just pushing him around or making him conform to your arbitrary whims, but without having to tell him just what the considerations which make it more than that sort of thing are'⁸. His idea is that telling someone they ought to do something is telling them, in part, that there is a good argument for doing it, even though one isn't telling them what this argument is.

As I indicated above, this is only the beginning of a 'design characterization' of OUGHT considered as 'an intellectual device,' and even more would need to be done to extend it to all normative concepts. However, Millgram argues that we can already begin to see how it has resources to meet his three desiderata outlined above, without resting on any semantic analysis of what irreducible normative property or relation 'ought' is about or what noncognitive stance or attitude it expresses. The reason some ought-claims seem to be 'just (really!) true', in Millgram's view, is that they are the ones where we are sure there is an good argument for doing something, but – because of the way 'ought' screens off the content of the arguments it is used to say exist – we don't necessarily know or agree about what the arguments are. That's why they have a kind of foundational feel; practical justification bottoms out when anything else we could say to support a claim about what someone ought to do is less secure than the claim itself.

According to Millgram, this is not a quirky or unfortunate feature of our central ought-concept. It's a feature of the 'design' of this concept: Since almost any ought claim one

⁵ *Ibid.* 131.

⁶ *Ibid.* 132.

⁷ *Ibid.* 142-43.

⁸ *Ibid.*

might really give in support of doing something is defeasible, knowing whether any particular practical argument really supports doing it requires understanding the conditions under which the argument would be defeated, which is often not straightforward and can require a great deal of domain-specific practical wisdom. It would be impossible for us to keep track of all of the defeaters in every case, and it is much better for managing the cognitive load of practical deliberation for us to be sure that there is an argument for doing certain things, even if we're not sure what the content of this argument is. If asked for the argument why we ought to, e.g. take the means we believe to be necessary for our ends, do what a trusted adviser has advised us to do, or refrain from doing something that causes a lot more harm to oneself and others than it benefits oneself, often the best we'll be able to say is that it just seems (really!) true that this is what we ought to do.

That's how Millgram sees things working with the 'central' ought-concept, the one caught up directly in practical reasoning and prescribing action, but what about the other less central uses of 'ought'? Millgram argues that while these penumbral ought-claims make systematic semantic analysis of 'ought' in other terms impossible – a claim which I'll challenge below – by pursuing a design characterization, there's a way to see how this concept could have been exapted for use that's not directly prescriptive but involves a weaker form of gesturing at the existence of arguments for doing things. Those ought-claims that seem to advert to the verdicts of informed preference (e.g., 'Supermarket sushi ought to be avoided unless it is made fresh on the premises') are, according to Millgram, ways of saying that people in the know think there is a good (though defeasible and domain-specific) argument for doing what they say one ought to do. Ought-claims that seem to note the dictates of some sovereign rule (e.g., 'One ought not drive over 70mph in New Jersey') are, according to Millgram, a way of saying that the institutionalized wisdom is that there is a good (though defeasible) argument for behaving in a particular way. Ought-claims that seem to convey domain-specific expert know-how (e.g. 'The chain ought to be replaced when the front sprocket is replaced') are, according to Millgram, ways of saying that the relevant expert would think that there is a good (though defeasible) argument in a particular domain for doing what the claim says we ought to do.

In each case, the exapted ought-concept still serves to existentially quantify over good practical arguments for doing something, but such arguments have a less direct role in practical reasoning than the central case. In this way, Millgram finds unity in the intellectual function of the ought-concept where it looks impossible to find a unified definition or semantic analysis. Moreover, his view there is no indubitable foundation to practical reasoning, a first premise which grounds all practical inference. So he thinks it is wrongheaded to attempt to explain the foundational feel of some ought-claims in metaphysical terms, e.g., in terms of the irreducibly normative character of normative properties, or in expressivist terms, e.g. in terms of some fundamental noncognitive stance. Rather, he claims, this foundational feel is a consequence of the fact that ought-claims existentially quantify over practical arguments for doing things, but we don't always know what these arguments are. In any case, he thinks these arguments are always defeasible in any realistic cases of practical reasoning, which means that there is always the possibility of further discussion of whether any defeaters are in place. This is what makes sense of the usability of ought-claims in practical reasoning within the context of a commitment to the non-monotonicity of practical inference. If one thinks

there is a good argument for doing something, that can be a key premise in reasoning towards a decision to do it, even if one doesn't know what the argument is or what sorts of considerations would defeat it.

Returning to the four possible aims of a theory of normative concepts that I outlined above, it is clear that Millgram thinks it's pointless to try to explain the meaning of normative terms such as 'ought', whether by identifying what property or relation 'ought' is about or by characterizing the noncognitive stance this word and ones like it are used to express. So his 'design characterization' of the ought-concept conceived as an 'intellectual device' is much more focussed on the third and fourth projects to the exclusion of the first and second projects. That is, regarding the project of explaining why good forms of practical reasoning can (and often do) involve an ought-claim as a premise, Millgram provides a cognitive function analysis: making an ought-claim is a way of stopping practical justification when the chain of any piece of practical reasoning is potentially infinite. In making the ought-claim, one commits to the existence of a good (though defeasible) argument for doing what the ought-claim prescribes, without commitment to any particular view about (or even to know what) the content of this argument is.

Millgram doesn't go so far as to claim that the existential-quantifying function of OUGHT is what marks out normative concepts from nonnormative concepts, but regarding the project of explaining what normative concepts have in common which distinguishes them from nonnormative concepts, he does clearly think that OUGHT is a cognitive device that is crucial for simple forms of practical reasoning and which can be exapted for other cases deriving from increased specialization of human (practical) knowledge. And this is what he thinks distinguishes this concept and other closely connected normative concepts from other sorts of concepts.

As I said above this can be only the beginning of a bigger account of normative concepts, but we already see enough to understand how it is radically different from the sorts of nonnaturalist and expressivist approaches to explaining normative concepts based on projects 1 and 2 outlined above. Rather than pursuing an account of the distinctive piece of reality 'ought' is about or the distinctive mental state it functions to express, Millgram provides a cognitive function analysis, which begins to address project 4 of explaining the distinctive role normative concepts have in thinking and reasoning about what to do. In his view, good forms of practical reasoning can (and often do) involve an ought-claim as a premise. This, he seems to think, is the distinctive role for OUGHT in its central use.

Regarding project 3 of explaining what distinguishes normative concepts from other concepts (while capturing what they have in common qua concepts), Millgram doesn't go so far as to claim that the existential-quantifying function of the central ought-concept is *the* distinctive feature of normative concepts, but he does clearly think of OUGHT as a cognitive tool for practical reasoning in the context of pervasive defeasibility of practical reasons; and he thinks this tool can and has been adapted to other related purposes. So if OUGHT is foundational in the sense that other normative concepts are fraught with ought, we have the beginnings of a theory of normative concepts that would account for what distinguishes them as normative: they somehow function centrally to quantify existentially over arguments for doing things. I suspect that

wouldn't be Millgram's full and final account of normative concepts, but it does let us see how his cognitive function analysis of an intellectual device offers a distinctive approach to the theory of normative concepts.

4. Intensional Compositional Semantic Analysis of 'Ought'

Millgram's cognitive function analysis of OUGHT is highly original both in its approach and its results. It also raises many questions I don't know how he would answer, but rather than trying to develop some of the details on his behalf in order to consider how well the view fares against potential objections, I want to turn now to a different kind of account of OUGHT, the one which might be gleaned from my book *The Meaning of 'Ought'*. This is an account that, in some ways, is apt to seem much more traditional in its approach, but I want to explain some of the results in a way that highlight how my account shares more with Millgram's non-traditional approach than might be initially apparent.

In the book, I focus on the single term 'ought'. I apply and refine the standard treatment of English modal auxiliary verbs in lexical and compositional semantics to this term so that the standard semantic analysis of this term one sees in linguistics becomes more interactive with the concerns that motivate philosophers to ask about the meaning of 'ought'. And I think this can be seen as a way of developing foundations for a theory of normative concepts. In particular, I am interested in whether anything like the standard analysis in linguistic semantics can help us to understand why 'ought' seems distinctively prescriptive in its core normative use but also has other non-prescriptive uses, and also why ought-thoughts seem semantically and logically distinct from is-thoughts, both in their cognitive role and inferential potentials.

What we find in lexical and compositional semantics is that the English word 'ought' is regarded as an intensional operator, more specifically a (weak)⁹ necessity modal with diverse practical, evaluative, and epistemic applications. For example, each of the following are perfectly legitimate uses of 'ought':

- 'Charlie ought to apologize to Sarah.'
- 'It ought to be the case that no one exceeds 70mph in New Jersey.'
- 'Tomorrow it ought to rain.'

In linguistics, this diversity is not treated as an ambiguity but a consequence of the contextualized and intensional aspect of the semantics of modal verbs quite generally. We also see similar diversity in legitimate uses of other modal verbs such as 'may', e.g.:

- 'Charlie may leave the house.'
- 'Cars without mufflers may be driven only outside the city of Newark.'
- 'Tomorrow it may rain'

Similar diversity can be observed with 'must' and 'can' and 'should', and this variety seems to be a common feature of languages that lexically encode modality (rather than use mood markers on verbs for a similar purpose).

⁹ In the book, I provisionally propose an explanation of the relative weakness of 'ought' compared to 'must'. I won't attempt to explain this here. There are other attempts in the semantics and metaethics literature; most of these are consistent with the idea that 'ought' is some kind of necessity modal universally quantifying over some range of possibilities.

As a result, we need some general and flexible way to model the semantic contribution of modal verbs. And the contribution 'ought' makes to the semantic content of the declarative sentences in which it takes widest scope is typically modelled in semantics with a universal quantification over a contextually fixed range of possibilities. This quantification then determines a compositional articulation of the truth conditions of the sentence. Philosophers usually know this best from the possible worlds model-theoretic semantics for modal logic, made famous by Kripke; but drawing on Lewis and Kratzer, this is now also the standard view in linguistic semantics for ordinary language modal verbs such as 'ought'. In the most abstract formulation 'ought Φ ' is interpreted as saying that Φ holds across all of some range of possibilities. For example, as a crude first pass, a sentence such as 'Charlie ought to apologize to Sarah' might be analyzed as saying that Charlie apologizes to Sarah in all of the morally best of the relevant possibilities.

From the point of view of 20th century metaethics, it's natural to wonder what this semantic analysis of the word 'ought' means for the question of what ethical/normative thought is about: what putative piece of reality does 'ought', modelled as a quantifier over possibilities, refer to? From the point of view of compositional semantics, however, this looks like a misguided question. Not all words are referring terms, and we shouldn't think the word 'ought' (considered as an intensional operator) is about some putative piece of reality any more than we think the words 'all' or 'necessary' are about some putative piece of reality. These words are assigned 'denotations' or semantic values (complex functions) in standard truth-conditional semantic theories, but they are not typically assumed to refer to things in the same manner that names, definite descriptions, or ordinary predicates might be thought refer putative pieces of reality.¹⁰

Again, from the point of view of 20th century metaethics, it's natural to wonder if the idea that 'ought' is nonreferential doesn't commit us to some kind of expressivist semantics for 'ought', and if it does, what distinctive noncognitive attitude is being claimed to be expressed by ought-sentences. However, my view is resolutely truth-conditionalist in its articulation of the contribution modal words such as 'ought' make to the compositional semantic content of the sentences in which they figure. To be sure, there are interesting and difficult questions about how to interpret the idea that, e.g., 'Tomorrow it ought to rain' is true iff it rains tomorrow in all of the most likely possible futures. But that idea is manifestly different from the idea that this claim's meaning is to be explained in terms of the distinctive noncognitive attitude it expresses. Indeed, in my view, a semantic analysis of ought-sentences that treats 'ought' as a universal quantifier over possibilities is simply silent on what kind(s) of attitude are conventionally expressed by ought-sentences. Or, if you like, we might say quite blandly that ought-sentences express (inter alia) ought-beliefs whose content is usually identical to the sentence used to express them. That is, if we hold that 'ought Φ ' can be interpreted most abstractly as saying that Φ holds across all of some range of possibilities, there's a perfectly good sense in which we can accept that this interpretation also applies to the content of the belief expressed by an assertive utterance of the sentence. If, for instance, someone asserts 'It

¹⁰ Here's where it might seem important that I've focussed on 'ought' rather than 'good'. If we're understanding normative concepts as those which are 'fraught with ought', I think that's a fair starting point for developing a theory of normative concepts. But I acknowledge that it does raise questions about evaluative concepts and their similarities and differences with normative concepts. See fn. 3 above and sec. 7.3.3 of Chrisman *op. cit.*

ought to be the case that no one exceeds 70mph in New Jersey,' we might say this expresses the belief that it ought to be the case that no one exceeds 70mph in New Jersey – a belief whose content can be represented in terms of the proposition that no one exceeds 70mph in New Jersey holding in of the relevant possibilities.

5. Towards a Theory of Normative Concepts that is Interactive with 20th Century Metaethics?

So far I've described an account of the semantic value of the word 'ought' rather than an account of the concept OUGHT. But I think an account of the ought-concept can be gleaned fairly easily from it, and from that we can see the beginnings of a different theory of normative concepts. The account is one that treats ought as a modal concept, whose use like other modal concepts is not primarily for thinking about how things are in actuality but for displacing our thought from actuality to some range of possibilities in order to think about what holds in all, most, or some of these. Thinking Charlie ought to apologize to Sarah is clearly different from thinking that Charlie is apologizing to Sarah; and, as we all know, we cannot assume that Charlie will do as we think he ought to do.

I suggested above that my approach is apt to seem much more traditional than Millgram's is because I pursue the foundations of a theory of normative concepts by engaging in semantic analysis, and semantic analysis is the methodology Millgram encourages us to eschew in thinking about normative concepts. However, as I think should now be clear, I don't engage in semantic analysis in the way characteristic of much of 20th century metaethics. That is, I don't see semantics for 'ought' as an attempt to explain what distinctive kind of (putative) thing in the world ethical/normative concepts are about or what distinctive kind of mental state ethical/normative terms function to express. Moreover, I don't treat ought as primarily a conceptual device for thinking prescriptively about what actions oneself or others are to perform, rather I see it as primarily a modal term for reasoning about nonactual possibilities.

That being said, I think there's a way to bring the standard intensionalist treatment of the semantics of the word 'ought' standard in linguistics into more direct contact with some of the concerns of traditional metaethics. In the book, I make two key suggestions in this regard. First, I suggest that the standard compositional semantic treatment of 'ought' as a weak necessity modal can be usefully refined and enriched to make sense of the difference between thoughts about what propositions ought to be true and what people ought to do, think, and feel. Second, I suggest that there's a way to reconceptualize traditional metaethical debate between nonnaturalists and expressivists as two different 'metasemantic' interpretations of the compositional semantic treatment of 'ought' that is standard in linguistic semantics; and I argue that once we do so, a third metasemantic interpretation comes into view and looks at least as plausible as those two.

How can we account for ought-to-do thoughts? Standard intensionalist semantics for modal words was designed primarily for alethic modality, that is, thought and talk about what is necessarily/possibly true or false. As a result, formal semanticists tend to assume that modal words are propositional operators used to articulate thoughts about what propositions are true not in the actual world but in various possible worlds. There

is a straightforward way to apply this to claims about what someone ought to *do*: we treat 'ought' as an operator on the proposition that the relevant agent does the relevant thing, and we analyse the sentence in terms of this proposition being true across all of some range of possible worlds (not including the actual world), e.g. the morally best worlds.

That might be the right view, but I have come to suspect that there is an important conceptual difference between thinking that a proposition ought to be true and thinking that a particular agent ought to do something, at least in the prescriptive way that is distinctive of normative thought and talk. It would take me too far afield to motivate that suspicion here, but the basic idea is the thought that someone's actively doing what they ought to do is only one of many ways to make true the proposition that they have done what they ought to do. Hence, thinking or saying that, e.g., Charlie ought to apologize to Sarah (in the prescriptive way that is characteristic of genuinely normative discourse) might be obscured by standard intensionalist treatment of the semantics of 'ought'.

If that is right (and, again, I'm not arguing here that it is), we might worry that it threatens the usability of the intensional semantics for 'ought' as the foundations for a theory of normative thought. In the book, however, I argue this worry can be met. For one can generalize from the idea that modal words are propositional operators quantifying over possible worlds where they hold true, allowing that they could operate on other kinds of content quantifying over other kinds of possibilities where they hold. More specifically, I develop an account of 'ought' that allows for different kinds of objects ('propositions' and 'practitions' or, roughly, declarative and imperative content) and in either case says that the relevant object 'holds' across a range of possibilities. In the latter case, I suggest that imperative content can be said to hold *legitimate* across a range of possible norms much like declarative content is usually said to hold true across a range of possible worlds. (The details are a bit more complex than this, as imperatival legitimacy is usually dependent on declarative truth in a way that declarative truth is not usually dependent on imperatival legitimacy.) And this allows me to generate an account of a distinctively prescriptive ought-concept as one whose application is correct just in case some imperative is legitimate across all of a contextually determined range of possible norms (given the truth of various propositions).¹¹

This isn't the place to get into the details of the view. My point in alluding to it is just to point out that I think there is a way of incorporating some kind of distinctive prescriptivity into an account of the ought-concept based on the idea that it is fundamentally a modal concept rather than a prescriptive concept. OUGHT is a modal concept, and as such functions like other modal concepts to displace our thinking from actuality to various possibilities. This idea helps to make sense of the diversity of uses of 'ought' without positing ambiguity, but it is consistent with the idea that one distinctive use of 'ought' (perhaps the central 'normative' use) is prescriptive in a way that other uses are not.

What about projects 1-2 mentioned above; should we simply reject traditional metaethical debate about what distinctive (putative) piece of reality normative

¹¹ See Chrisman *op. cit.* sec. 5.3-4.

concepts are about or what distinctive attitudes they express? In the book, I argue that the standard intensionalist analysis of the content of ought-sentences and beliefs allows for a variety of both representationalist and nonrepresentationalist interpretations. That is, we can view the belief expressed by an ought-sentence as getting this semantic content from the way it represents reality as being, or we can view the belief as getting this semantic content in some other way. And I think there is a way of rehabilitating some nonnaturalist and expressivist ideas not as a semantics for normative terms but as metasemantic accounts of how normative sentences/beliefs get their content.¹² More specifically, I envisage a form of nonnaturalism which says that ought-sentences get the content, which is modelled by the standard intensionalist semantics in terms of a universal quantification over possibilities, by virtue of the fact that they represent nonnatural features of reality (roughly, normative relations between possible worlds). But equally, in an expressivist vein, one could argue that ought-sentences get the content they have (roughly, that something holds across all of some contextually determined range of possibilities) by virtue of expressing some noncognitive attitude with distinctive motivational potential rather than a representational belief about the way reality is.

In the book, however, I argue that neither of these rehabilitated 'metasemantic' forms of nonnaturalism or expressivism are preferable to a version of inferentialism about modal thought and talk. Inferentialism about modal thought is the view that these thoughts (and the sentences expressing them) have the content they have by virtue of their cognitive role in affirming the acceptance of inference rules, rather than by virtue of a role in directly representing the modal way reality is or in expressing an attitude with distinctive motivational potential. So, in the end, I defend a non-representationalist but non-expressivist metasemantic account of that in virtue of which ought-sentences get the truth conditional compositional content standard linguistic semantics predicts that they have. This is a view that locates the most fundamental explanation of meaning in the inferentially articulable commitments one undertakes in affirming or accepting a sentence rather than in the attitude the sentence expresses or in the piece of reality the sentence represents. More specifically, ought is treated as distinctive for how it allows to think about and articulate the inferential commitments we accept between other thoughts. In this regard OUGHT is similar to other modal concepts, but I also allow that there may be a special function in some distinctively normative uses of ought where it contributes the content that it does by virtues of how it allows to think about and articulate the inferential commitments between practical thoughts.

6. Conclusion: Compare and Contrast

Regarding projects 3 and 4, Millgram's view and my view have striking similarities. When it comes to explaining the distinctive role normative concepts do, could, and should have in thinking and reasoning about what to do, think, or feel, both of us focus

¹² There are several connected metasemantic projects. The project I am mainly interested focuses on a question of grounding: what is it in virtue of which a particular piece of language would have the semantic content that our semantic theory predicts it to have. But that can be cast as an account of what we are up to in developing theories that predict semantic contents, and it is connected to the philological-*cum*-genealogical project of accounting for how various pieces of language came to have the meanings that they have and the cognitive-functionalist project of explaining what various words and the concepts they express do for us in our speaking and thinking.

on the second-order role of ought-thoughts in reasoning, especially (in 'core' or 'central' cases) about what to do. Normative ought-claims don't, on our respective views, provide first-order considerations in favor of doing something, rather they function as ways of committing to or acknowledging something like reasons or inferential support for doing something.

For both of us, OUGHT is a concept, like other concepts, in virtue of its role in thinking thoughts; we both think of OUGHT as a constituent of thought. However, the second-order function of OUGHT in reasoning about what to do marks this concept off from other concepts. It is not immediately clear how to extend that to all normative concepts to generate an explanation of what differentiates them qua normative concepts from other concepts. If, however, there is any promise in the idea that normative concepts are normative in virtue of being 'fraught with ought' Millgram and I at least have the beginning of an answer to project 3 in our similar accounts of the second-order and quantificational role of ought-thoughts.

Importantly, however, that's only part of the story about OUGHT for both of us, as we both think this concept has other related functions. For Millgram, these are other cases where someone might need to reason by appeal to the existence of an argument for doing something without articulating or even knowing precisely what that argument is. For me, these are other cases where someone might need to think in a displaced way, characteristically facilitated by intensional operators, not about actuality but about what holds in all of a range of possibilities. I think this includes cases where one's reasoning appeals to the existence of arguments for doing something without actually concluding in a prescriptive thought about what someone in particular is to do (e.g. Millgram's examples of gestures at informed preference, statements of sovereign dictates, and summaries of expert know-how), but it also includes probabilistic reasoning about what to think (e.g. about the likelihood of rain tomorrow).

Moreover, both Millgram and I recognize a flexible function for OUGHT whose exaptability into other contexts where it could prove useful for people to be able to think second-order thoughts about something like reasons or inferential support for doing (thinking or feeling) something. So even if our accounts don't capture everything the ought-concept could and should do for us as thinkers, there is flexibility within our accounts to see how the concept might develop or evolve as humans face new circumstances.

So that is how our accounts are similar in their response to projects 3 and 4. One more point of similarity is negative, regarding project 1. Because we both think 'ought' plays an essentially quantifying role in normative thought and talk, we both reject the referentialist presumption of the traditional project in metaethics to explain what normative thought is about, in order to frame investigation into the nature and source of normativity (and ultimately the possibility of knowing it and being rationally moved by it). Both of us think it's wrongheaded to think we could explain normative concepts by identifying what in reality they are about.

There are, however, two related differences between Millgram's view about 'ought' and my view about 'ought'. The first is that he treats 'ought' as an existential quantifier over arguments for doing something, whereas I follow standard intensional semantics in

treating it as a universal quantifier over possibilities. My motivation for this is mainly that analysing 'ought' in terms of a universal quantification over possibilities (rather than in terms of an existential quantification over arguments) allows us to align normative uses of 'ought' with deontic notions such as 'obligatory' and contrast it with normative uses of 'may' which can then be aligned with 'permissible'. It is standard in deontic logic to analyse the former in terms of universal quantification and the latter as existential quantification.

As natural language semanticists took up ideas from modal and deontic logic in the attempt of developing a compositional semantic model for natural language modals, they followed this idea treating 'ought' as a (weak) necessity modal. They do this in large part because it provides a way to unify the various 'flavours' of 'ought' (and other modals), especially the use of 'ought' as an epistemic and normative modal. When faced with sentences such as 'It ought to rain tomorrow,' it is tempting to think this 'ought' is a homophone of the use in 'Charlie ought to apologize to Sarah.' However, this dual use of 'ought' is manifested in other modal words such as 'must' and 'may', and it is a feature of translations of 'ought' in many other languages. This suggests we're not dealing with homophones or even some other sort of ambiguity; and this is precisely what the standard treatment of 'ought' as a (weak) necessity modal is designed to make sense of. This kind of systematicity is – I think – an advantage of my account over Millgram's.

It means that the account of the meaning of 'ought' I adapt from intensional semantics is as compositional as standard intensionalist semantics, and it fits nicely with a broader account of modal words as intensional operators. Moreover, this can be seen as deriving from the cognitive function of all of these words in engaging in a kind of 'displacing' thought, characteristic of the use of intensional operators, whereby one thinks some range of possibilities rather than actuality.

Millgram argues that the variety of uses of 'ought' shows that 'we better not think we know what 'ought' *means*'¹³. At best, he seems to think, we could generate an open-ended disjunction for the meaning of 'ought', paraphrasing each of the different functions currently in use. This is why he thinks 'our model shouldn't be telling us what an ought *means*...we're would do better to think of it as a preliminary characterization of an intellectual *device*, one which may be much more open-ended than an attempt at a definition would tolerate'¹⁴. To my mind, however, this line of thought manifests a misunderstanding of the point of a compositional analysis of the semantic contribution a word such as 'ought' makes to the sentences in which it figures. I agree that we cannot explain what 'ought' means by giving a definition of it. (If you look it up in a dictionary, it's notable that its meaning is usually explained by reference to its linguistic function, and no explicit definition is given.) However, that's true of a lot of words, e.g. 'the', 'is', and 'all'; and yet we don't conclude that we had better not think we know what these words mean. We surely do know what they mean (tacitly), and compositional semantics seeks to articulate this meaning explicitly. What this branch of linguistics doesn't attempt to do is to say what in reality these words are *about*. That's because they aren't referring expressions – not because we don't have any idea what they mean.

¹³ Millgram *op. cit.* 139.

¹⁴ *Ibid.* 141.

So when it comes to project 2 of explaining the meaning of normative terms (ideally in nonnormative terms so as to 'reduce' the normative to the nonnormative or at least show how normative thought and talk fits into the natural world), Millgram and I reach importantly different conclusions. We are both sceptical of the reductionist ideal, but I don't reject the ambition of explaining the meaning of 'ought'. Semantic analysis conceived as, in part, explaining the semantic contribution of a word to the sentences in which it figures so as to model the implicit know-how of competent users of the language seems to me to be a perfectly worthwhile project. Moreover, it's not just something worthwhile to do. I think it can provide important insights into the cognitive function of the concepts expressed by the words we semantically analyse. From compositional semantics, we learn that modal words play a function of displacing our consideration from actuality and quantifying over possibilities; and that is what I think can be used as the basis for an account of the cognitive function of the concepts they express, including the ought-concept.

To be clear, I think Millgram's account of the cognitive function of OUGHT is fascinating and possibly foundational for a highly original approach to the theory of normative concepts. In this paper, I have sought to explain why I see the account of the meaning of 'ought' I developed in *The Meaning of 'Ought'* as animated by a similar methodological orientation to the traditional projects in metaethics, even though that might not be immediately apparent. I think there is important progress that can be made by focussing on explaining what differentiates normative concepts from other concepts and explaining the distinctive role normative concepts have in thinking and reasoning about what to do, rather than explaining what normative terms are about or what kind of attitude they express. I also think, however, that this alternative orientation is consistent with projects in linguistic semantics to explain the contribution words such as 'ought' make to the meaning of the sentences in which they figure. Indeed, advances in the semantics for our normative terms offer various insights for philosophical reflection on the way the concepts expressed by these terms function.

So far, I have suppressed any discussion of what cognitive function OUGHT could and should perform. I think this is the next step, and it brings attempts in metaethics to account for what 'ought' does now mean into contact with projects in foundational normative theory. Millgram's idea of a core cognitive function which can be exapted for other purposes in evolving ecological and social circumstances seems right to me. And I suspect further backwards-looking theorising about the historical precedents of our normative thinking (e.g. genealogical and evolutionary) can combine fruitfully with forward-looking theorising (e.g. conceptual engineering and architecture) to improve on our thinking about what to do, think, and feel. Perhaps a theory of normative concepts framed around the questions of what in reality they are about or what distinctive kind of noncognitive attitude they express can aid in this improvement. But I've found it fruitful to try to move away from those methodological assumptions in projects 1 and 2 by focussing, like Millgram, on projects 3 and 4 instead.