'Ought'-Contextualism Beyond the Parochial

Alex Worsnip

Forthcoming in *Philosophical Studies*
Penultimate version; please cite published version where possible

One of the major developments in the metaethical literature of the last decade or so is the increasing popularity of, and attention to, a view about the meaning of 'ought' known as "contextualism".[1] The basic feature of a view that makes it contextualist is that it claims that the semantic content of 'ought' varies depending on the value of one or more parameters that are determined by the context in which it is uttered. On its own, this is a rather weak and unspecific claim. It says nothing about *how* the semantic content of 'ought' varies from context to context, or what the relevant parameters are, or how they are determined. As I'll argue in a moment, construed in this generic and unspecific way, contextualism – "generic contextualism", let's call it – should not really be all that controversial.

Yet – and this has been the source of significant confusion within discussion of contextualism – the term 'contextualism' is often associated in the metaethical with a much more specific *kind* of contextualist view. This more specific view is characterized by two claims. The first claim concerns what one of the relevant parameters *is*: it is a parameter for a set of normative *standards*, or similar. The second, and particularly distinctive, claim concerns how the value of the standards parameter is determined by context: the relevant standards in any context of utterance will be those that the speaker, or others in the speaker's environment, *actually subscribe to*. Thus, simplifying, 'A ought to Φ' means something like 'by standards S, A ought to Φ', where S are the standards of the speaker, or others in the speaker's environment. And so this utterance will be true so long as the standards of the speaker, or others in the speaker's environment, require Φ-ing.

Call this view "parochial contextualism", since it makes the semantic content of 'ought'-claims dependent on the local or *parochial* standards of the speaker (or those in her environment). Unlike generic contextualism, parochial contextualism should be, and is, controversial among metaethicists. According to parochial contextualism, normative claims can be made true simply by a speaker (or others in her environment) subscribing to a set of standards according to which it is true. Thus, parochial contextualism makes the truth of normative utterances radically mind-dependent. As I'll suggest later, while parochial contextualism does not, strictly speaking, entail metaethical anti-realism, it sits unnaturally with realism. Parochial contextualism is naturally thought of as a way of combining anti-realism about the metaphysics of the normative with a cognitivist, non-error-theoretic view of normative thought and talk, according to which moral utterances express beliefs that can be straightforwardly (but mind-dependently) true or false. No wonder parochial contextualism – and

For helpful discussions and/or comments related to this paper, I'm grateful to Janice Dowell, Steve Finlay, Daniel Fogal, Chris Howard, Josh Knobe, John Pittard, Emmanuel Viebahn, Ralph Wedgwood, and three anonymous referees. I'm especially grateful to Daniel Wodak for very helpful written comments on a previous draft.

[1] See Wedgwood (2006, 2007: ch. 5, 2016); Brogaard (2008); Björnsson & Finlay (2010); Finlay (2014); Dowell (2012, 2013); Chrisman (2015); Silk (2017); Khoo & Knobe (2018). Contextualism has for some time been the dominant view of 'ought', and modals more generally, within linguistics – largely due to the influence of Kratzer (1981, 1991, 2012). For earlier forerunners of contextualism in metaethics see Harman (1975, 1996) and Dreier (1990).

contextualism more generally, to the extent that it's associated with parochial contextualism – is regarded with suspicion and hostility by many metaethical realists.

Yet we should expect there to be non-parochial forms of contextualism available. On one view, the relevant normative standards are just whichever particular set of objective, mind-independently true normative standards are conversationally salient. I'll call this view "aspirational" contextualism, since it says that normative utterances *aspire* to objectivity, rather than merely attempting to make claims about what's required by the local, parochial standards. As I'll argue, aspirational contextualism is still a form of *contextualism* in a meaningful sense. While the majority of self-described contextualists are parochialists, a non-trivial minority are aspirationalists[2] – though the distinction between the views is rarely, if ever, explicitly drawn.

This paper has two purposes. The first, which roughly occupies sections 1-2c, is to get clearer on the logical terrain around contextualism and, in particular, the differences between parochial and aspirational contextualism. The second, which roughly occupies sections 2d-3, is to introduce (and endorse) a new view, which I call "ecumenical" contextualism, and to explore its relationship to metaethical theory. This view is a flexible compromise between parochial and aspirational contextualism, according to which *some* normative utterances are parochial, and *others* aspirational – and whether a normative utterance is parochial and aspirational is itself determined by context (in particular, speaker intention).

1. Generic contextualism, and why we should all accept it

As I said above, generic contextualism is the view that the semantic content of 'ought' varies depending on the value of one or more parameters that are determined by the context in which it is uttered. 'Semantic content' is a term of art here. One might substitute 'meaning' for it, but this is potentially misleading. Following Kaplan (1989), we can distinguish two different notions in the neighborhood of meaning, which come apart for context-sensitive terms: character and content. The character of a context-sensitive term remains fixed across contexts, whereas the content varies. The character gives a kind of recipe for *how* the content varies across contexts, or to speak slightly more technically, is a function from context to content.

An example will help here. Take the indexical 'here'. The *character* of 'here' is, very roughly, this: 'here', as uttered by a speaker in location L, refers to L.[3] The content of 'here', by contrast, will be the concrete value of L that is instantiated on a particular occasion of utterance. So, for example, suppose that Sheldon is in Honolulu, and says 'it's hot here', while Ivan is in Moscow, and says 'it's cold here'. The *character* of 'here' as uttered by Sheldon and Ivan is the same: in both of their utterances, 'here' refers to the location in which they speak at the time of utterance. But the *content* of 'here' as uttered by Sheldon and Ivan differs: as used by Sheldon, 'here' refers to Honolulu, whereas as used by Ivan, 'here' refers to Moscow. We can thus say that 'here' is a context-sensitive term, for its semantic

[2] See fns. 23 & 24, respectively, for references.

[3] It's more complicated, in reality, since there are "deictic" uses of 'here', where it refers to a location other than the one in which the speaker is located at the time of utterance, for example, when the speaker is pointing at a map.

content varies depending upon the value of a contextual parameter (L) that is supplied by conversational context.

This example is helpful in illustrating a further point. The fact that 'here' shares the same *character* both as it is used by Sheldon and as it is used by Ivan amounts to an important sense in which the meaning of 'here' is the same in both contexts of utterance. This distinguishes a contextualist theory of the term 'here' from an *ambiguity* theory, according to which 'here' is simply ambiguous between many different meanings: 'here' sometimes means 'in Honolulu', sometimes means 'in Moscow', sometimes 'in New York', and so on, with no unified character to explain how these different meanings get selected in different contexts. In the case of 'here', the latter view is manifestly implausible. 'Here' is not just *ambiguous* between all the different places in the world: it has a single character across contexts, and this character plus the value of the parameter (L) that features in it explains systematically the differences in its content across contexts.

More generally, it is a sound methodological principle that if you find yourself positing ambiguity between a huge number of different potential meanings of a term, you should look for a potential way to unify these meanings under a single (or, at least, fewer) character(s), and to switch from an ambiguity view to a contextualist one. Now, it is worth noting here that linguists standardly distinguish two kinds of ambiguity: homonymy and polysemy. Homonymy occurs when the same string of letters or symbols can have completely unrelated meanings – for example, the difference between 'bat' as in the animal, and 'bat' as in a piece of sports equipment. Polysemy occurs when a word can have different meanings that are clearly analogically or structurally related – for example, the difference between 'batted' as it occurs in 'she batted 2-for-4 with a home run' and 'batted' as it occurs in 'she batted away a fly'. Polysemy is a widespread phenomenon that it is not semantically implausible to posit quite extensively. That said, it is still implausible to say that a *single* term is polysemous between *innumerable* different meanings, as with 'here'.

Given the definition of contextualism above, a contextualist could (though need not) accept that 'ought' is polysemous in certain respects.[4] Perhaps it will prove impossible to unify *all* the usages of 'ought' under a single character, and some polysemy will remain. What is distinctive of contextualism, though, is its claim that differences in the semantic content of 'ought' are not *solely* a result of polysemy (or ambiguity more generally). In other words, there is at least one character of 'ought' that itself allows for further semantic variability in content, depending on the value of one or more parameters that are determined by context. If that's so, then 'ought' exhibits context-sensitivity.

Generic contextualism is a very weak claim. For, in and of itself, it says nothing about *how*, or indeed *how much*, the semantic content of 'ought' varies according to context. If one allows that the semantic content of 'ought' varies according to context to any extent, in a way that is not a result of ambiguity, then one is a generic contextualist. And almost everyone *does* accept that the semantic content of 'ought' varies according to context, to a minimal extent. For consider utterances like

You ought to pass the bread basket only to the right.

[4] Viebahn & Vetter (2016) argue that at least some modals are both polysemous and context-sensitive.

Most of us, I think, would accept that the semantic content of 'ought' as it occurs in this sentence can differ according to context. In most contexts, the sentence is naturally interpreted as expressing a claim about what you ought to do according to the norms of etiquette (or, perhaps more specifically, according to the norms of traditional British etiquette); here it is the "'ought' of etiquette" that is in play. But there could be contexts where it is clear that the sentence is being used to make a (no doubt false) claim about what you morally ought to do; here it is the "moral 'ought'" that is in play. That is already to accept that the semantic content of 'ought' can differ across contexts.

Now, on its own, this doesn't show that 'ought' is context-sensitive; the variability could be due to ambiguity (specifically, polysemy). However, if we want to posit an ambiguity of 'ought' between the moral 'ought' and the 'ought' of etiquette, we cannot just stop there. There are also usages of 'ought' distinctively connected with many other bodies of norms: epistemic normativity, the law, self-interest, instrumental rationality, aesthetic norms, and so on. Moreover, there are also plausibly usages of 'ought' connected with innumerable *particular* cultural practices and systems of norms: the 'ought' of *British* etiquette, the 'ought' of *French* etiquette, the 'ought' of *Japanese* etiquette; the 'ought' of *American* law, the 'ought' of *Ancient Roman* law; the 'ought' of Mafia morality; and so on. Ultimately, these different 'ought' proliferate in a way so extensive as to make a pure ambiguity theory (even a polysemy version) implausible in the same sort of way that it was for 'here'. It's more attractive to accept a theory that unifies different usages of 'ought' under a smaller number of characters, explaining much of the variability as due to context-sensitivity rather than polysemy. And that gets us to generic contextualism.

Let's pause to consider two objections. A first objection might be as follows: not all of these purported usages of 'ought' are robustly *normative*. That is, many of these "systems of norms" do not have genuine normative authority over our actions (at least, not in and of themselves). This is virtually undeniable of Mafia morality; it is also plausible for systems of etiquette, and quite plausible for bodies of positive law; indeed, for each putative 'ought' mentioned above (morality, self-interest, instrumental rationality, epistemic norms, aesthetic norms, etc), it has been claimed by *some* philosopher or other than the "source" of the 'ought' in question lacks genuine normative authority. So perhaps there are only a handful of genuine normative 'ought's – few enough to make a polysemy theory manageable.

I'm open to the possibility that there are only a handful of genuine sources of normativity, but this doesn't, in fact, help the ambiguity view. To start with, we should be careful to separate the question of whether (e.g.) the law genuinely, in and of itself, has genuine normative authority – whether it is a genuine "source of normativity" – from whether there is a robustly normative *usage* of the legal 'ought'. The former requires the law to *actually* possess normative authority, whereas the latter only requires there to be speakers who *take* the law to possess normative authority.[5] So even if only a handful of the above 'ought's reflect a genuine source of normativity, many more of them might nevertheless be robustly normative *usages* of 'ought'.

Moreover, even if there were only a handful of robustly normative *usages* of 'ought', the other usages are still usages of 'ought'. Moreover, they are usages of 'ought' that are in a broad sense deontic

[5] Compare error theories about morality (Mackie 1977; Joyce 2001), which are distinctive precisely by combining the view that there are no categorical moral truths with the view that ordinary moral judgments unavoidably presuppose that there are such truths.

– connected with systems of norms or requirements – they are not, for example, epistemic (that is, "expectational") usages of 'ought'. As long as 'ought' can refer to (for example) what one ought to do according to traditional British etiquette, this is still something that a semantic theory of 'ought' has to accommodate, whether or not this usage deserves to be called 'robustly normative'. So the objection at hand doesn't mitigate the need to appeal to contextualism. It does, however, set a desideratum for particular contextualist theories, which is to adequately explain in what way some deontic usages of 'ought' *are* robustly normative and others are not. I'll come back to this in part 3c.

The second objection is this: I've implied that contextualism and the ambiguity view are the only options, but aren't there other theories, for example relativist treatments of 'ought', that compete with the contextualist view? The answer is that there are indeed relativist treatments of 'ought', but such relativist treatments are not in fact incompatible with a (certain degree of) contextualism about 'ought'. The dispute between relativists and what we might call "thoroughgoing" contextualists (see below) concerns not whether generic contextualism is true, but rather the extent to which 'ought' is context-sensitive, and to what degree the context-sensitivity of 'ought' can be used to explain certain puzzling phenomena. Relativists do not dispute that context of utterance influences the semantic content of 'ought' in certain basic ways; for example, by determining whether it's the moral 'ought' or the 'ought' of etiquette that's in play.[6] Rather, they deny that contextualism can be leveraged to explain other phenomena, such as the apparent information-sensitivity of 'ought'-judgments, holding that such phenomena can only be explained by the additional, distinctly relativist claim that the *truth* of 'ought'-judgments is in some respects relative to a circumstance of evaluation or assessment. For all that, relativists accept generic contextualism.

Though I think the conclusion of this section – that we should all accept generic contextualism – is worth being aware of, I'm not claiming it as some hugely significant result for metaethics. The very feature of this claim that makes it so weak and easy to accept – that it says nothing about *how* 'ought' is context-sensitive – also means that it is not really, in and of itself, a full-fledged *view* about the semantics of 'ought'; it's at most a *kind* of view. Moreover, generic contextualism is compatible with a range of views about the *degree* of context-sensitivity of 'ought', and the degree of contextualism that I've claimed that *everyone* is committed to is relatively minimal. What I've argued is that everyone should acknowledge that context contributes to the semantic content of 'ought' by selecting what kind of normativity is at play – moral, prudential, legal, etiquette, etc. That's a fairly limited way for context to contribute to the semantic content of 'ought', and is compatible with there being no contextual variability *within* such broad normative categories – so that, for example, there's only one possible semantic content for the *moral* 'ought'.[7]

Let's call a view on which context *only* contributes to the semantic content of 'ought' in this way "non-thoroughgoing" contextualism. By contrast, *thoroughgoing* contextualists hold that there is

[6] Prominent relativists Kolodny & MacFarlane (2010: 131), for example, make it a feature of their semantics that the "deontic selection function", i.e. the relevant set of norms that select the deontically ideal possible worlds given an information-state, is "generally supplied by context".

[7] Silk (2017: 210) thinks that this kind of view shouldn't count as contextualist at all. Obviously, this is a terminological dispute, but I prefer my terminology. Even non-thoroughgoing contextualism still contrasts with an ambiguity view about the different "flavors" of 'ought'. Moreover, my way of talking preserves the simple rule that we should call a view of a particular term 'contextualist' if it posits context-sensitivity with respect to that term. Later in his paper (Silk 2017: 235-6), Silk appears to slip into my way of talking.

contextual variability in the semantic content of 'ought' *within* such broad normative categories: there are different, contextually determined semantic values of the moral 'ought', of the prudential 'ought', of the legal 'ought', and so on.[8, 9] From here on, I'll set non-thoroughgoing contextualism aside. My aim is not to argue for throughgoing contextualism over other views but to explore different ways of pursuing thoroughgoing contextualism.

2. Forms of thoroughgoing contextualism

In considering forms of thoroughgoing contextualism, it helps to introduce the basic semantic framework that most contextualists take their cue from, due to Kratzer (1981). In this framework, there are two parameters: a *modal base* and an *ordering source*. The modal base consists of a body of propositions that are held fixed in the context. It might contain information about the circumstances the agent finds herself in, about how things would turn out given various possible courses of action on the agent's part, and (potentially) about what the agent actually is going to do in the future. The modal base then determines a set of worlds, namely the worlds that are consistent with the modal base; those in which all the propositions in the modal base are true. One could call these the "live" worlds, since they are the worlds that are possible (as opposed to ruled out) given what is held fixed. The ordering source consists of a set of standards, norms or expectations, which can be satisfied or not in each of the worlds. It orders, or ranks, the worlds by how well they satisfy these standards. So we arrive at a ranked set of worlds. On the Kratzer semantics, "S ought to Φ" is true iff all the top-ranked worlds are ones in which S Φ's. In other words, given what is being held fixed, the only way for S to satisfy the relevant norms (to the greatest degree possible) is to Φ. "S may Φ" is true iff *some* of the top-ranked worlds are ones in which S Φ's. In other words, given what is being held fixed, S *can* satisfy the relevant norms consistently with Φ-ing.

Most contextualists hold either some form of Kratzer's view, or a similar view. Kratzer's view still leaves a lot open, primarily because it leaves open how context determines the values of the two parameters. Considering different ways in which it might do so will allow us to consider different contextualist views. To fix ideas, I'll assume the basic Kratzerian framework as common ground, but the distinctions between views I'll describe could survive migration to other frameworks.

With the Kratzerian framework explained, I need to clarify a few points about terminology that are liable to confuse the debate. Sometimes I will talk of 'information-sensitivity' and 'standards-sensitivity'. It would be natural to assume that 'information' corresponds to Kratzer's modal base parameter (which is, in a good sense, filled by a body of information) and that 'standards' correspond to Kratzer's ordering source parameter (which is, in a good sense, filled by a set of standards). However, this actually isn't right, for changes in information can take effect not just on the modal

[8] Some contextualist views might deny the reality or significance of these categories entirely, holding that there are simply many different potential 'ought's, for many possible sets of norms or standards, and that it is unnecessary or unhelpful for a semantic theory to try to group them into 'moral', 'prudential', 'aesthetic', etc usages. I count these views as thoroughgoing versions of contextualism. If the categories mentioned are not real or significant, that's a problem for non-throughgoing contextualism, since it's that view that has to rely on such categories to keep the extent of its contextualism in check.

[9] It's worth noting that only a thoroughgoing form of 'ought'-contextualism deserves the name meta*ethical* contextualism, since a non-thoroughgoing form of contextualism holds that there is only one semantic value of the *moral* 'ought'.

base but on the ordering source.[10] Specifically, consider an 'ought' connected to what would make things go (expectedly) best relative to some body of information I1, and an 'ought' connected to what would make things go (expectedly) best relative to some body of information I2. For the former, the worlds will be ordered by how well the agent's actions maximize expected utility given information I1; for the latter, they will be ordered by how well the agent's actions maximize expected utility given information I2. These orderings can differ based on differences between I1 and I2. So the ordering source parameter can vary with changes in information.

Consequently, I will reserve the term 'standards-sensitivity' for changes to the ordering source that are *not* purely a result of changes in salient information. So, if the worlds were *always* ordered by what maximized expected utility (relative to the salient set of information), there would be no standards-sensitivity in my sense.[11] Standards-sensitivity would only enter if the worlds are sometimes ordered by something else entirely (for example, whether the agent's actions conform to rigid set of deontological norms). This, logically stronger, way of construing 'standards-sensitivity' ensures that information-sensitivity and standards-sensitivity are kept distinct.[12] But there could be other permissible ways of talking; the important thing is just that we are clear about what we mean by our terms.

Finally: when I talk about 'information-sensitivity', I am just referring to the phenomenon of 'ought' (and other modals) taking different semantic values as which body of information is relevant shifts. This contrasts with Dowell's (2013: 158) terminology, on which *some* uses of 'ought' are "information-sensitive", and other uses of 'ought' (viz. "objective" – or as I prefer to call them, *fact-relative* – usages relativized to the totality of the facts, epistemically unconstrained) are "information-insensitive". I don't find that a perspicuous way of talking; the totality of the facts is still, in a perfectly good sense, a body of information, even if it is one that is not actually possessed by any particular individual or group. On my way of talking, if the semantic value of 'ought' is always (partially) determined by which body of information is relevant, then *all* usages of 'ought' are thereby information-sensitive. The fact-relative 'ought' is no exception: it is just the value of 'ought' where the relevant body of information is the totality of the facts. The difference between Dowell and myself on this point is purely terminological, but is liable to confuse if not marked.[13]

(a) A very simple view

[10] This is clear in Dowell (2012, 2013).

[11] Silk (2017: 209-10) appears to build standards-sensitivity specifically into his definition of 'contextualism'. This is surely a mistake (even granting Silk's exclusion of non-thoroughgoing views from counting as contextualist; cf. fn. 7 above). Whether a view of a term is contextualist is a matter of whether (and, perhaps how thoroughly) it makes the semantic value of that term sensitive to contextual parameters, not of *what kind of* contextual sensitivity it posits.

[12] That said, it doesn't require thinking of information and standards as the two semantic *parameters*. We can still understand those as being the modal base and ordering source, with the orthodox Kratzerian picture. Compare the debate between Björnsson & Finlay (2010) and Dowell (2013: 174-176).

[13] Dowell (2013: 175) herself criticizes Björnsson & Finlay's (2010) view on the grounds that, since they think that all uses of 'ought' are information-sensitive, they cannot account for "objective" uses of 'ought'. However, I suspect that Björnsson & Finlay may be using 'information-sensitive' in my sense, rather than Dowell's. If that is so, her criticism misses its mark.

Here's one very simple proposal that can seem initially natural. On any contextualist view, it is supposed to be the *speaker's* context that determines the values of the contextual parameters. So maybe the natural thing to say is that the modal base consists of the speaker's knowledge, and the ordering source consists of the speaker's normative standards.

It's worth noting right away that this is far from the only view available to contextualists. It's true that contextualists say that it's the speaker's context that matters. But this general claim is a loose one with various possible interpretations. One might hold that sometimes, a speaker's context is such that the relevant body of information, or the relevant standards, are not those of the speaker herself. Indeed, even the very simple view just described will have to say something more than that the relevant standards are "the speaker's" standards. For that doesn't explain how context selects whether it is the speaker's *moral* standards, or her *epistemic* standards, or her *prudential* standards, or whatever. Even with that qualification, few contextualists accept the very simple view as described. Nevertheless, suspicion of the contextualist view amongst metaethicists may be linked to a mistaken assumption that making the semantic value of 'ought' relative to the speaker's context unavoidably means making it relative to the speaker's own knowledge and standards.[14]

(b) Parochial contextualism

Nevertheless, there is a family of contextualist views that retain a degree of similarity to the very simple view. I'll call these views, as indicated in the introduction, *parochial* forms of contextualism. Parochial forms of contextualism hold that the ordering source parameter is generally occupied by the positive norms or standards that the speaker, or others in the speaker's environment, actually subscribe to. The very simple view is a form of parochialism. But so are views which resemble the very simple view for the ordering source parameter but give a more complex account of how the value of the modal base parameter is determined. And so are some views that allow somewhat more flexibility in how the value of the ordering source parameter is determined, allowing, for example, that it can be the standards of a locally salient *group*, or even of a locally salient individual who is not the speaker, that fill this parameter.[15]

However, what the parochial view crucially *doesn't* say is that the ordering source is filled by the *objectively true* or *objectively correct* normative standards. Consequently, parochial contextualism offers a way to accept a descriptivist, truth-conditional semantics for 'ought', that allows that 'ought'-claims can be straightforwardly true or false, while avoiding realist metaethical commitments. Strictly speaking, parochial contextualism doesn't entail metaethical anti-realism. It could be that there *are* objective, mind-independent normative standards, but these never occupy the ordering source parameter. But the combination of realism and parochial contextualism is an odd one. On such a view, though there might be (for example) objective, mind-independent moral standards that require us not to murder, the fact that there is such a requirement could never be picked out by the sentence "you

[14] A similar confusion arises in discussions of contextualism about 'knows' in epistemology: see e.g. Hawthorne 2004: 85-91, and, for a clarificatory response, DeRose 2009: 246.

[15] The possibility of such flexibility is noted even in Harman's early version of the view: see Harman 1975: 10-11.

ought not to murder".[16] On such a view, there *is* a mind-independent normative reality, but we're imprisoned within a language where we can never make claims about what it's like (at least, not using our most common, ordinary normative vocabulary, such as 'ought') – instead, all our 'ought' claims are just claims about what is required at our own, local normative standards. (Note that this view is the precise inverse of an error theory, on which there is no mind-independent normative reality, but we're imprisoned within a language which ineliminably presupposes that there is one.) So parochial contextualism is much more natural on an anti-realist metaethical view.

That said, there is a very simple objection to parochial contextualism that I think is fatal, which is that it is simply too liberal with truth. This objection, or something close to it, is often framed as a problem about disagreement. To simplify, we'll consider it as it applies to the very simple kind of parochial contextualism on which the ordering source is filled by the speaker's normative standards, though it can be adapted to apply to more sophisticated parochial views. Framed as a problem about disagreement, the worry is that parochial contextualism fails to explain how speakers with conflicting normative standards disagree. When a speaker S1 says "A ought to Φ", and another speaker S2 says "A ought not to Φ", parochial contextualism (in its simple form) appears to say that the two speakers are not disagreeing, for their claims express different propositions: crudely, S1's claim is that given *S1*'s standards, A ought to Φ, while S2's claim is that given *S2*'s standards, A ought not to Φ. Plainly, both those things could be true, and they do not contradict each other. But this seems to be the wrong result; S1 and S2 *do* seem to be disagreeing. This objection, of course, is hardly a new one: it is the classical objection to traditional moral "subjectivism", of which parochial contextualism is a sophisticated kind.

In fact, however, I think that the problem for parochial contextualism is not really best framed as a problem about disagreement. The real underlying problem – that the theory is too liberal with truth – is in one way broader, and in one way narrower, than the problem of making sense of disagreement. It is narrower in the sense that the issue of making sense of disagreement is one that all thoroughgoing contextualist theories, and not just parochial ones, need to deal with. Specifically, the problem of disagreement is also urged as an issue for any version of contextualism that allows for information-sensitivity.[17] It is broader, on the other hand, because even if the problem of disagreement is solved, the problem of being too liberal with truth remains. There are sophisticated contextualist proposals for how to deal with the problem of disagreement, generally centering on the idea that there can be meaningful disagreement between two speakers without them expressing claims that contradict one another.[18] Even if that's right, though, we still don't want a theory that attributes truth to

[16] To clarify, parochial contextualism doesn't preclude the standards that *happen* to be the mind-independent, objective ones from occupying the ordering source parameter: after all, the relevant speaker or group might subscribe to those standards. But this is incidental; what makes them the operative standards are the speakers' subscribing to them, not their mind-independent truth.

[17] Cf., e.g., MacFarlane (2014: 284-5), who presses the disagreement problem for contextualism with reference only to information-sensitivity. However, it's possible that the disagreement problem is easier to solve with respect to information-sensitivity than with respect to the sort of standards-sensitivity envisaged by the parochial contextualist, since it's more clear (in my view) that there *is* a deep disagreement between those who have different normative standards than that there's a deep disagreement between those who have different background information.

[18] Cf. Björnsson & Finlay (2010); Plunkett & Sundell (2013); Finlay (2014: ch. 8, 2017); Silk (2017); Khoo & Knobe (2018); Bolinger (ms.).

utterances that are plainly false – which the theory might still do if it resolves the disagreement problem in a way that does not appeal to inconsistency.

The objection that parochial contextualism is too liberal with truth is even more flat-footed than the disagreement objection. According to parochial contextualism, provided that one subscribes to standards according to which one ought to Φ, one's utterance of "I ought to Φ" will be true.[19] So, when very evil people say that they ought to do very evil things, and subscribe to normative standards that require them to do such things, we are forced to concede that they speak truly.[20] (I'll spare you the filling-out of this schematic objection-form with Hitler examples.[21]) And more generally, any of us can (in one, very real sense) make all of our normative utterances true just by subscribing to the relevant standards.[22] This simply makes truth for normative utterances too cheap. Notably, this bad result is *not* delivered by other anti-realist-friendly semantic theories such as relativism and expressivism.

(c) Aspirational contextualism

Though committed parochial contextualists aren't moved by this objection, many metaethicists are, and I suspect that hostility toward contextualism is largely driven by the association of contextualism with parochial contextualism specifically. Indeed, the majority of prominent contextualists do seem to be parochial contextualists.[23] However, there are exceptions to this. A minority of contextualists

---

[19] Of course, we have to finesse the objection a bit to apply to views that don't mechanically make the relevant standards always depend on the individual speaker – but it doesn't take much.

[20] Could the points made to defend contextualism against the disagreement objection be generalized to deal with this one too? I think not. The best candidate is Khoo & Knobe (2018), who use experimental data to show that in at least some moral exchanges between two speakers where the two speakers express apparently contrary moral claims, subjects are less inclined to say that one of the speakers must be "incorrect" than they are to say that the two parties disagree. However, two points limit the upshot of this in the present context. First, though subjects are *less* inclined to say that one of the parties are incorrect than they are to say that the two parties disagree, their responses to the former question are still around the midpoint of the scale used. Thus, though the results show some capacity for judgments about disagreement and judgments about incorrectness to come apart, they don't suggest that subjects are strongly inclined, in absolute terms, to deny that one of the speakers has to be incorrect. Secondly, even for those subjects who *do* deny that one of the speakers has to be incorrect, this denial does not entail the claim, endorsed by parochial contextualism, that both subjects speak *truly*. Some of these subjects might instead be operating on a folk theory whereby the notions of truth and falsity (and correctness and incorrectness) are out of place in (some) normative disputes.

[21] No doubt it's partly a desire not to want to foot-stomp about how Hitler's normative utterances were false that leads to the framing of the objection in terms of *disagreement* rather than in terms of any one particular party speaking falsely. But even if it's less elegant, I don't think we should ultimately be reticent about foot-stomping about how Hitler's normative utterances were false.

[22] It should be readily conceded to the parochial contextualist that we don't make any underlying *proposition* true by subscribing to a standard; the underlying proposition has the form "given standard S, one ought to Φ", and one doesn't make that proposition true by subscribing to standard S. Instead, one affects the truth of one's utterances by affecting which propositions those utterance express. Still, we can still object to the claim that one can make all of one's normative *utterances* true by subscribing to the relevant standards. If it seems like our utterances sometimes don't get to be true this cheaply, something is wrong with a theory that says they do.

[23] E.g. Harman (1975, 1996); Dreier (1990); Brogaard (2008); Björnsson & Finlay (2010); Khoo & Knobe (2018); also Finlay (2014), *modulo* his relativization to ends rather than standards. Silk (2017) is officially neutral between parochialism and aspirationalism (2017: 207-8, 236), but many aspects of his presentation and positive view reveal parochialist assumptions (*ibid.*: 207, 209-10, 212, 218, 226).

endorse an alternative view that I will call "aspirational contextualism".[24] On this view, normative claims typically "aspire to objectivity": that is, they are claims not about what is required by the positive local standards that are in operation "around here", but rather by the objective, mind-independently true normative standards.

Can an aspirational contextualist still be a thoroughgoing contextualist? The answer is that they can be, but the thoroughgoing aspect of contextualism is going to enter as a result of information-sensitivity rather than standards-sensitivity.[25] One might hold that it is against the spirit of (thoroughgoing) contextualism to suggest that the value of any semantic parameter floats free of the speaker's own control: shouldn't speakers be able to determine, through their own intention, what occupies the ordering source parameter? But the aspirational contextualist can, in one sense, accommodate this point. The aspirational contextualist can say that whenever a speaker uses the moral 'ought', for example, the speaker *intends* to talk about what is required by the objective moral standards, where this intention-ascription is given a *de dicto* reading. Thus, in a case where the speaker is mistaken about or ignorant of the objective moral standards, the *genuine* objective moral standards are still what occupy the ordering source parameter (and so the speaker will be prone to speak falsely).

This is highly analogous to what all contextualists will say about the information-sensitivity of 'ought', given the possibility of the aforementioned *fact-relative* 'ought'. When a speaker uses the fact-relative 'ought', she wants to make a claim about what she (or someone else) ought to do, not merely given the information she herself possesses at the time, but given the totality of the facts. In a normal case, the speaker doubtlessly can't fully identify what the totality of the facts consists in. Nevertheless, she can identify the salient body of information under the general description 'the totality of the facts', and intend to pick out whatever body of information fits that description. The aspirational contextualist makes a similar move for 'the objective moral standards'.

I am persuaded that speakers do sometimes intend to talk about what one ought to do according to the objective moral standards. They are not always merely making a claim about what salient local standards require; sometimes, they are intending to make a bolder claim that does not get to be true that cheaply. However, the aspirational contextualist says more than this: that *all* usages of 'ought' (or at least of the moral 'ought') take the objective standards for the ordering source parameter. This is, in my view, unduly restrictive: there is no reason to say that speakers can *never* use 'ought' simply to talk about what the local moral standards require. Indeed, anyone is going to have to admit that 'ought' is sometimes used to talk about what local, conventional *non*-moral standards require: this is the only plausible treatment of the 'ought' of etiquette, or the 'ought' connecting with playing a game like chess. If that's so, what principled ground is there for denying that it can also be used to talk about what the local moral standards require? Moreover, if it's speaker *intention* that primarily determines what fills the ordering source parameter, then the aspirational view requires the claim that speakers *cannot* intend to use 'ought' to talk about what's required by the local moral standards. This is an implausible *a priori* restriction on what speakers can intend.

[24] See Dowell (2012: 283). Wedgwood (2006, 2007, 2016) also seems to be an aspirational contextualist.

[25] As clarified in section 2, however, this doesn't mean that their thoroughgoing contextualism will take effect only on the modal base parameter, since information-sensitivity can also take effect on the ordering source.

(d) Ecumenical contextualism

Consequently, I think we should endorse a compromise between parochial and aspirational contextualism, which I'll call "ecumenical" contextualism.[26] Though I'm not aware of anyone else who endorses this view, the idea is simple: there are some uses of 'ought' that are parochial – picking up on the local standards accepted by some salient group – and others that are aspirational – picking up on the objective standards (if any).[27] Ecumenical contextualism borrows from aspirational contextualism's account of how the latter possibility can occur: the speaker intends to talk about what is required by the objective standards, where this is given a *de dicto* reading, such that the speaker can have this intention even if she has not correctly identified what the objective standards are, and such that it is the objective standards as they *are*, not the objective standards *as the speaker takes them to be*, that fill the relevant parameter. However, it adds that speakers do not always have this intention: sometimes they intend to talk only about what is required by local standards. It is speaker intention, then, that determines whether a usage of 'ought' is parochial or aspirational.[28] Ecumenical contextualism thus allows 'ought' to be more sensitive to speaker intention than either across-the-board parochial contextualism or all-the-board aspirational contextualism does.

Notice that ecumenical contextualism is more thoroughgoing in its contextualism that aspirational contextualism. It acknowledges a degree of standards-sensitivity of 'ought' that goes beyond distinguishing the different flavors of 'ought' (viz. moral, prudential, aesthetic, epistemic, etc). In particular, it acknowledges a degree of standards-sensitivity *within* moral usages of 'ought'. Yet, ecumenical contextualism avoids the extreme liberalism about truth that parochial contextualism entails. It allows that a significant proportion of usages of 'ought' are aspirational, and these usages don't get to be true as easily as parochial usages. True, the account allows that when evil people say that they ought to do evil things, they will be speaking truly *if* their usage is parochial. But that, on reflection, is what ought to be said: if they really are just claiming that *according to their standards*, they ought to do those evil things, they speak truly. What is objectionable about parochial contextualism, we can now clarify, is not its recognition of that fact, but the way it interprets all normative utterances as parochial in this way. For there are times when the evil people do intend to claim something more

[26] The name may call to mind Ridge's (2014) "ecumenical expressivism", but there's no particular similarity between the views. Ridge's view is ecumenical in that it's a *hybrid* (of descriptivism and expressivism), holding that all usages of the normative 'ought' have both descriptive and expressive content. My view is not exactly a hybrid (of parochialism and aspirationalism), but rather a view that allows for *some* usages of 'ought' that are (purely) parochial and some usages that are (purely) aspirational; it's ecumenical in the sense of acknowledging and accommodating both usages, and not trying to assimilate one to the other.

[27] This may be a respect in which ecumenical contextualism is "flexible" (cf. Dowell 2013). But I am not completely clear on what it means for a particular form of contextualism to be "flexible" rather than "inflexible". A first pass at the distinction would be this: inflexible forms of contextualism say that while the particular value of a contextual parameter changes across contexts, there's a more general level of description at which the parameter is always the same, or always filled the same way. For example, saying that the ordering source parameter is always filled by *the speaker's standards* would be a kind of inflexible contextualism. However, it's not obvious that many views will count as "flexible" on this characterization. Even ecumenical contextualism might be parsed as saying that the ordering source parameter is always filled by *the standards that the speaker intends to talk with reference to*. Does that make it inflexible?

[28] The focus on speaker intention is shared with, among others, Dowell (2013). But Dowell is not herself an ecumenical contextualist; she is an aspirational contextualist (see Dowell 2012: 283). Her discussion of speaker intention concerns its production of *information*-sensitivity in uses of 'ought', not standards-sensitivity (in my sense; see section 2).

– that they objectively ought to perform the evil acts, independently of the standards they happen to hold. But the parochial theory still interprets those utterances in a way that makes them come out true, by continuing to relativize their semantic content to the speakers' standards. And *that* is the wrong result.

3. Ecumenical contextualism and the traditional theories of metaethics

Above, I suggested that parochial contextualism sits oddly with realism in metaethics. What about ecumenical contextualism – how does it interact with realism and anti-realism?

(a) Anti-realist (error-theoretic) ecumenical contextualism

In presenting both aspirational and ecumenical contextualism, I stressed that, for aspirational usages of 'ought', it's the mind-independent objective standards *as they are* that fill the standards-parameter, not the mind-independent objective standards *as the speaker takes them to be*. This seems to presuppose that there actually *are* mind-independent objective standards, which contradicts anti-realism as I understand it. However, in fact, one can be an ecumenical contextualist without thinking that there exist objective, mind-independent normative standards. The result is a kind of attenuated error theory: error-theoretic about *aspirational* usages of 'ought', but not about parochial usages. On this view, aspirational usages of 'ought' presuppose that they are objective normative standards, but this is a false presupposition, and as such aspirational usages of 'ought' are faulty – so that (depending on one's views about presupposition) either they are all false, or they are neither true nor false.[29]

The error-theoretic ecumenical contextualist will have to explain how this integrates with the semantic theory of her choice. One proposal, on the broadly Kratzerian semantics, might be that when there *are* no standards to fill the ordering source parameter (as seems to be so, on the anti-realist view, for aspirational usages), all the worlds are (vacuously) top-ranked. This yields the result that most aspirational 'ought'-claims are false,[30] which comports with many[31] versions of error theory. But it also has the result that aspirational 'may'-claims – about what one *may* permissibly do – will tend to come out true, by default. That's a more unorthodox result for error-theorists, who tend to hold that both 'ought' and 'may'-claims are equally tainted by false presuppositions that make them either false or truth-valueless. Such a view is interestingly unusual in that it vindicates the sometimes-popularly-assumed, but usually-rejected-by-philosophers, claim that if there are no objective (moral) standards, then everything is (morally) permitted.

---

[29] Error theory is often associated with the former view, but the latter is also a possible development of it. See Joyce (2001: 6-9); Perl & Schroeder (forthcoming).

[30] The exception would be claims of the form 'you ought to Φ' where Φ-ing is something that one does in every possible world left live by the modal base. It's a bit odd that the error-theoretic view would have to say that such utterances are true, but this is actually a more general problem for the Kratzerian theory: even non-error-theoretic views seem to have the result that what one does in *all* the (live) possible worlds is *a fortiori* something that one does in all the top-ranked possible worlds and thus, on the orthodox semantics, something that, given *any* standards that fill the ordering source parameter, one "ought" to do. So perhaps the error-theorist can borrow whatever more general solution is in the offing to finesse this problem.

[31] But not all; see fn. 29 above.

A different proposal would hold that when there are no standards to fill the ordering source, there are, in the relevant sense, *no* top-ranked worlds. This might be thought to have the truly strange result that all aspirational 'ought'-claims are true, but all aspirational usages 'may'-claims are false – if we allow in our semantics that when there are no top-ranked worlds, vacuously one Φ's in all the top-ranked worlds. But we might deal with this by building a non-vacuity requirement into the semantics, such that 'one ought to Φ' is true iff one Φ's in all the top-ranked worlds, *and* there is at least one top-ranked world. This would then yield the result that both aspirational 'ought'-claims and aspirational 'may'-claims would come out false. Alternatively, the error-theoretic ecumenical contextualist could claim that when there are no top-ranked worlds, a presupposition fails in a way that makes the normative claim in question truth-valueless. So there are various options here.

Either way, the attenuated error-theory suggested by combining anti-realism and ecumenical contextualism is an interesting one, and contrasts interestingly with anti-realist versions of the other forms of contextualism. Combining anti-realism with parochial contextualism yields a view that is not error-theoretic at all (since it allows our normative utterances to be *true* merely in virtue of their comporting with our own standards), whereas combining anti-realism with aspirational contextualism yields a more wide-ranging, traditional error theory, according to which *all* normative 'ought'-claims, or at least all moral 'ought'-claims, are false. It's a strength of the ecumenical contextualist account, in my view, that it yields a less extreme result, allowing that it is *possible* for speakers to use normative 'ought'-claims, even moral 'ought'-claims, in ways that are not aspirational, and so that will not be false or truth-valueless, even if there are no objective normative standards.

(b) Realist ecumenical contextualism

Combining ecumenical contextualism with realism is a simpler affair. On this view, aspirational usages of 'ought' can be true or false, as determined by the objective, mind-independent standards (again, as they are, not as the speaker takes them to be). This contrasts with parochial contextualism, which, as I argued above, has the odd result when combined with realism that although there are objective, mind-independent standards, speakers never succeed in making claims about them with the ordinary normative 'ought'. It also contrasts with aspirational contextualism, however, in acknowledging that there are some normative (again, indeed, moral) usages of 'ought' – the parochial ones – that are true in a way that is not mind-independent. This allows it to be more semantically flexible and to recognize a wider range of ordinary usage. I still think of it as unequivocally realist, however, since it affirms that there *are* mind-independent, objective normative standards.

(c) An expressivist insight

Contextualism, at least in its broadly Kratzerian form, appears to be a *cognitivist* theory. It gives a semantics for 'ought', and other deontic terms, on which they can be straightforwardly true or false. However, in this final subsection, I want to suggest that contextualism is strengthened by the incorporation of an insight that is at least traditionally associated with expressivist theories.

In what has preceded, I've been glossing over what it takes for a usage of 'ought' to be *normative*. Some deontic usages of 'ought', on the contextualist theory, appear to be purely descriptive. In these usages, I am just *reporting* what some set of standards requires: for example, I might just be reporting what the (conventionally fixed) rules of chess say, using 'must' and and 'may' to say what they require and forbid. Similarly when I'm just reporting what you ought to do *given, or according to,* 19th century British etiquette. These usages may be normative in a very broad sense,[32] but there is also clearly a sense – that which I was earlier calling 'robust' normativity, and which I'll try to identify more precisely in a moment – in which they are not (necessarily) normative. I'll use 'normative' in this narrower way, and 'deontic', by contrast, for the broader sense that captures any usage of 'ought' that is connected with rules, standards, what is required, etc – as opposed to non-deontic usages such as the 'ought' of expectation.

One might think that there ought to be a fundamental difference between normative and non-normative usages of 'ought' (and other modals) – whether the latter includes both usages that are broadly deontic but mere descriptive reports of what sets of standards require, and usages that are non-deontic entirely. The contextualist theory, on its own, may seem to gloss over this difference, since it gives a single character for as many usages of 'ought' as possible, and then explains differences in terms of variability in the two parameters. It's not obvious where the difference between normative and non-normative usages of 'ought' would show up here. Indeed, for contextualists, even entirely non-deontic usages differ from deontic usages only in what is occupying the ordering source parameter. For example, for the 'ought' of expectation, the worlds get ordered by how normal, or conforming-to-expectations, they are.[33]

One might suggest that a usage of 'ought' is normative just when the standards that occupy the ordering source parameter really do have genuine normative authority.[34] But this seems like the wrong thing to determine whether a *usage* of 'ought' is normative. A speaker might *take* a set of standards to have genuine authority when they do not have such genuine authority, or vice versa. It seems like whether a speaker is *using* 'ought' normatively should depend on whether *she* takes the relevant set of standards to have normative authority, not whether they actually *do* have such authority.

This suggests a way forward. It seems that what makes the purely descriptive usages of 'ought' purely descriptive, and non-normative in the relevant sense, is that they need not be accompanied by any kind of *endorsement* of the relevant standards. I can tell you what you "ought" to do according to 19th century British etiquette, while entirely rejecting that set of standards as silly and archaic, and taking them in no sense to be authoritative with respect to your action. If that's so, perhaps what's distinctive of normative usages of 'ought' is that the speaker in some sense *endorses* or *accepts as authoritative* the standards that occupy the ordering-source parameter.[35] This incorporates into the theory a dimension of expressivism, construed first and foremost as a theory in philosophy of mind

[32] Cf. the usage in, e.g., Finlay (2014).
[33] Indeed, Knobe & Szabó (2013) plausibly argue that there are "impure" usages of 'ought' where there's no sharp line between a set of standards and a set of expectations, with the worlds being ordered by this standards-expectations hybrid.
[34] Cf. Fogal (2016: 283).
[35] The idea here, and its development over the next couple of pages, is similar in a number of respects to that of Silk (2017: 210, 227, 232-3), though see fn. 42 below.

rather than language.[36] On this view, it's (partially) constitutive of making a *normative* judgment that one be in a conative state of norm-endorsement or norm-acceptance.[37] Such a conative state will typically be accompanied by motivation to comply, hence the association between expressivism and motivational internalism.

This contextualization of such a theory helps to make sense of some expressivist and/or internalist claims that, out of context, can appear somewhat *ad hoc*.[38] For example, R.M. Hare notoriously claimed that any apparent moral 'ought'-claim *not* accompanied by corresponding motivation was a mere "inverted commas" moral judgment, meaning something like "by conventional standards, one ought to Φ".[39] In the abstract, this can seem like an attempt to define away any potential counterexamples to motivational internalism. But in the present context, it seems better-motivated. We need some way of distinguishing the genuinely normative usages of 'ought' from non-normative ones that *are* just descriptive reports like Hare's "inverted commas" judgments, reports (for example) of what conventional morality says. The proposal is that this is to be done according to whether the speaker really herself accepts the relevant standards as authoritative. Now, maybe Hare's proposal is too strong – perhaps actual motivation is not required, but merely tends to accompany what is actually required, namely genuine acceptance of the norm as authoritative. And perhaps his conception of what contrasts with genuine normative judgment needs broadening beyond "inverted commas" judgments specifically relativized to conventional moral standards. Nonetheless, the core idea is the same.

Derivatively on the *judgment* being normative just when the judger accepts or endorses the norm, we can also say that the 'ought'-claim is normative when the speaker is in this underlying state; in one sense, the 'ought'-claim *expresses* this underlying state of normative judgment. Unlike standard expressivist theories, however, the contextualist theory doesn't build this idea into the *semantics* of 'ought'-claims.[40] Consider two speakers who each say "you ought to pass the bread to the right". It might be that both speakers are only making a claim about what one ought to do according to British etiquette. As such, their utterances have the same semantic content, and the same truth-value. But it might still be that one speaker accepts the standards of British etiquette (rightly or, more likely,

---

[36] For that way of thinking of expressivism, see Schroeder (2008: 3).

[37] See e.g., Gibbard (1990). Could the relevant attitude instead be a cognitive one, a belief to the effect that the relevant standards are authoritative? I worry that this overintellectualizes things, and requires too much sophistication in order to count as making a normative judgment. A conative attitude of *endorsing as authoritative*, by contrast, involves affective and motivational dispositions to treat the standards as authoritative in certain ways, without requiring the same cognitive sophistication that a belief would.

[38] Proto-contextualists Harman (1975: 8) and Dreier (1990) saw their theories as ways of accounting for the *datum* of motivational internalism, without going expressivist. This reflects a time when motivational internalism was more widely accepted. Additionally, though, it only makes sense for a simple version of contextualism that is both parochial and focused narrowly on *moral* language. On such a view, the thought goes, the operative moral standards in play tend to be the *speaker's* standards, and so of course they are standards she endorses. The picture is complicated greatly when we allow standards that the speaker does *not* herself endorse to occupy the ordering source, and when we focus on a wider range of deontic language than the moral, including whole categories of norms (such as those of etiquette) that a speaker might reject the authority of.

[39] Hare (1952: 164-5).

[40] Finlay (2014: 130-1) appears to think that holding that *normative* judgments necessarily involve conative, motivating attitudes, but that non-normative 'ought'-judgments do not, requires us to posit *ambiguity* between normative and non-normative senses of 'ought'. I do not see why this should be so: the difference is plausibly not semantic at all, and even if it were, it could be accounted for by context-sensitivity and not by ambiguity.

wrongly) as genuinely authoritative, while the other does not. The first would then be in a state of mind of making a *normative* judgment to the effect that one ought to pass the bread to the right – we might say that they accept this content *under a normative guise*[41] – and thus of making a normative 'ought'-claim, while the latter would not.[42] Interestingly, then, we cannot say whether an 'ought'-claim is (robustly) normative or not just by examining its subject-matter: it depends on the state of mind of the person making the judgment. (Of course, that's compatible with the question of which standards *really are* authoritative being independent of such states of mind.)

Note that the distinction between normative and non-normative judgments (and usages of 'ought') cross-cuts that between aspirational and parochial usages of 'ought'. The first person described just now uses 'ought' in a parochial way, relativized to a set of conventional standards,[43] but – in virtue of genuinely accepting the standards in question as authoritative – still makes a normative judgment. One might hold that this is always a *mistake* – that one *should* only accept standards as authoritative when they are objective and mind-independent.[44] But whether it's a mistake or not, it's possible to do otherwise, and our theory of language and mind should acknowledge that.

A trickier question is whether there can be aspirational but non-normative usages of 'ought': whether one can take there to be objective, mind-independent standards, but *not* accept those standards as authoritative.[45] I'll leave that open. Either way, for aspirational usages of 'ought' that *are* normative, the present theory will have to once again invoke a kind of *de dicto* reading of the relevant attitude of acceptance of norms as authoritative. As I said earlier, aspirational usages of 'ought' take for the ordering source parameter the *actual* mind-independent, objective standards (if any), not the standards that the speaker *takes* to be the mind-independent, objective standards. So for aspirational usages of 'ought', the operative standards may be ones that the speaker herself *doesn't* accept (*de re*), simply because she isn't aware that they *are* the mind-independent, objective standards. Instead, she needs to accepts the authority of the objective, mind-independent standards, where this is given a *de dicto* reading: she accepts that, whatever the objective, mind-independent standards are, they have authority.

4. Conclusion

[41] Compare Dreier (1990: 18-19).

[42] Following Silk (2017), we might explore ways in which the first speaker *pragmatically* communicates (without semantically asserting) her acceptance or endorsement of the norm in some way, and use that to mark the way in which her usage of 'ought' inherits the *normative* status of her underlying judgment. However, I want to resist Silk's (*ibid*: 226) suggestion that the only normative dimension in our 'ought'-claims consists in implicated content about "what norms to accept". It would be a very odd result if the only normative dimension to our speech concerned the *second-order* normative question of what norms (about what to do) to accept, and could not concern the *first-order* normative question of what to *do*. So instead I say that when the speaker accepts the relevant norms or standards, they accept the first-order 'ought'-claim under a normative guise.

[43] Someone might try to equate normative usages with aspirational ones by claiming that the person who uses the 'ought' of etiquette normatively must be *thinking* of the etiquette standards as mind-independent and objective. If that's a psychological claim – as it needs to be for these purposes – I think it's false. It's psychologically possible to realize a set of standards is conventional but to treat it as having genuine normative authority – even if that's a mistake.

[44] Something like this seems to be implicit in the arguments of Enoch (2011: chs. 2-3).

[45] Some naturalist realists such as Brink (1989) seem to think that one can.

The idea behind ecumenical contextualism is simple, but the possibility of the view is strangely overlooked. More broadly, many metaethicists are unaware of the possibility of a form of contextualism that is not parochial, and that need not be bundled with anti-realism. Ecumenical contextualism shows how we can have such a view while still being thoroughgoing contextualists – even about standards. Instead of trying to shoehorn all our usages into a parochial mode, or all of them into an aspirational mode, it allows us to recognize the wide variety of intentions speakers can have in using normative language, as they are.

References

Björnsson, G. & Finlay, S. (2010). "Metaethical Contextualism Defended," *Ethics*, 121/1: 7–36.
Bolinger, R.J. (ms). "Metalinguistic Moral Disagreement."
Brink, D.O. (1989). *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
Brogaard, B. (2008). "Moral Contextualism and Moral Relativism," *Philosophical Quarterly*, 58/232: 385–409.
Chrisman, M. (2015). *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. Oxford: Oxford University Press.
DeRose, K. (2009). *The Case for Contextualism*. Oxford: Oxford University Press.
Dowell, J. (2012). "Contextualist Solutions to Three Puzzles about Practical Conditionals," *Oxford Studies in Metaethics*, 7: 271–303.
-------- (2013). "Flexible Contextualism about Deontic Modals: A Puzzle about Information-Sensitivity," *Inquiry*, 56/2-3: 149–78.
Dreier, J. (1990). "Internalism and Speaker Relativism," *Ethics*, 101: 6–26.
Enoch, D. (2011). *Taking Morality Seriously*. Oxford: Oxford University Press.
Finlay, S. (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University Press.
-------- (2017). "Disagreement Lost and Found," *Oxford Studies in Metaethics*, 12: 187-203.
Fogal, D. (2016). Book Review of Stephen Finlay's *Confusion of Tongues*. *Ethics*, 127/1: 281-88.
Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.
Hare, R.M. (1952). *The Language of Morals*. Oxford: Oxford University Press.
Harman, G. (1975). "Moral Relativism Defended," *The Philosophical Review*, 84: 3–22.
-------- (1996). "Moral Relativism," in Harman & Thomson, *Moral Relativism and Moral Objectivity*. Oxford: Blackwell.
Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
Joyce, R. (2001). *The Myth of Morality*. Cambridge, UK: Cambridge University Press.
Kaplan, D. (1989). "Demonstratives," in Almog, Perry & Wettstein (eds.), *Themes from Kaplan*. Oxford: Oxford University Press.
Khoo, J. & Knobe, J. (2018). "Moral Disagreement and Moral Semantics," *Noûs*, 52/1: 109-143.
Knobe, J. & Szabó, Z.G. (2013). "Modals with a Taste of the Deontic," *Semantics & Pragmatics*, 6/1: 1-42.
Kolodny, N. & MacFarlane, J. (2010). "Ifs and Oughts," *Journal of Philosophy*, 107/3: 115-143.

Kratzer, A. (1981). "The Notional Category of Modality," in Eikmeyer & Rieser (eds.), *Words, Worlds, and Contexts*. Berlin: de Gruyter.
-------- (1991). "Modality," in Von Stechow & Wunderlich (eds.), *Handbuch Semantik*. Berlin: de Gruyer.
-------- (2012). *Modals and Conditionals*. Oxford: Oxford University Press.
MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.
Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
Perl, C. & Schroeder, M. (forthcoming). "Attributing Error Without Taking a Stand," *Philosophical Studies*.
Plunkett, D. & Sundell, T. (2013). "Disagreement and the Semantics of Normative and Evaluative Terms," *Philosophers' Imprint*, 13/23.
Ridge, M. (2014). *Impassioned Belief*. Oxford: Oxford University Press.
Schroeder, M. (2008). *Being For*. Oxford: Oxford University Press.
Silk, A. (2017). "Normative Language in Context," *Oxford Studies in Metaethics*, 12: 206-39.
Viebahn, E. & Vetter, B. (2016). "How Many Meanings for 'May'? The Case for Modal Polysemy," *Philosophers' Imprint*, 16/10.
Wedgwood, R. (2006). "The Meaning of 'Ought'," *Oxford Studies in Metaethics*, 1: 127–60.
-------- (2007). *The Nature of Normativity*. Oxford: Oxford University Press.
-------- (2016). "Objective and Subjective 'Ought," in Charlow & Chrisman (eds.), *Deontic Modality*. Oxford: Oxford University Press.