

Metalinguistic Negotiations in Moral Disagreement*

Renée Jorgensen Bolinger

bolinger@princeton.edu

Abstract

The problem of moral disagreement has been presented as an objection to contextualist semantics for ‘ought’, since it is not clear that contextualism can accommodate or give a convincing gloss of such disagreement. I argue that independently of our semantics, disagreements over ‘ought’ in non-cooperative contexts are best understood as indirect metalinguistic disputes, which is easily accommodated by contextualism. If this is correct, then rather than posing a problem for contextualism, the data from moral disagreements provides some reason to adopt a semantics that allows contextual variance in the meanings of ‘ought’.

Keywords: ‘Metalinguistic Negation’ ‘Deontic Modals’ ‘Contextualism’ ‘Disagreement’ ‘Metalinguistic Denial’

An increasing number of theorists have proposed some flavor of contextualism about normative language, seen most clearly in the treatment of deontic modals such as *ought* that build at least one context-sensitive parameter into the semantics.¹ Because this implies that when the *ought*-claims are indexed to different parameters, the proposition expressed by one speaker’s assertion of ‘S ought ϕ ’ is logically consistent with another’s of ‘S ought not ϕ ’, deontic contextualism has been charged with failing to accommodate and explain substantive moral disagreement. After surveying the types of cases thought to raise the disagreement problem for deontic contextualism, and briefly glossing the dominant responses, I propose that the disputes that pose the biggest challenge are amenable to being analyzed as metalinguistic in form, and note that such an analysis would explain why the linguistic data in these disputes patterns the way that it does. I do not attempt to show that all moral disagreement is metalinguistic; I only draw attention to a family of features that strongly suggest a dispute is metalinguistic and demonstrate that these features are present in many of the cases that deontic contextualism has so far struggled to explain. This suggests that contextualists need not give a distinct account of such disputes, and may even be leveraged to provide some reason to prefer a contextualist semantics over an invariant one.²

*My thanks to Andrew Alwood and Stephen Finlay for helpful comments on earlier drafts of this paper. Thanks also to James Dreier and audience members at the 2015 meeting of the Central division of the American Philosophical Association, Fabrizio Cariani, Maegan Fairchild, Nathan Howard, Shieva Kleinschmidt, Howard Nye, Caleb Perl, and David Plunkett for helpful discussion.

¹Björnsson and Finlay (2010); Cariani (2016); Carr (2015); Charlow (2013); Dowell (2013); Katz et al. (2012); Silk (2014), and von Fintel (2012) are some among the many who build at least one speaker-sensitive parameter into *ought*. While the majority defend contextualizing to an information-state, a variety of other parameters have also been offered: von Fintel (2012) argues for contextually-set goals (ordering sources), Carr (2015) advocates a parameter for decision rules, while Charlow and Katz, Portner & Rubinstein (2012) propose to accommodate decision-rule type information by appeal to several parameters for various sets of goals and some type of merge operation on the original ordering source.

²The analysis I offer is most appealing to theorists who want to give a contextualist treatment both to information-relative

I review the motivation for going contextualist in §1 before moving on to present the disagreement problem, a preliminary response, and some troubles for the response in §2-3. In §4 I overview classic indicators of metalinguistic negation, and note in §5 that moral disagreements satisfy the necessary conditions for classification as metalinguistic disputes. §6 presents some new data that the natural intonation patterns of the disputes strongly parallel clearly metalinguistic disputes over context-sensitive concept terms, providing some positive reason for a metalinguistic analysis. Finally §7 discusses how this helps address the disagreement problem for deontic contextualism, and gestures at some other ways in which it can be theoretically fruitful.

1 Sensitivity Found

A contextualist semantics for normative terms like ‘good’ and ‘ought’ is attractive for many reasons. One of the most obvious is that it allows us to capture the fact that practical *oughts* can be relativized to a variety of parameters, including information states, goals, and decision rules. The relativization can be made explicitly, as in

He ought to leave now *to get to the party on time*.
Given what he knows, he ought to accept the bet.

or left implicit (by suppressing the italicized portions) in context. The same phenomenon recurs in more weighty moral contexts, like Kolodny and MacFarlane (2010)’s miners puzzle:

Ten miners are trapped [together] either in shaft A or in shaft B, but we do not know which. Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft, but not both. If we block one shaft, all the water will go into the other shaft, killing any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft, will be killed.³

Since we don’t know where the miners are, it seems clear that

1. We ought to block neither shaft.

It is equally clear that since the miners are in fact in one of the two shafts, it would be best if we block whichever shaft they are actually in. Supposing that they are in fact in shaft A, it seems that

2. We ought to block shaft A.

Furthermore, since we do not actually know which shaft the miners are in, there seems to be no conflict between (1) and (2). Of course, if we came to know that the miners were in shaft A, then we would deny (1) and endorse only (2).

A broadly Kratzerian contextualist semantics for normative terms like *ought* can easily explain this, and allows us to give a unified account of the practical and moral uses of the term. Simplifying a bit, on such a picture ‘ought ϕ ’ is true relative to an information state *I* and ordering source *O*

and standards-relative *oughts*, but it may have broader support. Those who embrace invariant foundational metaethical accounts may still want to acknowledge that sometimes speakers’ claims about *oughts*, reasons, etc., are indexed to differing parameters, and yet instantiate genuine disagreement (rather than mere talking-past). The discussion in this paper should be of interest for these purposes as well.

³Kolodny and MacFarlane (2010, 115). They credit the case to Derek Parfit (“What We Together Do,” unpublished) who credits Donald Regan, *Utilitarianism and Cooperation* (New York: Oxford, 1980), p. 265n1.

iff ϕ holds in all the highest-ranked-according-to- O possibilities consistent with I .⁴ (1) and (2) are relativized to information states such that they are propositionally equivalent to (3) and (4), respectively.⁵

3. We ought (relative to our current information-state) to block neither shaft.
4. We ought (relative to full information/the actual state of the world) to block shaft A.

There is no intuitive conflict when our current information does not include the fact that the miners are in shaft A. If we did come to possess that information before we had to decide which action to take, (3) would be false, so we would deny (1) and assert (2).

2 Disagreement Lost

These benefits come with a major cost: as a variety of critics have argued, the contextualist analysis seems to rule out the possibility of genuine disagreement when speakers differ in information, goals, or decision rules, because there is no single proposition that is the subject of disagreement.⁶ To illustrate, consider a simple deliberative case:

Advice

Deliberator is facing the miner case, and has just concluded

DELIBERATOR: We ought to block neither shaft.

when Advisor, who has just seen that the miners are in fact in shaft A, unexpectedly arrives by helicopter. Eager to ensure that all the miners' lives are saved, Advisor interjects

ADVISOR: No that's wrong; you ought to block shaft A.

The surface form of the disagreement appears simple: one speaker asserts P , the other rejects it and asserts an incompatible proposition. But the contextualist must deny this. According to her, Deliberator's *ought*-claim is true: relative to his information-state, blocking neither shaft is the best option. Advisor's claim is also true, relative to her richer information, but since it is indexed to a different information state, it expresses a proposition that is consistent with Deliberator's.

Accounting for disagreement thus presents contextualists with two problems: the first is accounting for the *sense of conflict*. If the two claims are indexed to different information states, they seem to simply report truths about those distinct states; but then Advisor is, at least at the propositional level, changing the subject rather than directly engaging Deliberator's assertion. The second problem is to explain the felicity of Advisor's *semantic assessment* of Deliberator's claim as wrong or false. After all, according to the contextualist, the proposition expressed by Deliberator is "relative to my current body of evidence, the highest-ranked possibilities given the aim of saving the most

⁴For the original statement of this style of view, see Kratzer (1977). To allow evaluation of actions rather than propositions, subsequent proposals often modify the original structure somewhat, holding that 'ought ϕ ' is true relative to a set of alternatives A , information state I , and ordering source O iff, relative to I , no other alternative in A does better than ϕ with respect to the ideals in O . This is typically taken to mean that ϕ satisfies a superset of the ideals satisfied by any other alternative, but some theorists have opted for a different way of understanding the relation, in order to handle cases where alternatives satisfy different elements in O .

⁵Adding an information state parameter may be the most popular contextualist move to solve this type of problem, but it is worth noting that many different parameters (varying deontic ideals, decision rules, etc.) can be used instead to similarly resolve the puzzle. The important point for my purposes is that puzzles like these motivate *some* type of contextualist semantics.

⁶MacFarlane (2007), Francen (2010), Kolodny and MacFarlane (2010) are a few among many who press this objection to contextualism either about moral terms or predicates of personal taste.

miners involve blocking neither shaft”—which is true, and which Advisor has no reason to think false.⁷

One prominent contextualist response to these challenges appeals to the conversational pragmatics of situations like *ADVICE*, and denies that propositional incompatibility is necessary for disagreement between speakers. Björnsson and Finlay (2010), for instance, reject the assumption that the aim in deliberative contexts is to determine the truth value of a particular ‘ought’ claim. Rather, agents are primarily interested in achieving their aims or promoting their goals. As a consequence, their interest in the truth values of oughts are “derivative and instrumental”: if Advisor’s utterance delivers information about how best to achieve the agent’s ends, then the primary purpose has been served and it is of little import that it *strictly speaking* concerns a proposition distinct from Deliberator’s original assertion. Since the aim of deliberation and advice in these cases is to determine how to promote the relevant end, rejections and acceptances function as disapprovals or recommendations of actions. The proposition ‘at issue’—to which the assessment applies—is just the most contextually salient proposition, and need not always be the one asserted by the previous speaker’s utterance. In advice cases, the salient proposition for the Advisor’s assessment is the one that Deliberator’s utterance would have expressed if relativized to Advisor’s informational context.

3 Fall From Grace: Non-Cooperative Disagreements

That’s well and good, but the deontic contextualist is not out of the disagreement woods unless she can meet the twin challenges in more robust disagreements as well. The practical *ADVICE* case is strange, as disagreements go, in that both parties presumptively share the same goals, values, and decision rules; they differ only in that one party is more informed. While we can describe them as disagreeing, the disagreement is unstable and asymmetric:⁸ knowing that Advisor has more complete information, Deliberator cannot felicitously stick to her guns and reply “No, you’re wrong; we definitely ought to block neither.” These features make it especially amenable to resolution by appeal to conversational aims and shared goals, but most cases of moral disagreement — especially between agents *not* presently facing an urgent practical problem — do not have these features. When agents are equally well-informed of the non-moral facts, but subscribe to different moral theories, we may expect that both will be inclined to stick to their guns in a dispute like the following:

(Non-Cooperative) Theoretical Moral Disagreement

Jack and Jill are arguing about a hypothetical involving a probabilistic variant on the miner case, where though ignorant which shaft (of A or B) all the miners are in, they know that there is a 97% chance that they are in A, and therefore a 97% chance that blocking A will save all 10. Suppose that Jack and Jill agree about all the non-moral facts (have the same information), but Jack accepts the decision rule R1, while Jill accepts R2:

R1: Do the action with the best worst outcome.

R2: Do the action with the highest expected value.

Jack will think that they ought to block neither shaft, while Jill will think that they ought to block shaft A, since that is the option most likely to save the most miners.⁹ It is natural that

⁷Björnsson and Finlay (2010) refer to these problems as the *integration* problem and the *semantic assessment* problem, respectively. Sundell (2011) distinguishes these as the ‘sense of disagreement’ and the ‘licensing denials’ problems.

⁸Finlay (2014b) notes this and builds it into his account of disagreement in cooperative cases.

⁹While I have here followed Carr (2015) in developing the case in terms of divergent decision rules, a parallel case for ordering-source variance can be generated just by stipulating that either Jill has a lower threshold for acceptability of options (*a la* the semantics proposed by Katz et al. (2012)), or some of the goals in her ordering source are distinct from those in Jack’s.

they would express this disagreement in something like the following dialogue:

JACK: They ought to block neither shaft.

JILL: No, you're wrong. What they ought to do is block shaft A.

These speakers may very well be aware that they are employing differing standards in making their *ought* claims, yet still take themselves — and be generally taken by others— to be disagreeing. It doesn't look like we can simply extend a solution from deliberative contexts to these non-cooperative disagreement cases. We could say that acceptance of some moral standard α involves having a general interest in promoting behavior that conforms to α , and so even in theoretical moral disagreements agents' interest in the truth of *ought* claims is merely instrumental, really a conflict over *what to do* no different from the advice cases.¹⁰ This isn't promising, for two reasons. First, absent a joint-action context, we lose the explanation for the semantic assessment problem, and as McKenna (2014) points out, attitudinal difference does not on its own license denials.¹¹ Second, deliberative contexts clearly focus on a question of *what to do*, but it is just not clear that that is what is at issue in NON-COOPERATIVE DISAGREEMENT. Jack and Jill's dispute seems best glossed as disagreement about *what reasons to take as relevant* to what they ought to do — whether to determine what they ought to do by reference to R1 or R2 — and only derivatively over whether they ought to ϕ , for any concrete action ϕ .¹² After all, the miners in their dispute are purely hypothetical; the only action question for Jack and Jill to resolve is what to believe about relevantly similar cases. More generally, it seems plausible that the disagreement in many moral disputes is over which goals or principles should determine the set of obligatory or recommended actions, *rather than* over whether some mutually accepted rules and goals imply that some particular action ϕ is in the set.

3.1 Absolutist or Objectivist responses

It's worth noting that deontic contextualism is compatible with a variety of metaethical positions, and the classical disagreement problem only arises for a subset of them. One can solve the problem handily while remaining a deontic contextualist by ensuring that agents in moral disagreements actually operate with the same indexed parameter. Perhaps speakers may coordinate on an index by definite description (e.g. by indexing to 'the true moral standards'). One could also think that speakers may advert directly to an ordering source like *real goodness*, despite epistemic limitation or false beliefs concerning the nature of goodness.¹³ These sorts of realist approaches will work to ensure that speakers in moral disagreement cases are in direct, truth-conditional conflict. There are two major costs to this sort of solution (though theorists antecedently attracted to these views

¹⁰This is the strategy pursued in Björnsson and Finlay (2010) and Finlay (2014a).

¹¹McKenna presses a narrower form of this objection, arguing that merely attitudinal disagreements do not standardly license the full range of disagreement markers. For this objection to have force, it must be that (i) the full range of disagreement markers are in fact licensed in moral disagreements, and (ii) they are substantially less licensed or felicitous in paradigmatically attitudinal conflicts. It's not clear to me that either of these conditions are met, but that data is delicate. McKenna presents Huvenes (2012)'s case to illustrate:

ALEC: Haggis is tasty.

TORFIN (1): Yuck! Haggis is not tasty.

Plausibly ALEC and TORFIN's assertions express conflicting non-doxastic attitudes. Here the preface 'yuck', or even 'no'/'you're wrong' seem felicitous, but this breaks down as soon as we attempt to use an alethic preface:

TORFIN (2): # What you said/think/believe is false. Haggis is not tasty.

If McKenna is right that these sorts of disputes fail to license the full range of disagreement markers, then explaining non-cooperative disagreement as conflict in attitude fails to account for the felicity of alethic markers.

¹²This can be the case even if Jack and Jill are each unable to articulate the precise rule by which they think *ought* should be governed, so long as they have a clear idea of what their privileged rule would dictate in this circumstance.

¹³My thanks to Ralph Wedgwood for suggesting this alternative.

may not consider them to be costs): (i) for the ought-claims to be successfully saturated, it must be that there really is some objective standard to be referred to, and (ii) this approach loses the ability to accommodate *faultless* moral disagreement, since on a realist view, at least one of the speakers must in fact be mistaken.¹⁴ The disagreement problem looms largest for theorists who either are ambivalent or want to avoid commitment to the existence of a single, correct moral standard. Since these are common motivations for going contextualist in the first place, the remainder of this paper will attempt to give an alternative solution to the disagreement problem that does not incur such costs.

3.2 The need for supplementation

In more recent work, Björnsson and Finlay each acknowledge the disanalogy between cooperative and non-cooperative disputes, and offer models of disagreement specifically for non-cooperative cases that do not require conflict in the propositions asserted. Finlay (2014b) explicitly sets aside the semantic assessment problem in order to explain the sense of conflict. He models Jack and Jill’s fundamental disagreement as conflict in expressed actual preferences over ends: their respective claims express preferences for incompatible contents. Björnsson (2015) argues that to *accept* an assertion is to perform one’s part in its conventional function, and since the conventional function of asserting ‘Ought ϕ ’ is to lead hearers to form a corresponding judgment (and make plans accordingly), two ought-claims conflict when the actions required to accept them are incompatible: forming the judgments corresponding to one precludes the judgments required by accepting the other. We count claims ‘correct’ when the judgment they express conforms to the fundamental standards appropriate for judgments of that kind, and incorrect—licensing denials like ‘No, you’re wrong’—otherwise. While this answers the *sense of conflict* problem, and lays a foundation for replying to the assessment problem, it does not explain why denials can take the form of alethic prefaces (e.g. ‘that’s false!’), and makes no attempt to explain the linguistic mechanisms involved.

Sundell (2011) tackles a parallel set of challenges for predicates of personal taste, noting that there are multiple models of disagreement found in natural language disputes, only some of which require conflict in *content* or expressed propositions. Most significantly for the deontic contextualist, Sundell observes that in ‘context disagreement’ involving gradable terms like ‘tall’, agents who are equally well-informed of the relevant height-facts may “disagree about what level of height is the salient standard for tallness” by making claims about whether a given individual counts, given the appropriate standards, yielding a dispute like

A: Ivan is tall.

B: Nuh uh. Ivan is not tall.¹⁵

Sundell follows Barker (2002) in observing that context disagreements may license denials metalinguistically, and Plunkett and Sundell (2013) offer a model of metalinguistic disagreement, but strongly suggest that the object of disagreement is *how to use terms*, rather than which moral standard is appropriate to use.¹⁶ So it doesn’t fit an intuitive gloss of Jack and Jill’s disagreement.

¹⁴The problem of accounting for faultless disagreement is foregrounded in many discussions of the disagreement problem for contextualism, among them Brogaard (2008) and Huenes (2014).

¹⁵This dialogue is Sundell’s case (27/28).

¹⁶Sundell (2011) asserts that “Context Disagreement seem[s] totally incapable of licensing metalinguistic negation” (fn 19) but does not elaborate. Plunkett and Sundell (2013) deny that all non-canonical disagreements can be glossed as occurring through metalinguistic negation, choosing instead to explain it as *metalinguistic negotiation*. They suggest that in many cases of interest to them, the denial utterance does not lend itself to contrastive intonation, does not display the typical licensing effects of metalinguistic negation, and is not internally inconsistent if interpreted as a descriptive negation (a feature they take to be essential to metalinguistic negation). I do not suggest that *all* non-canonical disputes should proceed via metalinguistic

The project I undertake here is to provide a linguistic mechanism to address the assessment problem for these sorts of models, connecting these theoretical defenses of how disputes instantiate genuine disagreement to an account of why we should expect the linguistic data to pattern the way that it does. In the next section I briefly review the main features of metalinguistic negation and denials, and show that a metalinguistic analysis of the denials in moral disagreement cases is not ruled out. I then present some new data that gives positive reason to think that it is the right analysis, and discuss how it helps address the semantic assessment problem while preserving our intuitive sense that Jack and Jill’s disagreement is fundamentally *about* the appropriateness of the moral standards invoked, rather than (merely) how to use terms.¹⁷

4 An (old) new kind of disagreement: Metalinguistic Negation

Laurence Horn (1989) describes metalinguistic negation as “a device for objecting to a previous utterance on any grounds whatever, including the conventional or conversational implicature it potentially induces, its morphology, its style or register, or its phonetic realization.”¹⁸ Typically the objection is introduced by a negation term (‘not’) but can also be fronted by a wide range of disagreement markers, including alethic prefaces (e.g. ‘it’s false that’, ‘it’s not the case that’, and even ‘it’s not true that...’). The preface scopes over an echoic utterance (that is mentioned or referenced), and this material is followed by a correction clause. The offending feature is marked by heavy vocal stress on it or its replacement.¹⁹ In the examples below, (a) rejects a scalar implicature by stressing the corrective ‘two’, (b) objects to the evaluative tinge of the term in the echoic clause, (c) corrects the pronunciation of the previous speaker, and (d) uses an explicitly alethic preface to reject a non-truth-functional aspect of the sentence (a conversational implicature).

- (a) This Birthday Card is NOT from one of your admirers.

It’s from TWO of your admirers.
Happy Birthday from both of us.

- (b) I haven’t DEPRIVED you of my lecture on negation; I’ve SPARED you it.

- (c) We don’t eat tom[a:t{uz} here, we eat tom[eiD{uz].

- (d) It’s not true that they had a baby and got married—they got married and had a baby.²⁰

When used in dialogue, explicit re-utterance of the offending material is not necessary for echoic reference. (b) can be transposed to dialogue form thus:

A: You deprived us of your lecture!

B: No, I’ve SPARED you it.

Metalinguistic negations often occur as rejoinders to an utterance of the corresponding affirmative, and are not truth-functional. The negation scopes over the echoed *utterance* instead of over the negation; only that a specific class of non-cooperative moral disagreements are best understood this way, and that this class of disputes *do* display the hallmarks of metalinguistic negation.

¹⁷Thanks to Andrew Alwood and Stephen Finlay for encouraging me to foreground this difference between my approach and Plunkett & Sundell’s strategy.

¹⁸Horn (1989, p 363.)

¹⁹Horn (1989, p 374.)

²⁰Horn presents (a) as an example in his Horn (1992) paper, (b) and (c) are given this form by Carston (1996) while variations are presented in Horn (1989) and Burton-Roberts (1989b), and (d) is from Horn (1985).

proposition asserted, since rather than ascribing a particular truth value, metalinguistic negations are simply all-purpose rejection markers.²¹

There are two reliable indicators of metalinguistic rejection: first, echoic reference to a reading to be corrected. In cases where the objectionable utterance is part of the immediate context, as in NON-COOPERATIVE DISAGREEMENT, the echoic reference may be accomplished by an anaphoric ‘that’, referring to the first speaker’s utterance. Second, if the sentence in question is *uttered*, there will typically be contrastive vocal stress on the corrected item.²² However, one can easily imagine speaker B in dialogue (b) above having a dry or understated sense of humor, and so opting to not stress ‘spared’ when delivering their reply. Since the utterance clearly employs metalinguistic negation, it appears that even when spoken, contrastive stress is more of a useful heuristic than a strictly necessary marker for metalinguistic negations.

These characteristics give metalinguistic rejections a distinctive prosody (typical pitch, stress and intonation pattern). In addition to forcing a metalinguistic reading with a descriptively inconsistent continuer, the prosody of (5b)—contrastive stress on the offending item ‘old’ and the offered replacement ‘antique’—clearly distinguishes it from the much flatter intonation typical of descriptive negations like (5a):

5a: 
That car isn't old *at all*; it's brand new.

5b: 
That car isn't OLD at all; it's *antique*.

One can use metalinguistic negation to reject a whole utterance, or just a specific part. The direct object of the rejection often receives heavy stress in the echoic clause, and the proposed replacement material is contrastively stressed in the corrective clause.

There are two diagnostic tests for metalinguistic negation: because the negation objects to the utterance, it resists neg-incorporation (morphological incorporation as a negative prefix like *un-* or *non-*), and it fails to license negative polarity items (items like *any* or *ever*, which normally occur only in negative environments) or inhibit positive polarity items (like *some* or *sometimes*).²³ That the negation in (6a) passes the first test is shown by the fact that while the first sentence in the pair below is felicitous, the second, which incorporates the negation with *n-*, is not:

6a. Maggie isn't EITHER patriotic or quixotic—she's both!

6b. # Maggie is neither patriotic nor quixotic—she's both!²⁴

²¹Burton-Roberts (1989a,8) suggested that the echoic clause must logically contradict the corrective clause if taken literally, but many of Horn's original examples—and some of Burton-Roberts'—do not have this feature. Burton-Roberts (1999) modifies this view to a pragmatic contradiction view, in response to Carston (1996)'s arguments that metalinguistic negations need not semantically contradict. Horn (1989:144) also emphasizes that metalinguistic negation often requires double processing (the hearer first attempts to understand the content as literal, descriptive, reprocessing as metalinguistic only when no consistent descriptive content is recoverable) to be correctly understood, but as Carston points out, this feature depends on the absence of vocal marking: a pronounced vocal stress pattern makes the metalinguistic reading immediately available.

²²Carston (1996) argues that there is no requirement that all instances of metalinguistic negation be utterances (and therefore no requirement that they have pronounced stress patterns).

²³Geurts (1998); Horn (1989).

²⁴Examples from Horn (1985):141.

If the negation in (6a) were operating descriptively, it would simply negate the ‘or’ claim, and so be logically equivalent to (and hence not resistant to formulation as) 6b. The failure to incorporate morphologically as *n-* shows that the negation in 6a occurs metalinguistically.

Positive polarity items normally do not occur felicitously within the scope of a negation; similarly, the occurrence of negation ordinarily licenses use of negative polarity items. Metalinguistic negations generally fail to have this effect because they operate at a different level from the echoic content. The otherwise surprising infelicity of the negative polarity items and felicity of positive polarity items indicate that the negations in 7 and 8 are operating metalinguistically:

7. Chris didn’t manage to solve {#any/some} of the problemshe solved all of them.
8. Chlamydia is not {#ever/sometimes} misdiagnosed, it is frequently misdiagnosed.²⁵

Chapman (1996) notes that this gloss needs a bit of nuance, since negative polarity items *can* sometimes occur felicitously within the scope of a metalinguistic negation. For example the negation in (5b) above is clearly functioning metalinguistically, objecting to ‘old’ and offering the replacement ‘antique’ in the corrective clause. Since being old is entailed by being antique, there is no consistent interpretation of (5b) that takes the negation to be descriptive. Nevertheless, the negative polarity item ‘...at all’ occurs felicitously within its scope.²⁶

The needed nuance is supplied by attending carefully to which elements of the rejection are occurring echoically, and which are functioning as part of the rejection. Speakers are free to quote or echo less than the full utterance to which they are responding; they need only echo the objectionable bit. Material not occurring echoically operates at the same level as the metalinguistic negation and is free to interact compositionally with it, and a speaker is free to encode the negation in any form she chooses. So in cases where a negative polarity item (NPI) occurs felicitously within the scope of a metalinguistic negation, “the NPI is part of the form in which the MN is encoded, and not part of what it is used to object to.”²⁷ Minimal forms for 5b’s rejection include “Not OLD; *antique*” and “No, it’s antique.” But speakers are also free, as Chapman notes, to encode the metalinguistic rejection more complexly; rather than (5b) to reject to an utterance a of “That car is old”, a speaker could felicitously use

- 5c. That’s not true at all. It’s *antique*.

Here ‘not true at all’ is part of the metalinguistic rejection, rather than a truth-functional evaluation of the utterance being rejected.

5 Testing the Possibility

Diagnosing non-cooperative disagreements as metalinguistic gains initial plausibility from the fact that the rejections in such exchanges typically display the two classic indicators of metalinguistic negations. The second speaker’s utterance in such exchanges typically makes echoic reference to the earlier assertion being rejected. The minimal form of such reference is just the use of a rejection marker like “no” or “that’s false”, but is often more explicit, actually echoing some portion of the rejectable utterance. This is immediately followed by corrective material, the speaker’s own suggestion for what we ought to do. They also exhibit the patterns of contrastive stress typical of metalinguistic disputes: contrastive stress on the rejected element and replacement. The obvious next question is whether NON-COOPERATIVE DISAGREEMENT passes the diagnostic tests.

²⁵Examples from Horn (1989):370, 374, 396.

²⁶My thanks to Caleb Perl for urging more development on this point.

²⁷Chapman 1996:391.

5.1 Polarity Items

Applying the diagnostic tests to disputes involving *ought* is not entirely straightforward. While the standard examples of a positive polarity item (PPI) are ‘some’ and ‘sometimes’, weak deontic necessity modals (*ought* and *should*) are themselves mobile PPIs, taking wide scope relative to the clausemate descriptive negation.²⁸ To illustrate, though the sentence structure in

9a. I don’t think John should marry Susan.

is [Not [I think [should [John marry Susan]]]], it is paraphrasable as

9b. I think John should not marry Susan.

in which the modal takes wider scope than the negation: [I think [should [not John marry Susan]]]. In fact it’s nearly impossible to obtain the narrow scope reading of 9a without building in a substantial and very specific background context.²⁹ *Should/ought* will take this wide-scope reading unless either (i) rescued, (ii) shielded, or (iii) the negation is metalinguistic, and so not a PPI-inhibitor.³⁰ To show that NON-COOPERATIVE DISAGREEMENT passes the polarity item test, it suffices to show that a narrow scope reading of *should* is available despite the absence of shielding or rescuing.

To test whether this reading is available, let’s consider a variant of the miner case:

Ten miners are distributed between two mine shafts, A and B. There are 9 in one and 1 in the other, and there is an 80% chance that the 9 are in shaft A. The shafts threaten to collapse; we have time to block one but not both. Blocking a shaft guarantees that all the miners in that shaft are saved, but it will kill everyone in the other shaft. If nothing is done, both shafts will flood completely and anyone in either shaft will die.

Since Jack subscribes to R1 he will judge that blocking A is just as good an alternative as blocking B (and both are better than blocking neither), while since Jill subscribes to R2 she would endorse the option which maximizes expected value. Jack may assert 10a or 10b, and Jill can use 10c or 10d to reject his claim:

10a. They should block *some* shaft.

10b. They ought to either block shaft A or block shaft B.

10c. No! They shouldn’t block {*some/#any*} shaft; they should block shaft A!

10d. No! They oughtn’t EITHER block shaft A or shaft B; what they *ought* to do is block shaft A!

Neither 10c nor 10d facilitate shielding or rescue. So if a narrow-scope reading of the modal is available, it’s because the negation is metalinguistic rather than descriptive. To make it easier to see that a narrow scope reading *is* available, I’ve embedded the PPI *some* in 10c. A wide-scope reading of the modal with a descriptive negation would inhibit positive polarity items and licence negative polarity items, requiring us to replace the PPI *some* in 10c with the NPI *any*. Notice though that in 10c the PPI *some* occurs felicitously; replacing it with *any* results in contradiction with the continuer, causing the utterance to crash. Negative polarity items are not licensed by metalinguistic negations, so if they can’t occur felicitously in the echoic content on its own, they shouldn’t be felicitous when

²⁸On a syntactic analysis of the kind Homer (2015) advocates, this means that even when it originates lower in the sentence structure, *ought* covertly moves to a position higher than the descriptive negation. However, there are other possible analyses, and it is not necessary for my purposes to embrace a syntactic analysis. My thanks to a referee for this journal for this point.

²⁹Homer (2015, p 34.)

³⁰A PPI is ‘shielded’ from negation when strong scalar terms like *always* or *necessarily* intervene (Homer, 2015, p 24). It is ‘rescued’ when there is another downward-entailing expression that outscopes the PPI (Szabolsci, 2004). Iatridou and Zeijlstra (2013, p. 532.) Iatridou and Zeijlstra demonstrate that for *must*, contrastive stress on the negation induces a metalinguistic reading, allowing the modal to take narrow scope. This point also holds for *should* and *ought*.

that content is embedded under a metalinguistic negation. The infelicity of ‘any’ in 10c shows that negation in this NON-COOPERATIVE DISAGREEMENT passes the polarity test.

There is a complication worth flagging here: metalinguistic negations can be used to target implicatures, and so one available explanation of 10c is that while the negation is definitely metalinguistic, it is applied to ‘some’, occurring echoically, rather than to *ought*. The result shown here is therefore a weak one, that the linguistic data in moral disputes is not inconsistent with a metalinguistic reading of the denial. It does not show that the *ought* itself is targeted by the denial.

5.2 Neg-Incorporation

If the negation in 10d were operating descriptively, it should be able to be felicitously incorporated. Since *not...either* is descriptively equivalent to *neither...* , incorporation would allow us to paraphrase 10d as 10e:

10e. No! They ought to neither block shaft A nor shaft B; what they *ought* to do is block shaft A!

In fact 10e is infelicitous, and not equivalent to the rejection 10d. So the negation in 10d resists neg-incorporation and passes the second diagnostic test.³¹

These considerations show that metalinguistic rejection is not ruled out as a viable model for the disagreement involved in NON-COOPERATIVE DISAGREEMENT. The model provides a schematic answer to the linguistic *assessment* problem: the speakers directly disagree in that the first speaker does something with his utterance — so far we haven’t specified what — that the second speaker rejects and counters with hers, and speakers are licensed to use a wide range of disagreement markers to introduce a metalinguistic rejection. As a defensive move, this is valuable enough to merit exploration on its own. The contextualist might even push the metalinguistic line further, leveraging it as positive evidence in favor of a contextualist semantics for *ought*. The remainder of this paper explores this more aggressive strategy.

6 A Suggestive Parallel

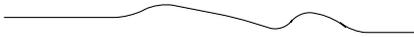
The second speaker’s utterance in non-cooperative moral disagreements often displays a distinctive prosodic pattern of focus stress that (i) is consistent with echoic reference to and correction of the first speaker’s statement, and (ii) closely parallels natural prosody for clearly contextualist disagreements (e.g. disagreements over tallness). In disputes over the appropriate parameters for context-sensitive terms, rejection often proceeds indirectly, by offering an object in the correction clause that forces selection of different parameter values. There are two readings available for the toy dialogue below, one on which the dispute occurs under a fixed value for ‘tall’, and a second under which the standards for ‘tall’ are themselves at issue:

Julia: Annie is tall.

Nathan: That’s false! Annie isn’t tall; Bobbie is tall.

The two readings can be distinguished by the prosody of Nathan’s rejection. A descriptive negation that merely disputes whether Annie falls under the extension of a common standard for tall has relatively flat intonation, stressing only the objects,

³¹Importantly, affixed contractions are not instances of morphological incorporation. So, the fact that it’s felicitous to phrase the rejections as ‘oughtn’t’ or ‘shouldn’t’ doesn’t speak to the neg-incorporation test. In fact, using an affixed contraction is the only way to make it clear that the negation is clausemate rather than embedded below the modal (Homer, 2015, p. 35 fn 25), and so necessary in order to screen out other explanations for a narrow scope reading (forcing by superordinate negation) or movement (resolution as an embedded negation).

Nathan (1): 
 That's false! Annie isn't tall; *Bobbie* is tall.

whereas adding heavy stress on *tall* and contradiction intonation (a final rise in the intonation of the rejected content) yields a clearly metalinguistic reading:

Nathan (2): 
 That's false! Annie isn't TALL; *Bobbie* is TALL.

The prosodic profile of this second reading is complex, in that rather than falling only on the rejected term 'tall', there is also some stress on the objects (Annie and Bobbie). This likely results from the way the discourse structure constrains the appropriate pitch and placement of stress (or what Jackendoff (1972) called A- and B-accent) in utterances like Nathan (2). As Buring (2003) details, while *focus stress* identifies the strategy for addressing a question under discussion, *contrastive topic* stress identifies the proffered answers. The placement of these stress patterns in an utterance indicates whether the speaker is contributing an answer to the conversationally presupposed subquestion, or changing the subquestion.

In ordinary object-level disagreement, the speaker accepts the setup of the question and simply contests the *answer* given. In this dispute, the question is "which children are tall?" Julia's answer, 'Annie', presupposes a subquestion: "who is tall-according-to-*these*-standards?" Nathan (1) accepts this question, and is simply contesting the answer. Focus stress on 'tall' indicates this, accepting the presupposition that we should evaluate tallness by the standard Julia invoked, and contrastive stress highlights his counter answer (Bobbie). But when the speaker's rejection is aimed at the presupposed parameter values, they are in effect disagreeing about the subquestion to be discussed. So in Nathan(2), we should expect to hear contrastive stress on the parameterized term ('tall'), and focus stress on the answers which presuppose the speaker's preferred parameterization. When they aim to change the question under discussion, speakers are not free to stress only the objects, or only the parameterized terms; to effectively reject the parameter they must reject the presupposed strategy. So if 'ought' patterns together with 'tall'—if speakers are sometimes disagreeing *by* rejecting the presupposed parameter values—we should expect a similar double-stress pattern, with contrastive stress falling on 'ought' and focus stress on the action options.

Provided that Nathan and Julia both have access to the height facts about Annie and Bobbie, Nathan's rejection serves to assert something about the relevant standards and range of application of 'tall'. In particular, it asserts that the relevant parameter for determining what counts as tall is such that Annie does not count, and Bobbie does. Nathan (2) does not deny that Annie is tall-according-to-Julia's-presupposed-standards. But the vocal stress pattern indicates that he is rejecting that concept as inappropriate, and offering a replacement.

Extending this model to moral disagreements, we should understand the second speaker's utterance as indirectly rejecting the original speaker's chosen parameters for *ought*. By asserting that the highest-ranked alternative, given the context, is ϕ , she communicates that the parameters appropriate to the context are such that ϕ is the highest-ranked. Since in non-cooperative disagreements the speakers have the same background information, the rejection utterance functions to assert the contextual appropriateness of a different ordering source or decision rule for *ought*. This contextualist analysis fits the intuitive understanding of what is at issue in non-cooperative disagreements.

If the disputes in NON-COOPERATIVE DISAGREEMENT-type cases are metalinguistic in this way, and used by speakers to replace the initial saturated *ought*-claim with an *ought* contextualized to the speaker's preferred parameters, we should expect a prosodic pattern that is closely parallel to the NATHAN(2) rejection. This predicts that speakers will often place noticeable stress on the *oughts* and set off the rejected claim with a final rise in intonation:



That's false! We OUGHTN'T block neither shaft; we OUGHT to block *shaft A*.

It's worth empirically testing whether these predictions are accurate. In a small-scale study, I asked participants to speak four toy dialogues in the way that felt most natural: two tallness disputes (one explicitly holding standards fixed, one varying them), a standards-invariant *ought* case, and the moral disagreement case.³² The first tallness case gave a baseline for disputes involving direct propositional disagreement: the speakers accept the same standards of application for the concept 'tall', and disagree only about which objects fall in the extension. Default intonation for the disagreeing utterance was relatively flat. Nearly all respondents (17/18) heavily stressed the replacement object while not stressing the second occurrence of 'tall', and the majority (14/18) contrastively stressed the offending object.

When presented with a dispute explicitly framed to involve standards variance, typical intonation patterns were far more dynamic. Speakers placed heavy emphasis on each contested term (both occurrences of 'tall', the rejected object *Annie* and the replacement object, *Bobbie*).

Ethel and Fred are arguing about which children on the playground to count as tall. Ethel points to Annie, a 3 ft. 6 inch first-grader, and asserts (E). Fred disagrees, and points to Bobbie, a 4 ft. 8 inch first-grader, while saying (F):

E: Annie is tall.

F: That's false! Annie isn't tall; the tall one is Bobbie.



That's false! Annie isn't tall; the tall one is Bobbie.

This is an indirect metalinguistic dispute. The speakers are not really disagreeing about whether, holding the meaning of *tall* fixed, a given individual is tall; they are disputing what parameters should determine what counts as tall.

The default prosody for the moral disagreement type cases closely parallels the profile for standards-variant contextual disputes. Almost every respondent (16/18) heavily stressed the contrasting *oughts*, and heavily stressed the replacement object *shaft A*, with a diminished but significant number stressing the initial object *neither shaft*:

G: We ought to block neither shaft.

H: That's false! We oughtn't block neither shaft; we ought to block shaft A!

³²This data is from a highly informal study, and records only a very small sample size (18 respondents.) Respondents were screened to ensure no prior contact to this literature, though some had taken an introductory course in semantics, and two were professional academics. All were native English speakers who had completed high school, and most (14/18) held at least a bachelors' degree. I think the data is nonetheless informative, and perhaps merits a more formal study. More complete information on the study, including the actual text of the prompts, is included in the appendix.



That's false! We oughtn't block neither shaft; we ought to block shaft A!

This exhibits double-stressing of objects and the concept term, with the heaviest stress falling on the second occurrence of the concept term, suggesting that the second *ought* is the offered replacement for whatever was objectionable in the rejected content. This pattern is exactly what we should expect of an indirect metalinguistic dispute focused on what parameters should determine what counts as what we *ought* to do, but is unexplained otherwise.

7 Disagreement Regained

Though this data is far from conclusive, there seems to be enough to justify exploring this strategy's potential to answer the *semantic assessment* problem for moral disagreement. When a dispute is metalinguistic, the conflict focuses on the utterances, rather than some at-issue proposition. So this picture does not require us to identify some single propositional content with respect to which Jack and Jill hold incompatible attitudes (alethic or otherwise). Instead, there must just be *something* about Jack's utterance that Jill finds objectionable; the considerations canvassed so far leave open *what*. If we are convinced by the close parallels with other contextualist disputes, we have reason to think that the relevant feature could simply be sensitivity to the wrong inputs for the context-sensitive parameters.

7.1 The main challenges, answered

You might worry that this trivializes moral disagreements, reducing them to pedantic linguistic quibbles. However, the fact that a dispute is linguistic—that the primary issue under contention is which concept to use in a context—does not imply that the disagreement is somehow meaningless, superficial, or trivial. Quite to the contrary, the choice of concept often determines which of our other attitudes apply to the object. In the case of moral disagreements over the salient *ought* concept, disputants can be understood as disagreeing over what sorts of considerations to count as reasons in favor of an action being the morally preferred or required action. While the dispute is *directly* over terms, it is indirectly over the fundamental structure of relevant moral concepts. Such a disagreement is in many ways deeper than disputes over how to promote a given end, or whether a target proposition is true, since resolving it requires alignment on not just the verdicts but also underlying reasons.

It's important to stress that diagnosing a dispute as metalinguistic is a claim about the form, rather than the content, of the disagreement. While I argue that the *form* of these disputes is metalinguistic, I do not claim that the disagreement is *about* the terms used. In the same way that metalinguistic rejections of scalar implicatures are *about* the inappropriateness of the implicature, *via* the rejection of the term generating the implicature, in the moral case the disagreement is at heart about the inappropriateness of using a certain set of standards, and proceeds via the metalinguistic rejection of the *ought* indexed to those standards. So while the dispute is metalinguistic and occurs at a level of ascent from the propositions expressed by the sentences/assertions, it is not fundamentally about language.

This kind of proxy-dispute—superficially expressed over the terms but *really* concerning the appropriateness of particular social practices which are cued by the term's application—is by no

means peculiar to moral disagreement. Disputes of this form have also been discussed for a range of other disagreements, for instance whether to apply a slurring term to a group (e.g. in Bolinger, 2017; Nunberg, 2017), whether to use a term like ‘athlete’ in accord with a concept that includes horses (Plunkett and Sundell, 2013), which objects (if any) should be included in the scope of thick terms like ‘pure’, ‘chaste’, or ‘lewd’ (Vayrynen, 2013), and which things to quantify over when making ontological claims (Thomasson, 2017). In all of these cases, the explicit subject of the metalinguistic dispute—which of several candidate concepts to term to mean—is somewhat independent of the real core of the speakers’ disagreement: the appropriateness of cued practices and attitudes. The connection to the latter is not, on most accounts, any part of the semantic content of the terms being contested; they are simply the conventional implications of it being appropriate to call something ‘lewd’, or an ‘athlete’, or to speak of it ‘existing’, or to say it is what we ‘ought’ to do.

This, I think, yields a more satisfying explanation for the *sense of conflict* in moral disagreements than those canvassed in §2. It agrees with Finlay, Björnsson, and others’ suggestions (reviewed in §3.2) about the location of the core disagreement, but provides us with new resources to simultaneously answer the *semantic assessment* problem. To explain the felicity of alethic prefaces, we need only recall that metalinguistic rejections license use of a wide range of disagreement markers, including ‘that’s completely false’ and ‘it is not true that’. Furthermore, since on this account these markers function to front the rejection of a previous utterance, the metalinguistic strategy is not vulnerable to the charge that it can only answer one of the problems at a time. The alethic prefaces are licensed to introduce precisely the kind of disagreement that the account identifies as the real conflict between the speakers’ assertions.³³ Paired with a conception of the state of disagreement as a disagreement in attitude over which moral standards to adopt, this account is well-positioned to explain why the intonation data patterns the way it does. When the disagreement is one over standards, rather than immediate practical action, it must express itself as disagreement over principles. We come at that sort of project slant, indirectly shifting the standards by insisting on shifted extensions.

7.2 Extensions to other cases

I have focused the discussion up to this point solely on non-cooperative disagreements, but one might wonder whether the metalinguistic analysis can be extended to cooperative advice contexts. While such an extension is possible, it does not fit well, since cooperative disagreements function (and prosodically pattern) differently. In non-cooperative contexts, an assertion of the form ‘No, x is F ’ (‘No, you ought to ϕ ’) serves to advocate for a conception of F foreign to the one the agent had been using. By contrast, in cooperative contexts the focus is on communicating the fact that ϕ is in the extension of F_2 , where F_2 is close enough to the concept F DELIBERATOR had been using that she can be expected to readily accept F_2 as the relevant concept. The point of asserting ‘No, you ought₂ to ϕ ’ in such a context is to highlight that ϕ is the thing *to do*, whereas in the non-cooperative cases it asserts that *ought*₂ is the concept to use.

It is also consistent with my analysis that some cases of moral disagreement are not metalinguistic. This is obviously the case for non-verbal *states* of disagreement, but given the mixed nature of actual disagreements we should not be surprised if some disputes with moral content are primarily

³³As noted earlier, it isn’t entirely clear how great a range of disagreement markers are felicitous in OUGHT-disagreements. One recurring (unsolicited) comment from survey respondents was that the preface ‘that’s false!’ felt unnatural. One participant noted ‘I probably would never say “that’s false”; I would say, ‘Oh I don’t agree’, or something else, so it was hard to come up with a natural way to say “that’s false”.’ Though this reticence is possibly due to speakers’ more general aversion to evaluating moral claims as being ‘true’ or ‘false’, given the delicacy of this data, it’s a significant strength of the metalinguistic interpretation that it permits speakers to use alethic prefaces, without predicting that they will be obviously felicitous to all speakers.

deliberative. In fact, one of the benefits of the account outlined here is how flexible it allows us to be. In each of our imagined disputes, there are two possible objects of disagreement: what to do, and what parameters to use. Speakers can align on one of these while diverging on the other. On this account, whether the interlocutors count as disagreeing in a context depends on how they choose to engage. If they foreground divergence on the *ought*-parameters, they're engaged in a metalinguistic moral disagreement. If they foreground divergence over options, they are engaged in a deliberative dispute. If instead they foreground convergence (perhaps because they diverge only with respect to the parameters, and primarily care about what to *do*), they count as at least partially agreeing.

As a consequence, the metalinguistic account is also well-suited to explain an additional problem, the problem of *Partial Agreement*. This puzzle arises in cases where two speakers have different values, goals, or information, but share a judgment about what ought to be done. Dreier (2009) notes that these cases are problematic for standard attitudinal accounts, which locate disagreement in a clash of practical aims. To make the problem vivid, he offers the following case: if you think people with headaches ought to just tough it out, and I think they should take painkillers, it seems that we disagree about what people with headaches ought to do. But if we both have headaches and there is just one dose of ibuprofen left, then my aims do not conflict with yours: I (*qua* person in pain) aim to take the painkiller, while you (*qua* person in pain) aim to just tough it out without attempting to take a painkiller. So, Dreier notes, 'my aims are met to the extent that you manage to act on yours. There is no practical clash. And indeed the clash would come precisely if we agreed that each of us ought to try to grab the Ibuprofen.' (Dreier, 2009, 105). The metalinguistic model provides a tidier solution to this sort of case. Whether the two parties count themselves as disagreeing in a partial-agreement case is determined by what their focus is: if they are engaged in a deliberative project, then alignment on the immediate question of what action to jointly take is all that matters, and they will report themselves as agreeing. If they are philosophers engaged in a theoretical moral debate, they will likely foreground their disagreement, changing practical questions if need be to find a case where their preferred standards no longer align in recommendation.

To see how this plays out linguistically, suppose Jill receives some additional information such that her preferred decision rule now recommends blocking neither shaft. Her dispute with Jack may now take the following form:

JACK: We ought to block neither shaft.

JILL (2): That's true, but not for the reasons you think.

The speakers agree in immediate practical plans, but disagree in attitude concerning which standards for *ought* to accept. So the "that's true" echoes the utterance of the previous speaker, not the proposition expressed; the corrective clause goes on to highlight that the speaker is making a metalinguistic move and advocating different value for a parameter of the context-sensitive term. Notice the parallel to a case of partial agreement over *tall*:

Nancy and Nathan are classifying basketball players as 'tall' or 'not tall.' Nathan knows that Nancy considers anything over 2 feet in height 'tall', but he is far more demanding in his standards, requiring an entity to be above average height for the relevant category to deserve the predicate. They might have the following exchange of partial agreement:

NANCY: LeBron is tall.

NATHAN: That's true, but not for the reasons you think.

Both Joan and Nathan's utterances can be understood as echoing the anaphorically referenced *utterance* picked up by 'that's true', in order to call attention (via the corrective '...not for the reasons you think') to the parameter they wish to modify. This can be felicitously made explicit by appending the continuer "...it's because it avoids risking killing anyone" (or in Nathan's case, "...because he's

over 6'7"). The perfect parallel between the felicitous assertion patterns of these two toy dialogues suggests that they should receive a unified explanation, and so speaks in favor of a contextualist semantics for *ought*.

8 Conclusion

Since metalinguistic rejections dispute which concepts to use, and how, rather than the truth-values of propositions, this frees us from having to find some proposition whose truth-value is the target of the speakers' disagreement. It also provides an explanation of (i) why speakers take themselves to be genuinely disagreeing about *oughts*, and (ii) why alethic prefaces can be felicitous in these disagreements. This strategy is available to any variety of contextualism, since it introduces no machinery and requires no substantive assumptions (beyond the assumption that metalinguistic negation is a genuine phenomenon.) Consequently, it appears that contextualists do after all have the resources to account for moral disagreement.

9 Appendix: Study

Participants were all native English speakers, recruited through facebook; there were 18 respondents, all but two were non-academic professionals. The prosodic pattern of each respondent is represented by an individual line in these graphs, while the patterns presented in the paper are the averages. Respondents were given the instructions, and presented with the five dialogues below:

I am conducting an informal study on the natural intonation patterns for various forms of disagreement. If you'd like to help out, use the web widget below to record yourself saying all of the dialogues below in the way that sounds most natural to you. (No need to repeat the stories, just the lettered assertions).

Dialogue 1- You and I are trying to count the number of tall trees in a field. We agree to count trees over 10 feet tall as 'tall', and any tree shorter than 10 ft as 'not tall'. There is a 7 ft fir tree, and a 12 ft redwood tree. I make assertion (A), and you correct me with (B):

A: The fir tree is tall.

B: That's false! The fir tree isn't tall; the redwood tree is tall.

Dialogue 2- You and I have agreed that what we ought to do is buy the cheapest cheese. There's a \$3 cheddar, and a \$5 brie. I make assertion (C), and you correct me with (D):

C: We ought to buy the brie.

D: That's false! We oughtn't buy the brie; we ought to buy the cheddar.

Dialogue 3- Ethel and Fred are arguing about which children on the playground to count as tall. Ethel points to Annie, a 3 ft. 6 inch first-grader, and asserts (E). Fred disagrees, and points to Bobbie, a 4 ft. 8 inch first-grader, while saying (F):

E: Annie is tall.

F: That's false! Annie isn't tall; the tall one is Bobbie!

Dialogue 4- Ten miners are all trapped in one of two mines shafts (shaft A or shaft B). We are 95% confident that the miners are in shaft A, but we're not sure. Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft, but not both. If we block one shaft, all the water will go into the other shaft, drowning any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft, will be killed. I think that we ought to do the action that has the best guaranteed outcome, while you think we ought to do what's most likely to have the optimal outcome. I express my judgment with (G), and you object with (H):

G: We ought to block neither shaft.

H: That's false! We oughtn't block neither shaft; what we ought to do is block shaft A!

Dialogue 1 should elicit intonation typical of descriptive disagreement, since it explicitly states that all parties to the conversation have the same standards of application for the concept 'tall'. Nearly all respondents (17/18) heavily stressed the replacement object *redwood* while not stressing the second occurrence of 'tall', and the majority (14/18) contrastively stressed the offending object *fir tree*:



That's false! The fir tree isn't tall; the redwood tree is tall.

The majority of responses to dialogue 2 closely paralleled the intonation pattern in dialogue 1; two respondents deviated from the rest in heavily stressing both the *oughts* and the objects. All contrastively stressed *brie* and heavily stressed the replacement object *cheddar*. A little over half still gave noticeable stress to the *oughts*, which may indicate that the metalinguistic intonation is natural enough that it is difficult to set up a context where no speakers will default to it:



That's false! We oughtn't buy the brie; we ought to buy the cheddar.

Dialogue 3 is set up to elicit a metalinguistic reading by making salient the fact that the disputants are using different standards of application for the concept *tall*. As expected, the typical intonation places heavy emphasis on each replacement term (the second occurrence of 'tall' and the replacement object, Bobbie), and also noticeably stresses the offending terms (tall and Annie) in the echoic material:



That's false! Annie isn't tall; the tall one is Bobbie.

Utterances of dialogue 4 prominently display intonation typical of the metalinguistic reading: almost every respondent (16/18) heavily stressed the contrasting *oughts*, and heavily stressed the replacement object *shaft A*, with a diminished number stressing the initial object *neither shaft*.



That's false! We oughtn't block neither shaft; we ought to block shaft A!

References

- Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25:1–36.
- Björnsson, G. (2015). Disagreement, correctness, and the evidence for metaethical absolutism. In Shafer-Landau, R., editor, *Oxford Studies in Metaethics*, volume 10. Oxford University Press.
- Björnsson, G. and Finlay, S. (2010). Metaethical contextualism defended. *Ethics*, 121:7–36.
- Bolinger, R. J. (2017). The pragmatics of slurs. *Nous*, 51(3):439–462. doi: 10.1111/nous.12090.
- Brogaard, B. (2008). Moral contextualism and moral relativism. *Philosophical Quarterly*, 58(232):385–409.
- Buring, D. (2003). On d-trees, beans, and b-accent. *Linguistics and Philosophy*, 26(5):511–545.
- Burton-Roberts, N. (1989a). *The Limits to Debate*. Cambridge: Cambridge University Press.
- Burton-Roberts, N. (1989b). On horn's dilemma: presupposition and negation. *Journal of Linguistics*, 25:95–125.
- Burton-Roberts, N. (1999). Presupposition-cancellation and metalinguistic negation: A reply to carston. *Journal of Linguistics*, 35(2):347–364.
- Cariani, F. (2016). Deontic modals and probabilities: One theory to rule them all? In Charlow, N. and Chrismann, M., editors, *Deontic Modality*. Oxford University Press.
- Carr, J. (2015). Subjective ought. *Ergo*, 2(27).
- Carston, R. (1996). Metalinguistic negation and echoic use. *Journal of Pragmatics*, pages 309–330.
- Chapman, S. (1996). Some observations on metalinguistic negation. *Journal of Linguistics*, 32(2):387–402.
- Charlow, N. (2013). What we know and what to do. *Synthese*, 190:2291–2323.
- Dowell, J. (2013). Flexible contextualism about deontic modals: A puzzle about information-sensitivity. *Inquiry*, 56(2):149–178.
- Dreier, J. (2009). Relativism (and expressivism) and the problem of disagreement. *Philosophical Perspectives*, 23:79–110.
- Finlay, S. (2014a). *Confusion of Tongues: A theory of normative language*. Oxford University Press.
- Finlay, S. (2014b). The pragmatics of normative disagreement. In Fletcher, G. and Ridge, M., editors, *Janus-Faced Judgment: Hybrid Approaches to Meta-Normative Theory*. New York: Oxford University Press.
- Francen, R. (2010). No deep disagreement from new relativists. *Philosophical Studies*, 151(1):19–37.
- Geurts, B. (1998). The mechanics of denial. *Language*, 74(2):274–307.
- Homer, V. (2015). Neg-raising and positive polarity: the view from modals. *Semantics and Pragmatics*, 8(4):1–88.
- Horn, L. (1985). Metalinguistic negation and pragmatic ambiguity. *Language*, 61:121–174.
- Horn, L. (1989). *A Natural History of Negation*. Chicago: University of Chicago Press.
- Horn, L. (1992). The said and the unsaid. In *Working Papers in Linguistics (SALT II Proceedings)*, number 40. Ohio State University.
- Huvenes, T. (2012). Varieties of disagreement and predicates of taste. *Australasian Journal of Philosophy*, 90(1):167–181.
- Huvenes, T. (2014). Disagreement without error. *Erkenntnis*, 79:143–154.

- Iatridou, S. and Zeijlstra, H. (2013). Negation, polarity, and deontic modals. *Linguistic Inquiry*, 44(4):529–568.
- Jackendoff, R. (1972). *Semantic Interpretation in Generative Grammar*. Cambridge, MA: MIT Press.
- Katz, G., Portner, P., and Rubinstein, A. (2012). Ordering combination for modal comparison. In *Proceedings of SALT 22*, pages 488–507.
- Kolodny, N. and MacFarlane, J. (2010). Ifs and oughts. *Journal of Philosophy*, 107(3):115–143.
- Kratzer, A. (1977). What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy*, 1(3):337–355.
- MacFarlane, J. (2007). Relativism and disagreement. *Philosophical Studies*, 132:17–31.
- McKenna, R. (2014). Disagreeing about ‘ought’. *Ethics*, 124(3):589–597.
- Nunberg, G. (2017). The social life of slurs. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*. Oxford: Oxford University Press. (unpublished; available online).
- Plunkett, D. and Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms. *Philosopher’s Imprint*, 13(23):1–37.
- Silk, A. (2014). Evidence sensitivity in weak necessity deontic modals. *Journal of Philosophical Logic*, 43(4):691–723.
- Sundell, T. (2011). Disagreements about taste. *Philosophical Studies*, 155(2):267–288.
- Szabolsci, A. (2004). Positive polarity-negative polarity. *Natural Language and Linguistic Theory*, 22(2):409–452.
- Thomasson, A. (2017). Metaphysical disputes and metalinguistic negotiation. *Analytic Philosophy*, 58(1):1–28.
- Vayrynen, P. (2013). *The Lewd, the Rude and the Nasty*. Oxford: Oxford University Press.
- von Fintel, K. (2012). The best we can (expect to) get? challenges to the classic semantics for deontic modals. Paper presented at the Central APA.