

Metaethical Contextualism Defended

Gunnar Björnsson & Stephen Finlay

ABSTRACT: We defend a contextualist account of deontic judgments as relativized both to (i) information and to (ii) standards or ends, against recent objections that turn on practices of moral disagreement. Kolodny & MacFarlane argue that information-relative contextualism cannot accommodate the connection between deliberation and advice; we suggest in response that they misidentify the basic concerns of deliberating agents. For pragmatic reasons, semantic assessments of normative claims sometimes are evaluations of propositions other than those asserted. Weatherson, Schroeder and others have raised parallel objections to standard-relative contextualism; we argue for a parallel solution.

NOTE: The paper by Niko Kolodny & John MacFarlane discussed below is still work in progress and has not been publicly circulated. Please enquire with them before referencing their paper.

Epistemic and deontic judgments (involving modals like ‘must’, ‘may’, and ‘ought’) seem to be essentially relative to different sets of information, and it is a perennially popular view that deontic judgments are additionally relative to different standards or ends. There are three schools of thought on how to accommodate this relativity in the semantics of modals: *contextualism*, *invariantism*, and *relativism*. For many reasons contextualism is a natural and appealing view, but recently it has been the target of a barrage of objections.¹ In this paper we defend a form of contextualism about deontic modals against these objections, and argue that they fail to show any advantage for contextualism’s rivals. Many of our points can be generalized for contextualist treatments of other kinds of terms, but here we limit our focus to the deontic case.

Contextualism, invariantism, and relativism about some set of terms involve contrasting views on how to understand (a) the *semantic content* of the sentences in which they occur (the propositions these sentences express), and (b) the *truth-conditions* of those propositions. According to contextualism, the terms at issue are semantically incomplete, having one or more open argument-places whose values are supplied by the *context of utterance* (if not explicitly).² In other words, speakers

¹ Cappelen & Lepore 2005, Weatherson 2008, Kolodny & MacFarlane ms, Schroeder ms, Streiffer 2003. Contextualism about modals is orthodox in linguistics, following Sloman (1970) and Kratzer (1977).

² Kent Bach (ms) objects to this division of theoretical space that it mistakenly assumes that every complete sentence expresses a proposition. He rejects ‘contextualism’ on the ground that it falsely claims that *sentences* are incomplete without contribution from context, and defends instead a view of modals he calls ‘radical invariantism’, according to which unrelativized sentences in which the modals occur invariantly express the same ‘propositional radicals’, which yield propositions when supplied with additional arguments (explicitly or by context). We do not challenge Bach’s basic point, but we

commonly *say different things* in uttering the same sentence. When an argument is provided, it yields a proposition that has an *objective* truth-value; i.e. the same truth-value for any assessor. According to invariantism, contextualism is right about the objective truth-conditions of the propositions expressed by those sentences, but wrong about the incompleteness of the terms, which rather make an invariant contribution to the proposition expressed; there is no relativity here at all. According to relativism, contextualism is wrong on both counts; the terms make an invariant contribution to the propositions expressed, but these propositions have relative rather than objective truth-conditions.³ In other words, every utterance of such a sentence *says the same thing*, but whether what it says is true or false depends on the standpoint of the person assessing it (the *context of assessment*).

We believe that deontic modals have *two* kinds of relativity to the context of utterance: they are relativized both to information and to some kind of standard or end.⁴ Both kinds of context-sensitivity have recently been challenged, and here we offer roughly parallel defences of each, focusing on ‘ought’—although our claims will generalize. These objections can all be understood as versions of a venerable kind of metaethical argument that turns on a theory’s inability to give a satisfactory account of moral *disagreement*.⁵ More precisely, they arise from difficulties faced by semantically relativistic theories in accommodating the interrelation of the moral claims accepted or asserted by people standing in different relations to the relevant information or standards. We first address problems for information-relativity, and then turn to problems for standard-relativity.

1. *Deliberation and Advice*

Practical questions about what to do are commonly resolved by reaching ought-judgments. In deliberation or first-personal practical reasoning this often leads to decision, intention, and action, while in second- and third-personal practical reasoning it often leads to advice and evaluation. In the real world, these questions typically have to be answered in less than ideal epistemic circumstances, in which the agent, advisor, or judge possess incomplete information. The course of action that is best *given available evidence* is often not the course of action that is best *given all the facts*. This

understand ‘contextualism’ more broadly, as the view (which Bach endorses) that modal verbs must be provided with additional arguments before any proposition is determined.

³ This characterization of relativism assumes a particular view of propositions, however. On an alternative view, relativism rather claims that the sentences at issue express different propositions relative to different contexts of assessment. To keep things (relatively) simple, we ignore this view here.

⁴ This is not necessarily to claim that these modals have separate argument-places for each (as in Sloman 1970, Kratzer 1977), because one argument-place can accommodate both kinds of relata. See Finlay 2008a for one suggestion of how this can be done.

⁵ This argument dates back at least to Moore 1912, ch. 3.

raises a dilemma: is what the agent ought to do a matter of what is best given the evidence, or of what is best given the facts? Both answers have had their champions, but it seems eminently reasonable instead to distinguish between a 'subjective' sense of 'ought' corresponding to *best given the evidence*, and an 'objective' sense corresponding to *best given the facts*.⁶

We cannot rest with a single subjective sense of 'ought', however, because the question remains: 'evidence available' to *whom*? When ought-judgments are offered as advice, they do not aim to tell agents which actions seem best *to them* or relative to their evidence, because they are sensitive to any additional information possessed by the advisor. But neither do they aim to tell agents which actions are objectively best, since advisors are often not in a position to have any confidence about that. 'Ought', the contextualist concludes, must have at least as many surface senses as there are different informational bases; an 'annoying profusion' of senses.⁷ To impose some order on these uses it is natural here to embrace a contextualist semantics for 'ought', as it is for other, nondeontic modals; the practical 'ought' always functions to select the best action, *relative to an informational base in some way determined by the context of utterance*.⁸

Difficulties for contextualism about 'ought' arise from the need to integrate the different ought-judgments issued in deliberation and in advice. These problems are presented by Niko Kolodny and John MacFarlane (ms)—hereafter 'K&M'—who believe that they require us to abandon contextualism for relativism. They illustrate these problems with the following scenario:⁹

Ten miners are trapped together in one of two shafts, A and B, and their lives are threatened by impending flooding. Another person, *Agent*, on whose evidence it is equally likely that the miners are in A as that they are in B, is able to block either A or B with sandbags (but not both), reducing the flooding in the blocked shaft. Agent's evidence suggests that if either shaft is blocked, then all the miners will be saved if they are in it, and all will drown if they are in the unblocked shaft—whereas if the shafts are left open then whichever shaft the miners are in, nine will survive and only one (who is deepest in the shaft) will die.

However, there is another person, *Advisor*, who possesses all Agent's evidence, but has additional information, and the totality of his evidence

⁶ Ewing 1952. Kolodny & MacFarlane (ms) provide a helpful survey of the literature, placing Prichard (1949), Ross (1939), Piller, Scanlon, and 'perhaps' Jackson (1991, 1997) on the side favoring the objective interpretation, and Moore (1903, 1912) and Thomson (1986, 1990) on the side favoring the subjective interpretation.

⁷ Jackson 1991; K&M ms; Finlay ms.

⁸ Sloman 1970; Kratzer 1977; Jackson 1985.

⁹ Adapted by K&M from Derek Parfit.

suggests rather that if A is blocked, then water will nevertheless flood both A and B equally such that whichever shaft the miners are in, nine will survive and one will die; if B is blocked then all the miners will survive if they are in B, but all will die if they are in A (which will be completely flooded); and if neither shaft is blocked then all will survive if they are in A, but all will die if they are in B (which will be completely flooded).

Assigning each life one unit of utility, the expected values therefore look like this:

| | <u>Agent's information</u> | <u>Advisor's information</u> |
|---------------|--------------------------------------|--------------------------------------|
| Block shaft A | $(.5 \times 10) + (.5 \times 0) = 5$ | $1 \times 9 = 9$ |
| Block shaft B | $(.5 \times 10) + (.5 \times 0) = 5$ | $(.5 \times 10) + (.5 \times 0) = 5$ |
| Block neither | $1 \times 9 = 9$ | $(.5 \times 10) + (.5 \times 0) = 5$ |

Intuitively, the appropriate outcome of Agent's deliberation is the judgment that he ought to block neither shaft, choosing the certainty of one death over an even chance of ten deaths. This is compatible with contextualism, interpreting it as the judgment that blocking neither is best relative to Agent's information. The appropriate advice for Advisor to offer is the claim that Agent ought to block A, again preferring the certainty of one death over an even chance of ten deaths. This too is compatible with contextualism, interpreting it as the judgment that blocking A is best relative to Advisor's information. So where is the problem for contextualism?

The puzzle, in K&Ms' words, is that advice 'is characteristically intended to help the advisee arrive at the correct answer to the question about which he is deliberating.' This is, they presume, a question about what the agent *ought* to do. But if the judgments made in deliberation and advice involve different senses of 'ought', then there really is no common question with which both are concerned, and deliberation and advice come apart in a puzzling way. If Agent's deliberation is concerned with the question, 'What ought I to do *given my information*,' then he has successfully reached the right answer when he judges that he ought to block neither shaft. He needs and receives no assistance in this inquiry from Advisor. Hence K&M write that according to contextualism, 'advisers do not address the question about which the agent deliberates. So they hardly seem to be advising at all.'

This problem is avoided, K&M believe, if we reject contextualism for *relativism* according to which ought-claims are relative not to the information of the speaker (the context of use), but to the information of the *assessor* (the context of assessment). Agent's claim expresses an unrelativized proposition, the truth-conditions of which vary according to the information possessed by the person assessing it. Agent is right to be confident of his deliberative judgment, because relative to his information it is indeed

true. But Advisor is right to challenge Agent's claim, because relative to Advisor's information it is false. Because of its ease in connecting deliberation and advice, K&M conclude that relativism must be correct and contextualism incorrect.¹⁰ We believe this conclusion is mistaken.

A preliminary step in defending contextualism is to pay careful attention to what we should identify as the relevant context and informational base. K&M's initial formulation of contextualism construes 'ought' as always sensitive to the *speaker's* information. The contextualism we defend is more liberal, and allows that context can determine the informational base in different ways, as conversational purposes dictate. K&M's claim that deliberators' ought-judgments are relativized to the deliberator's own information is certainly plausible in the case of ought-judgments that conclude *private* deliberation. But in their scenario, Agent asserts his ought-claim in a 'dialogue'; it is a piece of public communication. Plausibly, modal claims made in public communication are by default relativized to information possessed by the participants of the conversation *collectively* (i.e. the union of the information that each possesses individually). Consider that if Agent is aware that others in his conversation have more information than he, it is presumptuous and inappropriate for him to express confident ought-claims himself, and appropriate for him either to express uncertainty or ask for advice about what he ought to do. Agent's conversation includes Advisor, who has more complete information than Agent. The contextualist should interpret Agent as claiming that he ought to block neither shaft, given (at least) *what 'we' know*. Because of Advisor's information this claim is false, and Advisor correctly contradicts it. In this case, deliberation and advice are addressed to the same question.

K&M acknowledge the possibility of a 'sophisticated' contextualism of this kind. They believe that it also fails to escape their objection, because of a problem of 'advice from unexpected sources.'¹¹ Consider the following variation on the original scenario:

Having deliberated to the conclusion that he ought to block neither shaft, Agent is unexpectedly accosted by Physicist, who has been conducting experiments in a neighboring shaft and has new information that definitively places the miners in shaft A. He advises Agent that he ought to block shaft A.

The contextualist is thereby confronted with the following dilemma: either Physicist is, like Advisor, a member of the salient group whose information is relevant to Agent's judgment, or else he is not. If he is a member then, K&M claim, Agent's judgment would seem 'unwarranted and irresponsible'—which they find implausible, for he has reasoned 'well and appropriately'. On the other hand, if the truth of Agent's ought-claim is not sensitive to Physicist's information then we have K&M's original problem:

¹⁰ A similar view is defended in Horty 2006.

¹¹ The same kind of problem arises from the evaluations made by eavesdroppers.

Physicist's advice involves a sense of 'ought' different from that in the conclusion of Agent's deliberations, and therefore cannot be addressed to the same question.

This is an important challenge, but we think that contextualism can actually thrive on both horns of this dilemma. To anticipate in brief, the solution to the first horn is that ought-judgments can be warranted even when they are made in ignorance of some relevant information or source of information, while the solution to the second horn is that advice doesn't need to address the ought-question asked by the deliberator in order to be of assistance to him. We now explain these in turn.

There are two difficulties to overcome on the first horn. The first of these stems from the fact that Agent is oblivious to the existence of Physicist and his information; this may make it hard for us to see how Agent could have intended a judgment that would be sensitive to this new information, which was not salient in his context. In order to capture Physicist's information in its net, the informational base for Agent's judgment would have to be so broad that it could include new information from many sources: sources about which (as in the case of Physicist) Agent cannot be expected to know. A defense of contextualism along these lines would have to give a plausible account on which the informational base for Agent's judgment is of this kind. The second difficulty is that given a net this wide, it may seem that Agent would be unwarranted in reaching any confident judgment about what he ought to do—and we concur with K&M that Agent's judgment is in fact warranted.

In response to the first difficulty, we think that deliberative ought-questions can be (and typically are) sensitive not merely to information presently possessed by the deliberating person or group, but also to all information that is or will become accessible to them prior to their having to decide what to do; call these questions *news-sensitive*. This is because we are typically trying to make the best decisions we can, and other things being equal, decisions are better when based on fuller information. If figuring out what you ought to do were simply a matter of calculating what is best given your present information, then it would never require you to seek fuller information, which it clearly often does. We often think that it is too early to tell what we ought to do; in such cases we are clearly not just interested in what action is best given the information we already possess.¹²

There are a variety of ways of defining an informational base that accommodates this point. Intuitively and rather loosely, it seems that in deliberation we are interested in what action is best given the information we will or are able to acquire in the time

¹² Of course, sometimes the right answer to a deliberative question can be that we ought to seek more information, which presumably is made relative to present information. Not all deliberative questions can be of this kind, however, or it would never be legitimate to deny knowing what we ought to do on the ground that we need more information.

available before we have to make our decision.¹³ If Agent asks his companions what he ought to do, it is therefore plausible that he is asking what course of action is best, given all the information that they are able to gather before his decision must be made; this will include information from unexpected sources, like Physicist's.

It is not true that (as K&M suggest) such ought-judgments are always or typically unwarranted. We are often warranted in thinking that no relevant new information will be accessible to us in time, even where there remains a possibility of such news. It is plausible that Agent is warranted in assuming (falsely) that there is no source of information like Physicist available, and therefore that Agent is warranted in judging (falsely) that he ought to block neither shaft.

This solution drives us into the same issues that we face on the second horn of the dilemma. For the accessibility of Physicist's information depends on whether Physicist chooses to offer it. If he did not make it accessible, then Agent's ought-claim would not have been sensitive to that information, and would have been correct. It follows that Physicist *makes* Agent's claim false by making his information available. The contextualist needs to explain, therefore, how by his intervention Physicist is *helping* Agent answer the question he deliberates over.

Observe, however, that relativism faces a similar challenge here. On K&M's account, to deliberate is to try to determine the assessment-relative truth of ought-propositions *relative to one's own context of assessment*. Though one *could* try to determine its truth relative to some other context of assessment, that is not one's concern as one deliberates. For example, from the trapped miners' context of assessment, it is not true that Agent ought to block neither shaft; rather, he ought to block the shaft in which they are trapped. Agent might be painfully aware of this, but his concern cannot be with which ought-claims are true relative to *their* context of assessment, because that is a context he isn't able to occupy. His concern is to determine which claim is true from his context. The problem for relativism is to explain how Physicist's advice fits into this picture. For Physicist's ought-claim expresses an assessment from *his* context—a context that Agent isn't concerned with, unless Physicist provides his information. Although the relativist can say that Physicist is concerned to correct Agent's judgment, the incorrectness he removes is not one that Agent has been concerned to avoid. So the relativist also needs to explain how advice engages with deliberation. It is everybody's problem.

¹³ The notion of information that we 'will or are able to acquire' can be spelt out in different ways, and we (the authors) disagree amongst ourselves about how best to do this. However, since we think that a satisfactory contextualist account of advice needs to solve the problems of the second horn of the dilemma for reasons related to those raised by K&M, we leave this question open.

The puzzle that relativism and contextualism (on either horn) here face together is how offering advice from unexpected sources could be helping the advisee with the problem he is trying to address. K&M identify this problem as answering a question about what ought to be done, relative to some informational base. What this overlooks is that agents are concerned to make decisions based on the best information they can get. This is because an agent's fundamental concern is to promote and protect certain values or things that matter to him, and *not* simply to reach the correct answer to a predetermined ought-question. In the mine scenario, Agent's fundamental concern is presumably the preservation of life—the more lives the better. Fuller information is desirable because it enables decisions that are better with regard to these values. The contextualist ought to say, therefore, that deliberators have only a derivative and instrumental interest in reaching true and warranted ought-judgments relativized to any particular informational base. In deliberation we seek to determine what we ought to do relative to the information available to us simply because this is the best basis for a decision that is within our reach. But should we suddenly become privy to new information, the former question loses its importance to us and we abandon it for the question of what we ought to do given our new information. K&Ms' assumption that Agent's fundamental concern is with the truth of particular ought-claims is therefore false; it is rather with preventing the loss of lives. Physicist's information enables him to assist Agent in promoting this value.¹⁴

This means that K&M are mistaken in claiming that 'advice is most urgent when the advisee is in danger, without the adviser's help, of arriving at an incorrect answer' to the ought-question over which he deliberates. Even on relativism, the danger that the advisee will make (what is from the advisor's context of assessment) an incorrect ought-judgment cannot be what fundamentally motivates advice. This is because advice is typically inappropriate when the advisor is less well-informed than the agent. Suppose that Agent is better informed than we are about the circumstances. According to relativism, Agent is in 'danger' of arriving at what is from our context of assessment an incorrect answer. But advice is not appropriate in this case.

Once we have recognized that advice is driven by an interest in putting the advisee in a better position to promote his values, we can also see why the second horn of K&Ms' dilemma is not threatening. Suppose that Agent's ought-judgment is made relative to the information he possessed at the time. His judgment that he ought to block neither shaft was therefore correct, and not contradicted by Physicist's advice, which was relativized to a different informational base. But given his fundamental interest in saving lives—the more the better—he would be very much interested in acting instead on

¹⁴ Our treatment here of normative modals parallels Kent Bach's (ms) view of the behavior of epistemic possibility modals. He argues that our real interest is not in any static epistemic possibility proposition, but rather (as he puts it) in what is 'possible now'.

an ought-judgment based on more information. Since this is what Physicist can offer, advice is perfectly appropriate.

2. *Context-Insensitive Assessments*

By taking note of agents' fundamental interest in promoting and protecting what they value, and their consequent interest in basing decisions on better information, we have seen that contextualism is able to integrate deliberation and advice. However, this does not yet establish that contextualism gives a correct analysis of the ought-judgments that people actually make in deliberation and advice. Contextualism's opponents argue that the judgments we actually make do not behave as contextualism predicts. The practices in question involve expressed agreement and disagreement with ought-claims made in different informational contexts.

Consider first the practice of *denial*. It is natural for Physicist, entering unexpectedly, to reject Agent's claim that he ought to block neither shaft by saying,

(1) No, you ought to block shaft A.

This looks like a problem for contextualism. If Agent's and Physicist's ought-claims are relativized to different informational bases, then they are not contradictory. But intuitively it seems that unexpected advisors with better information often do *disagree* with agents' ought-claims. It may seem, therefore, that actual uses of 'ought' in contexts of advice always presuppose that there is just one proposition at issue, as relativism maintains.¹⁵

Of course, we have already observed the possibility of news-sensitive contextualism, which can accommodate this point about advice. But a similar problem arises for any form of contextualism from practices of (mere) *evaluation*, or eavesdropping, in which the assessor is unable to make her information accessible to the agent, and assesses his claims from afar. Suppose Physicist is observing Agent from such a distance, over closed-circuit television. Hearing Agent say that he ought to block neither shaft, Physicist tells his colleague,

(2) No, he ought to block shaft A.

¹⁵ Similar problems arise in diachronic disagreement with oneself (*retractions*). Suppose Agent first judges that he ought to block neither shaft, then acquires new information leading him to say, 'I was wrong/What I said was mistaken. I ought to block A'. Under news-insensitive contextualism, his new judgment does not contradict his former judgment. This seems less of a problem for news-sensitive contextualism, as it secures genuine disagreement for reassessments made prior to the act, while retractions after the act do not seem as natural.

Physicist has expressed disagreement with Agent, but according to contextualism their ought-claims are not contradictory.

Similar difficulties arise from the practice of *acceptance*, both for advice and evaluation. Suppose that Advisor has just told Agent that in order to distribute water equally between the shafts he must block B. But in processing this new information Agent becomes confused, and replies, 'Then I ought to block A.' At this point Physicist enters. Knowing that the miners are in A, he responds,

(3) Yes, that's true—you ought to block A—but not for the reason you think.

Or, in the case of distant evaluation (to press the problem against news-sensitive contextualism):

(4) Yes, that's true—he ought to block A—but not for the reason he thinks.

This is problematic for contextualism, because Physicist knows that the proposition which (according to contextualism) Agent asserted, namely that his blocking A is best given the information available to him, is *false*. Contextualism seems unable to identify a proposition that Agent asserted and with which Physicist agrees.

Finally, consider the case in which according to contextualism, Agent's claim that he ought to block neither shaft is (given his context) actually true. One might reasonably suppose that Physicist could acknowledge this. But it would definitely be odd for him to say,

(5) ??Yes, that's true, but you ought to block shaft A.

This can be handled by news-sensitive contextualism, again. But the corresponding claim in evaluation (eavesdropping) is equally unacceptable, and problematic for any form of contextualism:

(6) ??Yes, that's true, but he ought to block shaft A.

Apparently, the conjuncts of these claims contradict each other, yet contextualism seems to tell us that both can be true together. Relativism, however, respects the intuition of contradiction, as it tells us that the truth values of both conjuncts are determined by the same context of assessment.

Taken together, the phenomena of denial, acceptance, and contradiction of ought-claims suggest that contextualism cannot be the right account of our actual use of 'ought'. Moreover, similar examples can easily be constructed where what is assessed is Agent's belief rather than his utterance, for example replacing (1) and (4) with:

(7) You think that you ought to block neither, but really what you ought to do is to block A.

(8) Agent believes that he ought to block neither, and that is true, but not for the reason he thinks.

We agree with K&M that all these cases seem to undermine contextualism and suggest relativism.¹⁶ But the inference to the truth of relativism requires a significant assumption: that the acceptance, rejection, and semantic assessment of people's claims and judgments is always the acceptance, rejection, or assessment of the same proposition that the person asserted or accepted. A number of common linguistic practices cast significant doubt on this assumption. Consider the following:

(9) I think that Sally stole the money;

(10) ??Yes, that's true, you do/No, that's false, you don't;

(11) Yes, that's true, she did/No, that's false, she didn't.

(12) Tom is driving a rocket in this race;

(13) ??No, that's false, it's just a very fast car;

(14) Yes, that's true, he has the fastest car here.

It is natural to respond to utterances of (9) and (12) by offering assessments of propositions other than those that (on an orthodox account) they literally assert. Of course, these cases are dissimilar to the cases above involving ought-judgments, in that utterances of (9) and (12) plausibly also implicate or in some other way express the propositions that we assess. But they do suggest that the most salient proposition for assessment is not always the proposition that was asserted.

What then does determine salience for assessment? A plausible general *principle of assessment* is that salience is governed by relevance to conversational interests. In the case of (9), for example, what is of primary conversational relevance is the proposition that the speaker reports himself to believe, and not the reported fact of his believing it. It is usually the case that the proposition that is most relevant for assessment when responding to somebody's claim is the original proposition he asserted, but it is not always the case, and in Section 1 we found reason to think that it is often not the case with ought-judgments in particular.¹⁷

¹⁶ Similar arguments have been directed also against contextualism about epistemic modals (Egan, Hawthorne & Weatherson 2005; Egan 2007; MacFarlane ms) and taste predicates (Lasnik 2005).

¹⁷ The information-relativity of 'ought' is significantly like the information-relativity of epistemic modals, but also importantly like the *standard*-relativity of 'ought', as we will see in the next section. By contrast, it is quite unlike the context-relativity of typical indexicals. For example, 'I' refers to the speaker, but propositions asserted by 'I am X' typically continue to have relevance as contexts change.

We've already observed that the primary interest motivating deliberation and advice is not to reach a true answer to a predetermined ought-question, but the promotion and protection of certain values. Ought-propositions are ordinarily relevant to this purpose only when they are based on information that is the best available. For this reason we allow our conversational context to update automatically to include any new information that becomes available, without comment or protest. Consequently, an ought-proposition relating to an inferior informational base is just not relevant in the contexts of deliberation and advice. Ought-propositions relativized to supplanted contexts become conversationally moot, while ought-propositions relativized to the present and best available context acquire primary conversational relevance. For example, on Physicist's entry into Agent's conversation, the context shifts abruptly (on news-insensitive contextualism) to include his additional information. This makes the question of the truth or falsity of the ought-proposition that Agent originally asserted conversationally moot. For Physicist or Agent now to assess this original proposition—*while lives are at stake*—would manifest a perverse fixation on truth for truth's sake.

When the informational context improves in this way, expressing assent or dissent with the *original* ought-propositions that were asserted or accepted is pragmatically pointless. But there are still *relevant* ought-propositions, which are related to what the agent asserted or accepted in the following way: they are the propositions that the agent's utterance or judgment would have expressed relative to the improved context. Given our general principle of assessment, contextualism predicts that in contexts of deliberation and advice we assess previous ought-judgments *as if* they had been made relative to the presently best available information, and we evaluate the relevant propositions rather than the original propositions.¹⁸ This does not mean that context-insensitive assessments involve misunderstandings of the original claims, but simply that our communicative practices are governed by conversational interests.¹⁹

Since ought-judgments made relative to the best available information provide agents with the best available bases for decision, these judgments function pragmatically as recommendations. Sometimes recommendations need to be endorsed or rejected as new information becomes available; even if Physicist does not reject any

¹⁸ Similar conversational dynamics have implicitly been endorsed by philosophers who have noted that ought-claims involve implicit comparison classes (e.g. Sloman 1970; Jackson 1985). Here is Jackson: 'I say "It ought to be that Lucretia used less painful poisons". You retort "Oh no, it ought to be that Lucretia used painless poisons". I then retort "Oh no, it ought to be that Lucretia used political means rather than poison to obtain her ends". You then retort "Oh no, it ought to be that Lucretia never existed at all". I then retort "Oh no, it ought to be that Lucretia existed but made people happy". You then retort "Oh no, it ought to be that everyone was already happy so that there was no question of her making them happy"... Each retort seems a fair one, how so? What is happening is that the set of alternatives to which the 'ought' is relative is being implicitly changed at each stage of the conversation.'

¹⁹ In defending what we would classify as a contextualist account of epistemic modals, Kent Bach suggests that in assessing a prior epistemic judgment made correctly from an inferior perspective as false, we have to be mistaken about what proposition was accepted in that earlier judgment.

proposition that Agent accepts, he still rejects Agent's *decision*.²⁰ (Admittedly, mere evaluations or 'eavesdropping' cases have no direct relevance for decisions or advice, but they typically seem to be made in the *mode* of advice, or as simulated advice). A convenient way of endorsing or rejecting recommendations is to express agreement or disagreement with the ought-claims by which those recommendations were expressed. This is how the contextualist should understand (1)-(8).²¹

Contextualism can therefore accommodate the relativists' observation that assessments of ought-claims from positions of superior information are typically context-(of-utterance)-insensitive. Whereas the relativist diagnoses this as the context of assessment determining the truth-conditions of the (original) proposition that is assessed, we diagnose it as the context of assessment determining which (relevant) proposition is assessed.²² Well-motivated pragmatic considerations provide the contextualist with an explanation of why ought-judgments in practical reasoning, advice, and evaluation behave *as if* relativism were correct. Which diagnosis should be preferred? It might be thought that contextualism's dependence on this complex pragmatics puts it at a disadvantage, and favours the apparently simpler relativist account. This is not the case, because some of the phenomena seem rather to favour contextualism over relativism.

If relativism explains context-insensitive assessments directly by the semantics of 'ought' and 'true' (and 'false') while contextualism explains them rather by appeal also to pragmatic features of deliberation and advice, then these rival accounts will yield divergent predictions about assessments in cases where different pragmatic considerations are operative. On relativism, we might expect similarly context-insensitive assessments in the absence of these special circumstances, while on contextualism we might expect rather context-sensitive assessments. It is the contextualist prediction that is borne out by our practices.

One such kind of case involves *hindsight assessments*. Suppose that Agent acts in accordance with his judgment that he ought to block neither shaft, saving nine of the ten

²⁰ We think that this point receives some support from consideration of Physicist's choice between the following two responses to Agent: (a) 'No, don't do that!' (b) '??'No, that's false!' We find (a) natural, but (b) rather odd. However, (b) is arguably not *false*, and there doesn't seem to be any similar oddity about expressing semantic *agreement* (with 'true'), so this point needs care.

²¹ This supports what we might call a 'quasi-expressivist' interpretation of this discourse, explaining its expressivist-looking features with purely realist resources (turning the tables on the 'quasi-realism' of Blackburn 1993). The fact that they do actually express propositions allows us to express this attitudinal disagreement through evaluations in terms of truth and falsity, without having to embrace an expressivist semantics.

²² A contextualist position of this kind might seem self-undermining: isn't allowing context-insensitive assessments as 'true' incompatible with the contextualist's signature claim of context-sensitive *truth-conditions*? It is not, because by *truth-conditions* we mean the technical notion of the conditions of satisfaction of a proposition, which doesn't commit us to any stance on the meaning or use of 'true' in English.

miners. In the subsequent debriefing, despite our having learned that the miners were all in shaft A and would all have been saved if it had been blocked, it seems perfectly appropriate to say,

(15) Agent was quite right; blocking neither shaft was indeed what he ought to have done. Any other action would have been totally irresponsible.

Here, (15) is clearly relativized to the information Agent had at the time of his decision, just as contextualism predicts. With hindsight assessments we are no longer interested in guiding decision by the best information available, but rather in evaluating the actual decision in light of the information that the agent possessed. Here the relevant proposition for assessment is the original proposition that concerned the agent. But relativism may seem to be committed to rejecting (15) as false, since relative to the information available to the assessors, blocking shaft A was clearly the better option.

Relativism also seems to do worse than contextualism in accounting for *assessments from contexts of inferior information*:

Watching on closed-circuit television, *Observer* knows that the miners are trapped in one of the shafts but not which, and knows that blocking one will completely submerge the other. She also knows that Agent has just been told by *Physicist* where the miners are, but not where *Physicist* said they were, or what Agent concluded from this.

Given only *Observer*'s information, the best course of action would be to block neither shaft. But contrary to what relativism seems to imply, *Observer* could plausibly think that:

(16) If Agent concluded that he ought to block neither shaft, he was wrong.

This is what contextualism would predict: given Agent's information, it would be false that he ought to block neither. Superior information allows a better decision, and therefore we generally defer to the judgments of those who have it, so long as we trust that they aren't making poor use of it. As our account of contextualist pragmatics predicts, in general our assessments of ought-claims are context-insensitive only insofar as the context of assessment is informationally superior to the context of utterance.

It is important to note that these cases of context-sensitive assessments do not settle the issue in favour of contextualism, because there is a solution available to the relativist to explain why in these cases we behave *as if* contextualism were correct. Relativism need not identify the 'context of assessment' inflexibly as always constituted by the information available to the assessor, even if it has often been presented that way. Just as the contextualist's 'context of utterance' is not to be identified simply as constituted by the information available to the utterer at the time of utterance, but rather as

determined by the informational base *intended* by the utterer, so too can the 'context of assessment' be identified as determined by the informational base *intended* by the assessor—which can be the information available to people other than the assessor.²³ Observer can assess the truth of an ought-claim relative to his own information, or relative to Physicist's information, or relative to any other set of information. However, the point to emphasize here is that in order for relativism to achieve sufficient flexibility to accommodate these cases, it will presumably need to invoke pragmatics as complex as the pragmatics we have offered for contextualism, which therefore do not disadvantage contextualism at all.

We have, perhaps, a stalemate—and it might even be suspected that contextualism and relativism are mere notational variants, differing predominantly on the theoretical role that they assign to the technical term 'proposition'. Any argument for one over the other may have to turn on such abstruse considerations; contextualism's interpretation of assessment as directed at relevant rather than original propositions may be a cost that favors relativism (although we have argued that it is well-motivated and independently plausible), but this is counterbalanced by the cost incurred by relativism in embracing the radical notions of relative propositions and relative truth. But the significant conclusion here is that with regard to the information-relativity of normative ought-claims, contextualism is able to accommodate the data produced by its opponents. We now turn to consider the suggested standard-relativity of normative claims, where similar issues await.

Relativity to Standards

Many philosophers have been drawn to the view that normative 'oughts' are relativized to standards (norms, ends, etc.)²⁴ Relativity of this kind seems to offer comparatively unproblematic explanations of a variety of things, including how there can be such things as normative facts, why the normative domain divides into the moral, prudential, etc., why moral beliefs diverge between cultures, and why moral disagreement persists among well-informed competent judges. We believe that moral claims are indeed standard-relative and that contextualism is the correct semantic treatment of this relativity; every complete moral proposition includes a relation to a moral standard. Here also contextualism has recently received criticism, which is very similar to the

²³ Such a flexible version of relativism deflects another contextualist objection, presented against relativism about epistemic modals by Kent Bach (ms). This objection holds that it is clearly the case that *some* modal claims express relativized propositions, because sometimes we explicitly relativize them (e.g. 'Given what she knows, the keys might be in the door'; compare: 'Given what Agent knows, he ought to block neither shaft')—so the relativist seeks to introduce an unnecessary complication into the semantics of modals. But the relativist can offer an interpretation of 'given that *p*' as forcing a particular context of assessment.

²⁴ They include Harman 1975, 1996, Wong 1984, Dreier 1990, Copp 1995, and Finlay 2008.

criticism of contextualist treatments of information-relativity addressed above.²⁵ However, whereas information-relativity is alleged to tell in favour of semantic relativism, here the evidence is rather alleged to tell in favour of semantic invariantism; moral judgments make no reference to moral standards in their content, and their truth conditions are not standard-relative. In this section, we argue that contextualism has the resources to answer these objections. Some of these resources we have already observed, while others are peculiar to the present issue.

The criticism that we will discuss is directed at the contextualist treatment of fundamental moral disagreement, assessments of moral claims and beliefs, and propositional attitude reports. For illustration, consider Huckleberry Finn's belief that he ought to tell on Miss Watson's fugitive slave, Jim. Huck himself is wrestling with that belief throughout Mark Twain's novel, and as readers we reject it; it is not the case that Huck ought to tell on Jim. Since contextualism takes Huck's belief to consist in the acceptance of a standard-relative proposition, it accepts the possibility that the proposition that Huck accepted and the proposition that we reject are different: whereas Huck accepted the proposition that he ought-relative-to-standard-Y to tell on Jim, we reject the proposition that he-ought-relative-to-standard-Z to tell on Jim, where Y and Z are different standards.

If the standards are indeed different, this gives rise to a series of problems. The first is that it now seems that contrary to appearances, we are not really in disagreement with Huck: what he accepts is not what we reject.²⁶ This problem for the contextualist treatment of standard-relativity parallels the problem of integrating deliberation and advice, raised by K&M against the contextualist treatment of information-relativity. Moral judgment, like deliberation, aims at determining what ought to be done. Moral dissent, like advice, aims at correcting judgments about what ought to be done. So contextualism seems to have a problem: it has the result that moral 'dissent' addresses the answer to a question ('What ought-relative-to-Z one to do?') different from the one addressed by the original belief ('What ought-relative-to-Y one to do?')—and therefore it is hard to see how it is genuine dissent at all.

As with the problem of advice from unexpected sources, the contextualist can choose either to deny or to accept that the rival claims are relativized differently. Denying it immediately avoids the problems that we discuss here, but we shall not explore this option. We accept that there are genuinely some *fundamental* moral disagreements, where conflicting judgments are related to different standards.²⁷

²⁵ Weatherson 2008; Schroeder ms; Streiffer 2003.

²⁶ Lyons 1982: 210; Streiffer 2003: 8; Huemer 2005: 50.

²⁷ Appeal to common standards does have a cost: it neutralizes contextualism's supposed advantage in explaining the extent and persistence of moral divergence (Streiffer 2003: 14). However, we think contextualism is also (and better) motivated by other considerations.

Assuming the Huckleberry Finn case is of this kind, there are two independently plausible and complementary ways in which the contextualist can explain our sense that we disagree with Huck. One explanation is that we erroneously, but perhaps warrantedly, assume that there is a shared standard in this case, or in moral disagreements in general.²⁸ Robert Streiffer has argued that this defense of contextualism fails; in light of the ubiquity of moral disagreement it is highly implausible that we would assume common standards (2003: 14-15). But such an assumption could be reasonable even in the face of extensive and intractable disagreement; moral standards might reasonably be thought to be highly abstract and difficult to apply—perhaps even indeterminate—as would be the case if Kantians or Utilitarians were correct about the principles of morality. Moreover, if our sense of disagreement depends on the assumption of a common standard, that could explain why it is less clear when the standards are *strikingly* different. If we consider the moral beliefs of (e.g.) a New Guinean headhunter prior to ‘civilized’ contact instead of the moral beliefs of a 19th century American (like Huck), it is arguably much less obvious that we have an intuition that those moral beliefs contradict our own.

The other explanation available to contextualists parallels the explanation of the role of ought-judgments in deliberation and advice in the previous section. There we observed that the truth of any particular ought-proposition has only derivative importance to deliberators and advisors, which it owes to a more fundamental concern motivating them: the promotion of certain values.²⁹ This is plausibly also true in the case of fundamental moral disagreements like the one we (supposedly) have with Huckleberry Finn. However, while deliberation and advice are integrated by a common concern with promoting some value (like the preservation of human life), we must look elsewhere to find what integrates Huck’s and our own moral beliefs in such a way that there is something that can be called *disagreement* here. In fundamental moral disagreement the interlocutor is *contending* against his opponent, not trying to *help* him reach some common goal.³⁰

What is more fundamental in moral dispute than the truth of any particular ‘ought’ propositions is the issue of what *to do*. It is a well-worn point in metaethics that conflicts in moral attitudes need not always involve their having contradictory contents. Expressivists from Charles Stevenson (1944) to Allan Gibbard (2003) have argued that moral conflict consists in a disagreement of ‘attitude’ rather than a disagreement of belief. Contextualism can observe a conflict of this kind between ourselves and Huck.

²⁸ Finlay 2008b: 356-7; Wong 1984: 79.

²⁹ Contrast Streiffer (2003: 6), who assumes that the contextualist must say that the goal of each of the disagreeing parties is to establish the truth of its own position.

³⁰ It is possible to interpret moral disagreement in this more cooperative way (as, for example, factual disagreement over which are the objectively *best* or *correct* standards), but the possibility of such values undermines the rationale for standard-relativity.

Our moral judgments are relativized to the moral standards to which we *subscribe*. To subscribe to a moral standard is (inter alia) to have a strong preference that people conform to it in their conduct. Huck's belief and our own might not conflict by themselves, but in combination with subscription to conflicting standards they place us in conflict with regard to the practical matter of what to do. In virtue of his subscription to standard Y, Huck's moral belief commits him to support telling on fugitive slaves.³¹ In virtue of our subscription to standard Z, our moral belief commits us to oppose telling on fugitive slaves. These noncontradictory moral beliefs therefore precipitate a conflict of attitudes toward Huck's action.³²

This response only shows, however, that contextualism is able to accommodate a nonpropositional kind of disagreement between people who subscribe to different moral standards. It doesn't yet vindicate contextualism as an account of actual moral disagreement. Further disagreement problems emerge from propositional attitude reports and semantic assessments of moral claims and beliefs as true or false.³³ The most obvious problem is that contextualism would seem to prevent us from expressing our disagreement with Huck by saying that his belief is false, and force us to say that it is true. This is the problem of context-insensitive assessments again: our semantic assessments of people's moral beliefs are normally insensitive to the standards to which they subscribe.

One variation on this problem is raised by Brian Weatherson (ms), who suggests that only moral invariantism can account for how we attribute moral belief and moral knowledge. Weatherson points out that in our use of a sentence like

(17) Huck believed that he ought to tell on Jim

the relevant standard seems to be the one to which the believer himself subscribes, and not the one to which we subscribe (which he might never have contemplated).³⁴ By itself this is compatible with our liberal contextualism. But we can also assume (i) that Huck applies his moral standard correctly, so that according to contextualism the proposition that Huck believes is *true*, and (ii) that he is epistemically justified in believing it. He would therefore seem to meet the criteria for *knowing* it. But consider

(18) Huck knew that he ought to tell on Jim.

³¹ A commitment that is overcome by his humanity, as the reader knows.

³² Streiffer (2003) claims that contextualists have to reject ordinary intuitions about when our moral claims are and are not contradictory. We disagree; the contextualist can accommodate intuitions about when our moral claims are and are not in *practical conflict*, and we doubt that ordinary intuitions can be trusted to reliably discriminate this from genuine contradiction.

³³ Weatherson 2008; Schroeder ms.

³⁴ Weatherson's target sentence involves 'wrong' rather than 'ought' (and Jefferson Davis rather than Huck Finn): 'Davis believed that helping fugitive slaves was wrong'.

We agree with Weatherson that (18) is (usually) appropriately evaluated as false.³⁵ Huck couldn't have *known* that he ought to tell on Jim, because that is false. By itself, this also is compatible with our liberal contextualism, which allows us to interpret (18) as implicitly relativized to our moral standard, Z. But Weatherson claims that this combination of accepting (17) as true and rejecting (18) as false is a problem for contextualism, because it seems to require that normative terms behave differently (pick out different standards) in knowledge reports than they do in belief reports. Invariantism doesn't face this problem; Huck and we apply our different standards merely to lead us to conflicting judgments about what he ought (nonrelatively) to do.

Why is it a problem if belief and knowledge reports behave differently? Weatherson's main reason seems to be that belief and knowledge reports don't diverge like this in other cases, so moral terms would be anomalous if contextualism were true. We think this is mistaken; epistemic modals display the same pattern. When Physicist learns about Agent's predicament in the miners case, he can report that:

(19) Agent believes that the miners might be in shaft B.

Plausibly, the relevant context (informational base) here is that of Agent. But since Physicist knows that the miners are in A, he can also plausibly reject the following:

(20) Agent knows that the miners might be in shaft B.

His ground for doing so is that relative to the information available to him, it is not the case that the miners might be in B. This shows that the pattern isn't as general as Weatherson claims. Moreover, invariantism is patently false for this epistemic case.

The problem can also be motivated in the following way. What an agent S believes is a matter of the attitudes S has towards propositions. Hence it is natural to understand belief reports as involving propositions relativized to S's standards, since these are the propositions towards which S has attitudes of this kind. What S knows is a subset of what S believes. So it is natural to understand knowledge reports as concerning the propositions S believes. But then, to quote Weatherson, it 'would be a real shock if some term *t* behaved quite differently in belief and knowledge reports'.

Mark Schroeder (ms) raises a second variation on this problem, which he calls "the general problem with attitude-ascriptions" for contextualism. Consider

(21) Huck believed that he ought to tell on Jim.

(22) It is not the case that Huck ought to tell on Jim.

(23) Huck believed something that is not true.

³⁵ Unlike Weatherson, we *also* think that it can be appropriately evaluated as true, and heard as relating to the norms to which Huck and his society ascribed. See note 37.

The inference *seems* good; (23) seems to follow from (21) and (22), and both premises seem true. The dilemma facing contextualism is this: The 'ought' in (22) must relate to our standard Z, if the premise is to be true. If the 'ought' in (21) also relates to our standard Z then the inference seems valid, but (21) would be false: Huck didn't believe that he ought-relative-to-Z to tell on Jim. If it relates rather to Huck's standard Y, (21) is true but the inference is invalid, contrary to appearances.

These problems all involve insensitive semantic assessments, either explicitly in terms of truth and falsehood or implicitly in terms of knowledge, and we believe they can be solved by following the contextualist playbook for insensitive assessments proposed in Section 2. We simply need to extend the playbook from *information-insensitivity* to *standard-insensitivity*. As relativists have shown, some appropriate assessments of ought-judgments are insensitive to what the judge tried to get right, namely to identify the best action *given the available evidence*. We argued in defense of contextualism that assessment of an agent's claim or belief is concerned with a *relevant* proposition, related to what was asserted or accepted, which is not necessarily the *original* proposition asserted or accepted.

In the case of information-relativity, we observed that certain contexts (involving better information) were *privileged* because they put us in better position to promote our values, and that this supported a pragmatic explanation of context-insensitive treatment of others' normative claims (in advice and evaluation). Contexts involving lesser information are conversationally moot, as is the truth or falsity of propositions relativized to them. Conversational purposes are then better served by evaluating relevant propositions distinct from those that were originally asserted: the propositions that the uttered sentence would have expressed relative to the privileged context and information.

When we pay attention to the pragmatics of moral thought and discourse, we likewise find that some contexts (standards) are privileged: those standards constitutive of our *values*. Part of what it is to engage in a *moral* (or categorically normative) practice is to subscribe to a particular standard, to the exclusion of any rival, as determinative of what to do, both for one's own conduct and others'. Characteristically, what fundamentally matters to us when we make moral judgments is agents' conformity with the moral standards to which we ourselves subscribe. Our interest in the truth of moral (standard-relative) propositions is therefore derivative upon this concern, and the truth of propositions relativized to *other* standards is irrelevant to us. To address the question of whether Huck was right in thinking that not telling on Jim violates Huck's standards, *when the issue is how to act in circumstances like Huck's*, is therefore a perverse fixation on truth for truth's sake, in neglect of what is relevant to the purpose of moral thinking and moral discourse.

What would be relevant to our concerns, however, is to assess the moral judgments of others like Huck *as if* they had been made in relation to *our* standards and to express agreement or disagreement with the resulting ought-propositions, which *are* relevant to us. We observed in the last section that ought-judgments relativized to what is at the time the best available information function pragmatically as *recommendations*, and for the purposes of advice it is convenient to treat them as such, irrespective of the propositions they originally expressed. Similarly, ought-judgments relativized to the standard to which the judge or speaker subscribes function pragmatically as *imperatives*, and it serves the characteristic purposes of moral thought and discourse best to treat them as such. Since these purposes involve insistence on actions conforming to *our* standards, we can assess them simply as imperatives that we either accept or reject.³⁶ This is why we assess Huck's belief that he ought to tell on Jim as false rather than true.

It may seem odd to suggest that our assessment of Huck's belief is to be understood as an *imperative*, since we obviously cannot address this imperative to Huck or affect his actions. The situation here is parallel to the case of eavesdropping or evaluation from a distance, which we suggested takes the form of simulated advice. Likewise, we suggest that moral evaluation from a distance is expressed in the form of simulated demand.

We have claimed that semantically assessing Huck's belief can be assessing the proposition relevantly related to the original proposition that Huck accepted, a proposition relating to our standards rather than his. This solves the problem of context-insensitive assessments, and suggests the solution to Schroeder's problem. The inference that Schroeder challenges contextualists to validate should be understood as having the following form:

(24) Huck believed that he ought-relative-to-Y to tell on Jim.

³⁶ The dynamics of privileged propositions can be illustrated by the following kind of case. A: 'The killer ought to have used a silencer'; B: 'No, certainly not. He ought not to have killed at all'; A: 'No, of course.' As we interpret this exchange, speaker A is offering an instrumental evaluation, relativized to the killer's goals. Speaker B understands this, but substitutes an alternative (moral) standard as normatively trumping the killer's goals, thereby offering a context-insensitive assessment of A's claim. A accepts the retroactive change in context, since the new standard is more highly valued than the original one (Finlay 2006). The dynamics of conversations involving fundamental moral disagreement are obviously different, since the interlocutors have conflicting values and conversational purposes. See Finlay 2004, 2005, 2008; see also Barker 2000; Harman 1996.

An alternative proposal for accommodating these kinds of concessions suggests that ought-claims always identify something as the best out of a specific set of alternatives—a set that can always be expanded, sometimes introducing a better option, like *not* killing (Jackson 1985; see our note 18). But this doesn't work in cases where the change in context alters the ranking of the original options; for example it may be morally preferable that the killer failed to achieve his goals (by not using a silencer), given only the options {he succeeds, he fails}.

- (25) It is not the case that Huck ought-relative-to-Z to tell on Jim.
- (26) There is a proposition, p, such that Huck believed p and the relevant proposition for assessment that is related to p is not true. ('Huck believed something that is not true.')

As we have argued, the proposition that *Huck ought-relative-to-Z to tell on Jim* is relevantly related to the original proposition that *Huck ought-relative-to-Y to tell on Jim*. Given this, (26) follows from (24) and (25), as desired.

We now have a straightforward answer to Weatherson's challenge to explain why 'ought' behaves differently in belief and knowledge ascriptions. First, in attributing beliefs we are interested in understanding the subject's attitudinal state. Our contextualist proposes that in the case of moral ought-beliefs that state typically involves two attitudes: (a) an ordinary belief in a standard-relativized proposition, and (b) a subscribing/demanding attitude towards the standard involved. (This view has the virtue of offering an explanation for the intractable metaethical debate between cognitivists and expressivists: both are right). We therefore interpret (17) as the ascription to Huck of the compound attitude, subscribing to a slavery-tolerant moral standard Y, and believing the proposition that this standard calls for telling on Jim. Second, in attributing knowledge we are not only attributing belief but also engaging in semantic assessment; knowledge implies truth.³⁷ As we have already argued, assessing Huck's belief is assessing a proposition relevantly related to the proposition that Huck accepts, in this case the proposition that Huck ought-relative-to-Z to tell on Jim. Since we take that proposition to be false, we reject the attribution of knowledge to Huck in (18).³⁸

³⁷ As contextualists, we recognize that there are attributions of moral knowledge made with 'anthropological' rather than moral conversational purposes. For instance, it might make perfect sense to say, by way of retelling the story of Huckleberry Finn, 'Huck knew that he ought to tell on Jim, but his humanity weakened his will'. The reference to the standards of Huck and his social environment would be obvious enough, and no endorsement of those standards implied.

³⁸ Weatherson does have a further objection to differential relativizations of 'ought' in contexts of belief and knowledge reports: that it would implausibly allow an interpretation of (A) in which it isn't contradictory.

(A) S believes that he ought to ϕ . Indeed S knows it. But S doesn't know that he ought to ϕ . The reason that differential relativizations would make (A) non-contradictory 'in principle' is presumably that it would make it possible for the two knowledge reports to relate to different ought-propositions. The first knowledge report seems to ascribe knowledge of the proposition made salient in the previous *belief* report (and referred to by 'it'), a proposition thus related to S's moral standards; the second knowledge report ascribes knowledge of a proposition identified by an 'ought' which itself occurs embedded in that *knowledge* report, hence a proposition related to the speaker's standards. But this reading of (A) doesn't seem to be available.

This problem is dissolved by our contextualist account of context-insensitive assessments. Although the first knowledge report is concerned with S's belief, the assessment implied in that knowledge report is, like all such assessments, concerned with a proposition *relevantly* related to the proposition believed by S, viz. the corresponding proposition related to *our* standards. This gives the first knowledge report

We conclude that with the help of a rich understanding of moral pragmatics, the contextualist interpretation of the standard-relativity of moral discourse can be successfully defended against the objections recently raised against it—just as pragmatics came to the rescue of the contextualist treatment of information-relativity. It may again be asked, however, whether this pragmatic account isn't inferior to the simpler explanation offered by contextualism's rival—in this case, invariantism. The pragmatic account we have sketched allows contextualism to explain why it might seem as if invariantism about moral discourse were true. But then one might think that invariantism is the preferable theory, because simpler.

In our view, this advantage of invariantism is illusory. First, it is highly plausible that 'ought' (with other normative modals) sometimes *does* have an argument place for a standard, which is sometimes provided by context. We talk about what ought to be done relative to the rules of war, and what ought to be done relative to the rules of etiquette, etc.³⁹ Our contextualist account ascribes the moral 'ought' the same logical form as it has in these other uses, while invariantism seems to need to postulate ambiguity in logical form. Contextualism is therefore the simpler *semantic* theory. Second, we in fact do relativize moral claims when we take a nonmoral interest in reporting them (as anthropologists, for example), which shows that the invariantist behavior of this discourse is contingent on certain pragmatic considerations. Third, we suggest that contextualism provides a plausible explanation of what we are talking about when we make moral claims, why moral disagreements might be particularly intractable, how there can be such a thing as moral value, and why it reliably motivates and matters to us—while invariantism has no plausible explanation to offer.

These are claims that need a thorough defence, of course. But the following considerations alone show that invariantism gains no support from the phenomena dealt with here. The problems raised for standard-relative contextualism all depend on the fact that there are semantic assessments of moral ought-judgments that are

exactly the same content as the second; hence the contradiction. The contradiction could be avoided if the first knowledge report could be made with an anthropological interest and the second with a moralistic. But such shifts of conversational interest do not occur between conjuncts without significant contextual markers (change of voice, stressing the second 'know', etc).

Our solution also dissolves Streiffer's dilemma (2003: 9-12): Either contextualism holds that moral utterances are always relativized to the speaker's standards, or it holds that they can more flexibly be relativized to any contextually salient standard. The difficulty with the first horn is that moral sentences are not relativized to the speaker's standards when embedded in belief attributions. The difficulty with the second horn is that moral sentences are never relativized to any other standards except when embedded in belief attributions. But we now can explain why embracing this second horn is not (as Streiffer claims) ad hoc; in assertoric uses of a moral sentence or in semantic assessments, a speaker is fundamentally concerned with conformity of actions to her own standards, whereas in belief attributions (and anthropological attributions of knowledge) her concern is with the subject's attitudinal state.

³⁹ The standard semantics for modals widely accepted in linguistics (see Sloman 1970; Kratzer 1977) recognizes such an argument-place, which in deontic uses is filled by some system of norms.

insensitive to the standards endorsed by the judge. Traditionally this has been taken to support invariantism. But there are also assessments of ought-judgments that are insensitive to the informational base to which the judge was related, and invariantism is thoroughly implausible in that area. So context-insensitive assessments cannot *as such* favor invariantism. Furthermore, since contextualism invokes the very same sort of pragmatic considerations (in particular an appeal to *privileged* contexts) to account for both sorts of context-insensitive assessments, no significant extra cost is added to defend standard-relativity.

This concludes our support of contextualist treatments of the apparent information- and standard-relativity of normative and moral 'oughts'. We have argued that the objections against contextualism can be handled by accepting well-motivated views about the pragmatics of these discourses. We also suspect, although we have not argued for it, that many of our strategies generalize effectively for other normative terms and modal terms.⁴⁰

⁴⁰ We thank Niko Kolodny for sharing his and John MacFarlane's manuscript with us, and Mark Schroeder for helpful comments. Stephen Finlay's work on this paper was supported by funding from the Andrew Mellon Foundation/ACLS, and the USC Provost's office.

References

- Kent Bach (ms), 'Perspectives on possibilities: Contextualism, Relativism, or What?'
- Stephen J. Barker (2000), 'Is value content a component of conventional implicature?', *Analysis* 60: 268-79.
- Simon Blackburn (1993), *Essays in Quasi-Realism*. New York: Oxford University Press.
- Berit Brogaard (2008), 'Moral Contextualism and Moral Relativism', *Philosophical Quarterly* 58: 385-409.
- Herman Cappelen & Ernie Lepore (2005), *Insensitive Semantics*. New York: Oxford University Press.
- David Copp (1995), *Morality, Normativity, and Society*. New York: Oxford University Press.
- Jamie Dreier (1990), 'Internalism and Speaker Relativism', *Ethics* 101: 6-26.
- Andy Egan, John Hawthorne & Brian Weatherson (2005), 'Epistemic Modals in Context', in G. Preyer & P. Peter (eds.), *Contextualism in Philosophy*. Oxford: Oxford University Press.
- Andy Egan (2007), 'Epistemic Modals, Relativism, and Assertion', in *Philosophical Studies* 133: 1-22.
- Stephen Finlay (2004), 'The Conversational Practicality of Value Judgement', *Journal of Ethics* 8: 205-223.
- _____ (2005), 'Value and Implicature', *Philosophers' Imprint* 5 (4): 1-20.
- _____ (2006), 'The Reasons that Matter', *Australasian Journal of Philosophy* 84: 1-20.
- _____ (2008a), 'Oughts and Ends', *Philosophical Studies* (forthcoming).
- _____ (2008b), 'The Error in the Error Theory', *Australasian Journal of Philosophy* 86: 347-69.
- Kai von Fintel & Anthony Gillies (2008), 'Might Made Right', in Andy Egan & Brian Weatherson (eds.), *Epistemic Modality*. Oxford University Press.
- Allan Gibbard (2003), *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Gilbert Harman (1975), 'Moral Relativism Defended', *Philosophical Review* 85: 3-22.
- Gilbert Harman & Judith Jarvis Thomson (1996), *Moral Relativism and Moral Objectivity*. Cambridge, MA: Blackwell.
- John Horty (2006) 'Right actions in perspective', in *Modality Matters: Twenty-five Essays in Honour of Krister Segerberg*, H Lagerlund, S. Lindstrom, and R. Sliwinski (eds.), Uppsala Philosophical Studies, vol 53, Uppsala University Press
- Michael Huemer (2005), *Ethical Intuitionism*. Palgrave Macmillan.
- Frank Jackson (1985), 'On the Semantics and Logic of Obligation', *Mind* 94: 177-95.
- _____ (1991), 'Decision-Theoretic Consequentialism and the Nearest and Dearest Objection', *Ethics* 101: 461-82.
- Niko Kolodny & John MacFarlane (ms), 'Ought: Between Subjective and Objective'

- Angelika Kratzer (1977), 'What 'must' and 'can' must and can mean', *Linguistics and Philosophy* 1: 337-55.
- Peter Lasersohn (2005), 'Context Dependence, Disagreement, and Predicates of Personal Taste', *Linguistics and Philosophy* 28: 643-86.
- David Lyons (1982), 'Ethical Relativism and the Problem of Incoherence', reprinted in J. W. Meiland & M. Krausz (eds.), *Relativism: Cognitive and Moral*. Notre Dame, IN: University of Notre Dame Press.
- John MacFarlane (ms), 'Epistemic Modals are Assessment-Sensitive', accessed 11/3/08 at <http://johnmacfarlane.net/epistmod.pdf>
- G. E. Moore (1912), *Ethics*. Oxford: Oxford University Press.
- Mark Schroeder (ms) 'Hybrid Expressivism'
- Aaron Sloman (1970), "'Ought" and "Better"', *Mind* 79: 385-94.
- Charles L. Stevenson (1944), *Ethics and Language*. New Haven, CT: Yale University Press.
- Robert Streiffer (2003), *Moral Relativism and Reasons for Action*. New York: Routledge.
- Brian Weatherson (2008), 'Attitudes and Relativism', *Philosophical Perspectives* 22.
- David Wong (1984), *Moral Relativity*. Berkeley, CA: University of California Press.