

Metanormative Contextualism and Normative Uncertainty

JOHN PITTARD
Yale University
john.pittard@yale.edu

ALEX WORSNIP
New York University
alex.worsnip@nyu.edu

Forthcoming in *Mind*; please cite published version

Abstract. We offer a new argument in favor of metanormative contextualism, the thesis that the semantic value of a normative ‘ought’ claim of the form ‘*S* ought to Φ ’ depends on the value of one or more parameters whose values vary in a way that is determined by the context of utterance. The debate over this contextualist thesis has centered on cases that involve ‘ought’ claims made in the face of uncertainty regarding certain descriptive facts. Contextualists, relativists, and invariantists all have plausible ways of explaining these cases, and one could reasonably judge the debate between these views to be a stalemate. We argue that this stalemate can be broken by shifting focus to a case that involves normative uncertainty rather than descriptive uncertainty. While relativist and invariantist rivals of contextualism can give plausible accounts of the descriptive uncertainty cases, only contextualism can provide a plausible account of the normative uncertainty case.

Introduction

Jill says she ought to do action *A*, but Jack says that Jill ought to do some incompatible action *B* instead. These sound like inconsistent normative claims. But in some cases of this sort, cases where Jill is uncertain about something that Jack is not uncertain about, it seems that Jill and Jack both speak truly. The puzzle presented by such cases has recently received a number of competing solutions: relativist solutions, contextualist solutions, and non-relativist ‘invariantist’ solutions. Each of these approaches has resourceful defenders, and at present one could reasonably judge the debate to be a stalemate. In this paper, we argue that this stalemate can be broken if we focus on a type of case that differs in an important respect from the type of case that has been at the center of the discussion thus far: specifically, on a case involving *normative* uncertainty rather than *descriptive* uncertainty. Focusing on such a case reveals the inadequacy of the relativist and non-relativist invariantist proposals and allows one to mount a clear and compelling argument in favor of metanormative contextualism.

Normative uncertainty—and how one ought to respond to it—is in its own right the subject of lively current debate in normative theory. Yet cases of normative, as opposed to

descriptive, uncertainty have generally been passed over in the literature on contextualism. This is ironic since, as we will argue, such cases of normative uncertainty actually provide the key to securing the case for contextualism.

For the purposes of this paper, we stay largely neutral on how the details of a contextualist semantics for normative terms are best developed, and silent on how to address the various objections that have been leveled at contextualism. If our argument is sound, then some form of contextualism has to be correct; the details are challenges for future work.

1. The original descriptive uncertainty cases

Let us begin with a representative example of the sort of case of descriptive uncertainty that has been the focus of the existing debate. Consider the following case, adapted from (Jackson 1991, 462–3):¹

Sarah is an herbalist in training. While collecting herbal treatments in the woods, she encounters Sam, who has been bitten by a poisonous snake and who will die soon if not treated immediately. Sarah has three herbs at her disposal: herb A, herb B, and herb C. She knows that herb C will save Sam's life, but will leave him with mild hypersensitivity to the sun. She also knows that one of herbs A and B will completely cure Sam with no side effects, but that the other will kill him. Unfortunately, she doesn't remember which is which. 'Since I can't remember which of herbs A and B is the one that will cure Sam and which is the one that will kill him, I ought to give him a treatment of herb C,' she says to herself. The next day, she describes the episode to her herbology professor, who sighs and says, 'You ought to have administered only herb B. Now Sam will need to wear sunscreen wherever he goes.'

Sarah claims that she ought to administer herb C. Her professor claims that she ought to have administered only herb B. Because these are incompatible courses of action, the claims of Sarah and her professor appear on face to be incompatible.² Yet it seems that Sarah and her professor both speak truly when making these normative claims.

Metanormative contextualism offers a solution to this puzzle. According to metanormative contextualism (or, from here on, just 'contextualism'), the semantic value of a normative 'ought' claim of the form 'S ought to ϕ ' depends on the value of one or more

¹ For similar examples, see (Regan 1980, 265, n. 1; Parfit 2011, 1:159).

² Like others who have discussed cases like Sarah's, we take it for granted that Sarah's situation is *not* a deontic dilemma where it is true (from a single context of assessment and without any variation in the semantic value of 'ought') *both* that Sarah ought to give herb C *and* that she ought to give herb B alone. When we come to the case of normative uncertainty later, we will return to the possibility of deontic dilemmas (§3.3).

parameters whose values vary in a way that is determined by the context of utterance.^{3,4} The proposal is then that ‘ought’ has one semantic value in Sarah’s mouth, and another semantic value in her professor’s mouth. Hence, both claims can be true.

There are a range of competing contextualist proposals about how to make good on this general idea. On one proposal, there is one parameter for standards and a second parameter for information, so that the meaning of ‘*S* ought to ϕ ’ depends on which standards and information set are contextually relevant.⁵ Given this semantics, one way of interpreting the present example is to say that Sarah’s statement amounts to a claim about what action has the highest expected value in light of the information available to *her*, while the professor’s statement amounts to a claim about what action has the highest expected value given the information available to *him*. (Alternatively, the relevant standards might vary across the two statements if the professor’s statement amounts to a claim about which action has the highest *actual*—rather than expected—value.) To stave off potential misunderstanding, it is worth emphasizing that the contextually-relevant set of information need not be the information that is accessible to or known by the speaker in every case. The speaker’s interests and intentions could render some other information set contextually relevant.⁶

There are reasons for thinking that this simple proposal is not ultimately adequate.⁷ But we do not aim to take a stand here on the question of how exactly to develop the

³ There are also non-normative usages of ‘ought’, such as the ‘evaluative “ought”’, and the “ought” of expectation’. We think that the most plausible semantics will give a unified account of all normative usages of ‘ought’—both concerning how one ought to act and concerning what one ought to believe—and probably of non-normative usages too. However, here we focus on normative usages about how one ought to act. The contextualist semantics should also extend unproblematically to normative claims featuring other related terms (e.g., ‘should’ and ‘may’).

⁴ A contextualist semantics for ‘ought’ and related terms has recently been defended by, among others, (Björnsson and Finlay 2010; Brogaard 2008; Cariani, Kaufmann, and Kaufmann 2013; Dowell 2012 and 2013; Finlay 2014; Henning 2014; Wedgwood 2006 and forthcoming). Snedegar (2013) and Cariani (2013) defend a contrastivist semantics for ‘ought’, which for our purposes is a form of contextualism; their reasons for preferring the contrastivist variant do not bear on the kind of cases that will be the focus of our discussion. Strangely, while contextualism is non-standard amongst metaethicists, and explicit defenses of it in metaethics are mostly very recent, contextualism about modals such as ‘ought’ (including their deontic uses) has long been orthodox in philosophy of language, following the semantics of Kratzer (1981), which several of the above accounts draw and build upon. One forerunner of contextualism in the metaethical literature, however, is Dreier’s (1990) ‘speaker relativism’.

⁵ See, e.g., (Björnsson and Finlay 2010).

⁶ Cf., Dowell’s (2013) ‘flexible contextualism’; see also Kolodny & MacFarlane (2010: 122).

⁷ In particular, a catch-all notion of ‘information’ obscures two very distinct effects of information-changes, which may be best modelled separately in the semantics. To illustrate this, consider the orthodox Kratzer (1981) semantics, which includes two parameters: a ‘modal base’ and an ‘ordering source’. The modal base consists of a set of propositions, which restrict the worlds over which a modal like ‘ought’ quantifies (since the modal quantifies only over the worlds in which the propositions in the modal base are true). The ordering source consists of a set of norms, which rank these worlds by how well they satisfy the norms in question. So ‘*S* ought to ϕ ’ is true iff *S* ϕ ’s in all the highest-ranked worlds. Now, the difference in information between Sarah and her professor is best understood as resulting in a shift in the contextually relevant ordering source (cf. Dowell 2012). For Sarah’s claim, the ordering source is expected value given *her* information, while for the professor the ordering source would be expected value given *his* information (or

contextualist semantics. What all versions of contextualism have in common is that they posit different semantic values for different utterances of ‘ought’, in a way that is not explained by brute ambiguity but rather by a unified semantics with one or more contextually-determined parameters. Though this leaves a lot open, it is a distinctive enough claim for our purposes in this paper, and our argument is pitched at a level of generality that favors this general proposal over other competing proposals. We now turn to an explication of these competing proposals for dealing with cases of descriptive uncertainty like the one we have just given. This discussion will help make clear why our argument for metanormative contextualism, which appeals to cases of *normative* uncertainty, is stronger than arguments relying on cases like the one above.

2. Non-contextualist proposals for resolving the cases

2.1 *Soft invariantism*

The ‘invariantist’, as we will use the term, denies both that ‘ought’ is context-sensitive, as the contextualist holds, and that the truth of ‘ought’-claims is assessment-sensitive, as the relativist holds. Although this view may not be a popular one in semantics, we wish to stress that some kind of invariantism is still the dominant assumption in ethics and epistemology.⁸ So it is important for us to engage it in this paper.

Let us start by distinguishing two kinds of invariantist view. According to the *soft invariantist*, both Sarah and her professor speak truly, but this is because they use ‘ought’ in different *senses* rather than because ‘ought’ is context-sensitive; uses of ‘ought’ may

perhaps actual value). Although in one sense the ‘standards’—expected value—have remained fixed, one gets a different ordering of worlds depending on whose information is relevant in calculating such expected value.

On the other hand, there is a different category of information shifts that are best understood as taking effect on the modal base. Such cases involve information that bears on what options are possible or ‘live’ for the subject. Such information may include external constraints on the subject’s choice, but can also in certain contexts include the subject’s plans, intentions and future actions that, when held fixed, effectively rule out potential courses of action and thus restrict the range of ‘live’ options. For example, if Sheila is giving away \$1,000 to a charity and has narrowed her options down to charity A and charity B, Jim could truly say ‘You ought to give the money to charity B’ even if, on the standards that constitute the ordering source for Jim’s utterance, giving \$1,000 to charity C would be better than either A or B. This is because when we add the proposition that Sheila will not give to charity C to the modal base, we are restricted to worlds in which she does not give to charity C, and the best *remaining* worlds, given the same ordering source, might be those in which she gives to charity B. Depending on Jim’s purposes in speaking, he may wish either to restrict himself to the worlds in which Sheila does not give to charity C to express this ‘ought of the second-best’, or to talk about what Sheila *ideally* ought to do, not holding fixed her current plans and intentions.

In this paper, the first kind of information-sensitivity—that which takes effect on the ordering source—will generally be the one of relevance.

⁸ Many writers in the ethics and epistemology literatures assume that ‘ought’ is ambiguous between two ‘senses’: objective and subjective. See, e.g., the many references collected in Feldman (1988), including Brandt, Pollock, Goldman, Unger, Alston, Kornblith, Moser, and Kvanvig, as well as, more recently, e.g., Setiya (2004), Gibbard (2005), and Schroeder (2009). This is a ‘soft invariantist’ view, as characterized below, in contrast to a genuine contextualist view.

sometimes be ambiguous between these different senses.⁹ The *hard invariantist*, on the other hand, claims that Sarah and her professor really do speak inconsistently with each other, and that at least one of them speaks falsely, strictly speaking, though s/he may implicate something true in the course of so doing.

We begin with soft invariantism, according to which Sarah and her professor use ‘ought’ in different senses. An obvious place to start is with the common suggestion that ‘ought’ has an ‘objective’ sense and a ‘subjective’ sense. Perhaps Sarah uses the subjective ‘ought’, which concerns what she ought to do given her subjective position, and the professor uses the objective ‘ought’, which concerns what she ought to do from the point of view of all the facts.

Clearly this proposal has some affinities with the contextualist idea. However, the idea that there are just *two* ‘ought’s—objective and subjective—is subject to a serious, and in our view fatal, problem of regress.¹⁰ Suppose that we posit just two ‘ought’s: for example, one that pertains to what you ought to do given all the facts and another that pertains to what you ought to do given your doxastic attitudes (such as your credences and beliefs).¹¹ The latter ‘ought’ helps us to make sense of claims about what you ought to do *given your uncertainty about the facts*. But these two ‘ought’s are not enough to help us make sense of claims about what one ought to do given further layers of uncertainty. Perhaps, for example, what you ought to do *given your uncertainty about the facts* is to perform the action with maximal subjective expected utility. But clearly, you can be (justifiably) uncertain about which action satisfies this criterion, and sometimes you should not take the steps that would resolve all of this uncertainty. For example, circumstances might demand that you act quickly, and determining which action has the maximal subjective expected utility might involve lengthy expected utility calculations that cannot be completed before the moment of decision.

Naturally, we might want to speak about what you ought to do given this deeper level of uncertainty. We might say, for instance, that you ought to perform some action *A* even if you are certain that it is not the action with the highest subjective expected utility as long as (i) you believe that the expected utility of *A* is high, (ii) you do not believe of some alternative action that it will have a higher expected utility, and (iii) there are no alternative actions that are obviously worthy of consideration and that you would (upon initial consideration) believe to have a higher expected utility than *A*. Accounting for this claim requires that we posit an additional semantic value for ‘ought’; to accommodate this, the advocate of two senses of ‘ought’—objective and subjective—will either have to admit that at least *one* of those ‘ought’s is context-sensitive, or posit a third sense of ‘ought’.

⁹ Ambiguity and context-sensitivity should not be confused. An ambiguous term has multiple meanings, each of which needs its own semantics. A context-sensitive term, by contrast, has a single semantics, with one or more parameters which get supplied by the context.

¹⁰ For related discussions, see (Sepielli 2014; Kagan ms.; Wells ms.; MacAskill ms.).

¹¹ If it’s already occurring to you that there must be other ‘ought’s—like what you ought to do given your *evidence*—then this only serves to further illustrate our point. See Feldman (1988) for a nice overview of how the objective/subjective distinction gets drawn in various non-equivalent ways by different writers, such that the literature as a whole is dealing in more than two notions.

But that's not it: the problem will then iterate once more for that third 'ought'. For starters, whether an alternative action is 'obviously' worthy of consideration (in the way that clause (iii) refers to in the characterization of the third 'ought') itself depends on whether obviousness is understood in a more objective or more subjective way. For some purposes we might be interested in speaking about what you ought to do given a more objective construal of obviousness, and for other purposes we might want to speak of what you ought to do given what was and was not obvious *to you*. Additionally, one can sometimes be uncertain about the content of one's beliefs concerning the comparative expected utility of different actions, and we might want to speak about what you ought to do *given that you are uncertain whether you believe that the expected utility of A is high* (for example). Making sense of claims about what you ought to do given further layers of uncertainty (or given what does and does not obviously occur to you) will require positing further semantic values for 'ought'. One may accommodate this by positing additional senses, but the problem will reoccur for these additional senses, multiplying senses of 'ought' *ad infinitum*. It is not plausible that 'ought' is *ambiguous* between such countless readings. Contextualism avoids the need to posit such ambiguity: it posits one 'ought,' with one or more contextually-sensitive parameters. Ultimately, then, the best version of the objective and subjective multiple senses proposal collapses into contextualism, on which there are a range of 'ought's, some more subjective and some more objective (and perhaps on different dimensions of objectivity).

But there are alternative proposals for multiple senses of 'ought' that the soft invariantist can appeal to, and that are less easily dealt with. For example, one might claim that Sarah is using 'ought' in a moral sense, while her professor is using 'ought' in an *evaluative*, non-moral sense.¹² When I exclaim that one ought to be able to take a walk in the park without getting eaten up by mosquitoes, it seems that I am not making a claim about requirements (moral or otherwise) that pertain to any agent, but am rather making a claim about what states of affairs would be *best*. Similarly, perhaps the professor's statement should be understood as a claim about what would have been best rather than as a claim about what reasons Sarah had or what morality or rationality required of her.

Now, it seems to us that even this kind of soft invariantist will be forced to concede that at least one of the senses of 'ought' at play will require a contextualist semantics. Suppose that ingesting a certain amount of soil can also prevent the ill effects of snake venom, and that *unlike* herb B, soil does *not* have the property of neutralizing some of the mild positive health benefits that snake venom is known to have. So in an ideal scenario, snake bite victims would not be given herbs but would instead be given the right amount of soil. But suppose further that the amount of soil administered must be within a *very* narrow range to save the victim, a range whose values depend on the precise amount of venom in the blood, the patient's weight, the alkalinity of the soil, and many other factors that could not be reliably measured without specialized lab equipment, equipment that is *not* available in the woods. Ingesting a quantity of soil that is not in this narrow range will result in the certain death of the victim. If the professor knows all this, it still seems that he can truly affirm that Sarah ought to have administered herb B, even though he knows that there was

¹² See (Schroeder 2010) for a view that might inspire this kind of proposal.

some amount of dirt that, if administered by Sarah, would have resulted in a better outcome.

The soft invariantist response to the original case needs to explain how, in this modified case, the professor could still reasonably say that Sarah ought to have administered herb B even though he knows there is some action that would have been better. It seems to us that the soft invariantist must concede that evaluative ‘ought’ claims do not always express a view about what is best absolutely, but may sometimes express a view about what is best from among some set of options determined by conversational context. So there will be pressure to concede a contextualist semantics for ‘ought’ claims. Of course, the type of example that we would need to pressure the soft invariantist into contextualism would be different depending on how the soft invariantist proposes to disambiguate various ‘ought’ claims. But no matter what senses of ‘ought’ are appealed to, there will probably be a modified example that can be accommodated in a plausible way only by appealing to a contextualist semantics for at least one of these senses.

However, even if we are right that the soft invariantist will be pressured into accepting a contextualist semantics for ‘ought’ in at least one of its senses, the soft invariantist might respond to this pressure by allowing that the *non-normative* ‘ought’ is context-sensitive while denying that the normative ‘ought’ is context-sensitive. In this way, she could continue to appeal to multiple senses as a way of denying metanormative contextualism, which is a thesis about the normative use of ‘ought’ in particular.

2.2 *Hard invariantism*

According to the hard invariantist, there is just one sense of ‘ought’ in play in the kinds of cases that motivate contextualism, and it is not context-sensitive or assessment-sensitive. While this might seem like an extreme or rigid view, there are two maneuvers available to the hard invariantist to explain away the apparent variability of ‘ought’, which together can seem to adequately deal with the range of ways that we might develop the cases of descriptive uncertainty such as that of Sarah.

The first is to claim, drawing on the work of Jonathan Dancy (1977; 2000, chap. 3) and especially John Broome (1999; 2004; 2007; 2013, esp. ch. 3) that ‘ought’ can sometimes take wide scope over a conditional.¹³ Then, one can say that the following two claims are both true:

¹³ There is a considerable literature, somewhat orthogonal to our purposes here, on whether rational requirements are best-understood as having a wide-scope logical form. (See the aforementioned Broome and Dancy references, as well as e.g. Brunero (2010; 2012), and Way (2010; 2011) for the view that they are, and Schroeder (2004), Kolodny (2005) and Lord (2011; 2014) for the view that they are not.) As has only been explicitly appreciated fairly recently (cf. Dowell 2012; Silk 2014; Worsnip 2015a), this question is independent of the question of whether any particular class of ordinary utterances are best understood as *semantically expressing* wide-scope ‘ought’s, let alone whether this mitigates the need for a non-invariantist semantics. While some wide-scopers about rational requirements seem to suggest this semantic view (Broome 2013, Brunero 2010, Dancy 1977), there are powerful critiques of it going back all the way to (Hansson 1969; Fraassen 1972; Lewis 1973; Kratzer 1981), developed in the current literature in (Dowell 2012; Silk 2014; Worsnip 2015a). Still, this does not necessarily speak against the view that the fundamental requirements of rationality are wide-scope in their logical form; see Worsnip (2015a).

- (1) Sarah ought to administer herb B (rather than herb C).
- (2) Sarah ought [to make it true that] (if she is uncertain as to which of herb A and herb B is the optimal herb and which is the killer herb, then she administers herb C).

(1) and (2) are consistent, even on the assumption that the semantic value of ‘ought’ at issue in the two claims is identical. Moreover, we cannot derive any inconsistency from them even with the further premise that, as a matter of fact, Sarah is uncertain as to which of herb A and herb B is the optimal herb and which is the killer herb. That is because, as a general matter, the following scheme (often known as ‘factual detachment’ for the wide-scope ‘ought’) is invalid:

A
 Ought(A→B)
 Therefore, Ought(B)

How does this help the hard invariantist to account for the intuition that both Sarah and her professor speak truly? It is easy to explain why the professor speaks truly: his utterance expresses (1), which is true. As for Sarah, the hard invariantist has a few different options. She can claim that Sarah’s utterance actually expresses the wide-scope claim (2), despite the absence of any obvious syntactic marker that it does so. Less ambitiously, she can claim that Sarah’s utterance seems true to us because it is ‘in the ballpark’ of a truth, namely (2).

There is an important complication, however. While there is consensus that factual detachment fails for wide-scope ‘ought’s, a schema known as ‘deontic detachment’ is classically assumed:¹⁴

Ought(A)
 Ought(A→B)
 Therefore, Ought(B)

Standard deontic logic (specifically, the ‘K axiom’ of standard deontic logic) vindicates deontic detachment for wide-scope ‘ought’s. And we can develop Sarah’s case so that it is clear that she *ought* to be uncertain about which herb to prescribe. For example, if Sarah had learned that one of herbs A and B would save Sam and one would kill him, but had not yet been told which is which, then it would be the case that Sarah ought to be uncertain as to which is the killer herb and which is the optimal herb. In such a case, by deontic detachment, Sarah ought to prescribe herb C, so an inconsistency with (1) follows.

A hardline wide-scope may simply read such a result as a reason to reject the validity of deontic detachment. Indeed, Broome denies deontic detachment.¹⁵ We ourselves find

¹⁴ Cf., e.g., (Hare 1971; Greenspan 1975).

¹⁵ Broome (2013, 120–1, 127–8).

deontic detachment plausible, both in the abstract and in the particular case of Sarah.¹⁶ In any case, however, there is a second maneuver available to the hard invariantist that can deal with cases where it is clear that the agent's uncertainty is justified. Specifically, the hard invariantist can deny in such cases that the third party (in this case the professor) really speaks truly, by claiming that the third party mistakenly projects his or her own epistemic position onto the agent. The hard invariantist might suggest that once we clarify that the agent's uncertainty was justified, the tendency to endorse the third party's judgment dissipates at least somewhat.¹⁷ Moreover, to the extent that the intuition remains, we may *ourselves* be overprojecting our knowledge of the stipulated fact to which the third party is supposed to be privy in the example. The hard invariantist can appeal to the fact that Sarah could reasonably protest, in response to the professor, 'that's not fair! You hadn't yet told us which herb is the cure! I did exactly what I ought to have done.'

So we have a pleasing two-prong recipe for a hard invariantist resistance of contextualism in uncertainty-based cases like that of the herbalist. If it seems that the subject *ought not* to have been uncertain, then we say that it is the professor (or analogue to such a figure) that speaks truly, and that, while the subject speaks falsely, her assertion communicates something true, namely a wide-scope requirement governing the relationship between her doxastic states and actions. If it seems that the subject *ought* to have been uncertain, then we say that it is the subject who speaks truly, and that the professor-figure engages in mistaken projection that ends in a false utterance.

Again, we think that this attractive strategy fails in the case of normative uncertainty that we develop below. But before showing why, we have one more alternative position to explicate: that of the relativist.

2.3 Relativism

Like the hard invariantist, the metanormative relativist holds that 'ought' takes on the same semantic value in the 'ought' claims of both Sarah and her professor. But the relativist seeks to more fully accommodate the intuition that both Sarah and the professor speak truly. According to the relativist, (many) normative claims are neither true nor false absolutely, but are only true or false relative to some 'context of assessment' (Kolodny and MacFarlane 2010; MacFarlane 2014). A context of assessment may include a particular body of information, as well as particular standards of assessment. It is important to see that the *same* token utterance can be assessed as true or false at various different contexts of

¹⁶ See Worsnip (2015b, §3.8) for critical discussion of Broome's arguments against deontic detachment. Moreover, an even weaker principle than deontic detachment suffices to make trouble for the wide-scoping strategy. This weaker principle says that $Ought(A \rightarrow B)$, $Ought(A)$ and (the actual fact that) A all together entail $Ought(B)$. This weakened principle is immune to Broome's counterexamples to deontic detachment.

¹⁷ By analogy, consider the tendencies for post-hoc analysis of strategic decisions by sports fans. For example, roughly speaking, in baseball, attempting to steal a base raises the expected chance of the running team's scoring just if the baserunner is successful in stealing bases more than 75% of the time. Suppose a baserunner is generally successful 95% of the time; clearly, then, his decision to run is (given favorable circumstances) a good one. But, if he fails—if it turns out to be one of the unsuccessful 5%—sports fans will be liable to say, 'Argh! That was an awful decision! He shouldn't have tried to steal that base!' This seems a clear projection of what the fans know after the base-stealing attempt onto the epistemic position of the baserunner.

assessment. Although individuals making assessments occupy contexts of assessment, or ‘perspectives’, there is no determinate ‘assessor’ for an individual utterance in the way that there is a determinate speaker and subject; rather, there are as many possible contexts of assessment for the utterance as there are possible sets of information and standards.

In the case of Sarah, the relativist can say the following. Assessed from the professor’s vantage point, the proposition expressed by Sarah’s claim that she ought to administer herb C is false. But assessed from Sarah’s less informed perspective, this very same proposition is true. Some normative statements, at least, can be ‘true for you, but not for me’.

To be clear, the verdict is not straightforwardly that the professor’s claim is true and that Sarah’s claim is true. In fact, there is no *single* context of assessment from which we can say both of those things, since on the relativist account (by contrast to the contextualist account), the professor and Sarah express inconsistent propositions. But what we can say is that both speakers say something which is true at their own context of assessment. We are not by any means required to take up the same perspective that these speakers occupy when assessing the truth of their utterances; unlike contextualism, relativism gives no special privilege to the speaker’s context. However, the relativist can still explain why both claims seem sensible and true from their own contexts of utterance. This helps to explain why we feel tempted to endorse both claims as true, even though there is no single context of assessment from which we can actually do that.

Indeed, some relativists claim that relativism actually does *better* than contextualism in accounting for our intuitive reactions, since we may hesitate to say unequivocally that both Sarah and the professor speak truly, or that there is no disagreement between them. While their individual claims both seem reasonable, we may experience them as being in at least a certain amount of tension. Relativism can nicely account for this equivocal reaction.¹⁸

3. A new, normative uncertainty case

The cases that have thus far been the focus in the debate over metanormative contextualism have been cases involving what we will call *descriptive uncertainty* rather than *normative uncertainty*. A brief example will illustrate the distinction we have in mind. If you know that Abby ought to repay her lender, but do not know whether her lender is Beth or Carl, you may as a result be uncertain about whether or not Abby ought to pay money to Beth. While this is uncertainty about a normative claim, the uncertainty arises from uncertainty about a purely descriptive fact (the identity of Abby’s lender) and is compatible with your being completely certain about all of the general normative principles. If, on the other hand, you know that Beth loaned money to Abby but are uncertain about whether there is a *prima facie* obligation to repay lenders, then your uncertainty about whether Abby ought to pay money to Beth is fundamentally normative in character: your uncertainty could persist even if you were completely certain about all of the purely descriptive facts of the case.

¹⁸ See (Kolodny and MacFarlane ms.). In reply, many contextualists dispute whether the cases such as that of Sarah and the professor really amount to disagreements in the strong sense in which disagreement requires inconsistency. See e.g. Björnsson and Finlay (2010); Khoo & Knobe (ms.).

The case involving Sarah the herbalist that has been the focus here is a case where it is purely descriptive uncertainty that accounts for the utterance of multiple ‘ought’ claims that (superficially) appear to be inconsistent. Similarly, the most prominent example featured in recent debates is also clearly one featuring descriptive rather than normative uncertainty.¹⁹ Moreover, the theoretical development of the basic idea that there is more than one ‘ought’ in play in such cases generally assumes that it is descriptive rather than normative uncertainty that generates the need to appeal to multiple senses or semantic values for ‘ought’. The simple idea is that we have different possible bodies of descriptive information, but hold general normative principles fixed, as providing something like a function from that body of information to claims about what ought to be done relative to the information. So, for example, if utilitarianism is true, we just work out what maximizes expected utility relative to whatever body of information we are concerned with.²⁰

However, recent work elsewhere has suggested that *normative* uncertainty presents normative questions that may not be fully settled by first-order normative theories.²¹ So far, the case of normative uncertainty has not been brought to bear on the debate surrounding ‘ought’ and contextualism. In this section, we argue that shifting to an example involving normative uncertainty rather than descriptive uncertainty allows for a clear and compelling argument in favor of metanormative contextualism over its rivals.

The example we will focus on concerns Greg, a policy maker who has the power to enact one of three proposed tax policies: policy A, policy B, and policy C. Policy A significantly increases the total well-being of the population as a whole, but does nothing to improve the lot of the worst off. Policy B significantly improves the lot of the worst off, but does not significantly affect the total well-being of the population. Policy C significantly improves the lot of the worst off, though not quite as much as policy B, and it also significantly increases the total well-being of the population as a whole, though not quite as much as policy A. Different theories of distributive justice will rank these policies differently. Utilitarianism holds that total well-being ought to be maximized and thus implies that Greg ought to choose policy A. Lexical prioritarianism holds that the well-being of the worst off trumps other considerations and thus implies that Greg ought to choose policy B. Finally, weighted prioritarianism holds that the worst off ought to receive some priority, but not the sort of absolute priority that completely trumps considerations of total well-being. Many versions of weighted prioritarianism will imply that Greg ought to choose policy C.

Suppose Greg is certain that either utilitarianism or lexical prioritarianism is correct, but that he is not especially confident in either of these theories. His credence for utilitarianism is, say, 0.51, and his credence for lexical prioritarianism is 0.49. Deliberating in his office, he says the following to his colleague: ‘I know that the correct theory is either utilitarianism (in which case I ought to choose policy A) or lexical prioritarianism (in which case I ought to choose policy B). But since I have no idea which of these two theories is right, I ought to choose policy C. For on either theory, policy C is a close second’.

¹⁹ This is the ‘miners case’, due to Parfit (2011, 1:159). See discussion in (Kolodny and MacFarlane ms.; Björnsson and Finlay 2010).

²⁰ See, for example, (Jackson 1991).

²¹ See, e.g., (Sepielli 2009), and the references therein.

It is important not to misunderstand Greg's view here. His view is not that, since utilitarianism and lexical prioritarianism are both plausible, but somewhat extreme, theories, weighted prioritarianism is an even more plausible compromise between the two, and thus likely to be the true theory. Rather, he has credence *o* that weighted prioritarianism is actually the true theory of distributive justice: his whole credence space is divided between utilitarianism and lexical prioritarianism. (We can imagine that he has some particular reason for this: perhaps he thinks, for example, that weighted prioritarianism is very vague, and that no such vague theory could be the true theory.) Rather, his reason for following weighted prioritarianism is that he thinks it is a good compromise between the two theories he thinks may be correct, which is the appropriate way to act given his uncertainty. He chooses policy C because it is a close second on both of the metrics that he thinks could be correct.

Whether or not Greg expresses some mistaken views in his speech, the speech is entirely intelligible. Even if Greg is mistaken in thinking that either utilitarianism or lexical prioritarianism is correct, it is certainly intelligible for someone who is uncertain between these two views to think that the correct way to act in light of this uncertainty is to adopt some middle-ground action (one that happens to conform to some weighted prioritarian views). Of course, Greg might also be wrong about how to act in the face of uncertainty. Perhaps one should not adopt a course of action that gives some weight to both of the live possibilities, but should instead pick one policy (perhaps the one receiving the most credence) and perfectly conform to it. So Greg might be mistaken not only about his views on distributive justice, but also about the norms that describe how to act given uncertainty between these views. But mistaken or not, Greg's speech is understandable and seemingly coherent. It is not as though he has expressed some obviously inconsistent set of views.²²

And yet Greg has said *both* that either he ought to choose policy A or he ought to choose policy B *and* that he ought to choose policy C (an action that is acknowledged to be incompatible with choosing policy A or choosing policy B). Thus, Greg's case provides a challenge for proponents of the competing semantic theories of 'ought'—namely, explaining how it is that Greg speaks in an at least apparently reasonable way despite making two 'ought' claims that, when 'ought' is assigned the same semantic value, appear to be inconsistent. We will show how contextualism accounts for the apparent coherence of Greg's speech before arguing that relativism, soft invariantism, and hard invariantism are all incapable of doing so.

3.1 Contextualism

Unsurprisingly, the idea behind the contextualist account of what is going on with Greg's speech is that his two uses of 'ought' have different semantic values. As a result, they do

²² In fact, it's plausible that whatever the reasonable credences regarding moral theories are, they do involve some amount of uncertainty; and that given that uncertainty, one ought to hedge one's bets to at least some extent. These claims are defended at length in the literature on moral uncertainty; see, e.g., (Sepielli 2009). (Weatherson (2014) dissents here, but see §4 below.) But our argument does not actually depend on the claim that one ought to sometimes hedge one's bets in this way; all we need to claim is that Greg's utterances is understandable and not obviously unreasonable. Our argument starts from the apparent coherence of Greg's speech, not its truth.

not end up contradicting each other. Crudely, one might say that Greg's claim that he either ought to choose policy A or ought to choose policy B is a claim about what he ideally ought to do given the actually true theory of distributive justice, whereas his claim that he ought to choose policy C is a claim about what he ought to do given his state of subjective uncertainty about the theory of distributive justice.

Again, there are different possible proposals about how to implement this semantically, as there were in the descriptive uncertainty case. Perhaps the most plausible proposal is that there is a shift in the relevant standards. The first part of Greg's speech could be interpreted as a claim about the correct choice in light of the norms of distributive justice, while the statement that he ought to choose policy C could be interpreted as a claim about the correct choice in light of the norms that govern how to act under normative uncertainty.

Some moderate contextualists who are happy to allow that the semantic value of 'ought' is relative to information might want to resist the idea that the relevant standards can vary across different contexts of utterance. They may feel that the latter kind of contextualism results in a theory that is too promiscuous with truth, allowing agents to speak truly when they make normative claims simply by subscribing to the relevant standards that vindicate those claims. Those of a realist metaethical persuasion, who think that there are *correct* and *incorrect* standards independently of anyone's holding those standards, may balk at this.

But this worry is misplaced, since standards-contextualism is on closer inspection fully compatible with this realist position. The standards-contextualist can allow that some normative speech 'aspires to objectivity' in that it is a claim about what one ought to do *given the actually true normative standards*, not about what one ought to do *given the speaker's standards*. In other words, the standards-contextualist can deny that the standards parameter is always filled by the standards that the speaker subscribes to: when the speaker intends to make a more ambitious normative claim, the parameter can be filled by the actually true normative standards (if there be such).²³ Moreover, even if the standards-parameter is *never* filled simply by the subjective standards of the speaker, it may be filled by different kinds of standards on different occasions, compatibly with each set of standards being objectively correct. For instance, in Greg's case, it may be filled by the objectively true norms of distributive justice for one utterance, and by the objectively true norms of how to respond to normative uncertainty for the other utterance. So the choice between standards-invariantism and a view on which 'anything goes, always' is a false dichotomy.

However the contextualist semantics is to be precisely developed, it is clear that the contextualist needs to posit a change in context in the course of Greg's short speech that causes the semantic value of 'ought' to change. One might worry about this, given that Greg does not appear to have gained or lost any information in the course of his speech.²⁴

²³ This does not necessarily mean that the parameter is *always* filled by the actually true normative standards. On other occasions, the speaker might be making a less ambitious claim about what one ought to do given some set of standards that are assumed in the context.

²⁴ For various reasons, contextualists often hold that context-sensitive terms exhibit a certain 'stickiness', whereby it is difficult for either a single or multiple speakers to felicitously use the same context-sensitive

However, as we have just indicated, it is not obvious that the contextualist needs to understand the relevant parameter that shifts in value during Greg's speech as being information. More to the point, even if the relevant parameter is to be understood as information (specifically, *normative* information), recall the point we made in §1, namely that the contextualist is not committed to thinking that the information parameter is always filled by the information *available to* or *known by* the subject. When they intend to do so, it is possible for speakers to talk about what they ought to do relative to richer bodies of information, such as the full totality of the facts, even if these bodies of information are not (fully) available to them.

While the exact conversational mechanisms that govern shifts in contexts are hard to precisify, in Greg's case we do have a relatively clear marker of the context shift. Specifically, to introduce the 'normative uncertainty "ought"', Greg says, 'since I have no idea which of these two theories is right...'. This makes the perspective of normative uncertainty salient. Part of what governs shifts in context are the speaker's intentions to talk about different things, and here Greg signals an intention to talk about what he ought to do *given his uncertainty*.

Given this shift in context, what Greg says can be interpreted as consistent and cogent. It is important to see that the consistency contextualism wins us is not merely consistency in some impoverished sense, or in some specialized 'dynamic' or 'cross-context' sense. Utterances are inconsistent with one another by expressing inconsistent propositions. And when a term has different semantic values on different occasions of utterance, the propositions expressed are correspondingly different. So, if one says 'I ought to Φ ' and also says 'I ought not to Φ ', but 'ought' takes different semantic values on the two occasions, then the proposition one expresses with the second utterance is simply not inconsistent with the proposition one expresses with the first utterance. Any accusation of inconsistency is a simple mistake that is like the inverse of equivocation: it trades upon treating the term as having the same semantic value on the two occasions of utterance, when these semantic values are in fact different.

3.2 *Relativism*

One might be tempted to think that the relativist can simply mimic the contextualist's way of resolving the apparent inconsistency in Greg's speech. If the contextualist can say that the context of utterance changes during Greg's speech, why can't the relativist say that the context of *assessment* changes during Greg's speech?

It is not clear how to interpret this proposal. It is a determinate feature of any particular token utterance of a sentence that it occurs in some specific context. Individual token utterances in this sense *have* a context of utterance. When it comes to context-sensitive expressions, one can make a *mistake* by interpreting the utterance's meaning against the *wrong* context of utterance—that is, a different context of utterance than the one in which it is actually made. By contrast, individual token utterances do not *have* a context of

term with differing semantic values in close succession, at least not without a clear signal that there has been a context shift in the relevant values of the contextual parameters that determine these semantic values (see, e.g., DeRose 2004; Hannon 2015).

assessment (MacFarlane 2014, 61–2), and there are, on the relativist view, no semantic rules that determine a single, uniquely “relevant” context of assessment from the context of utterance. Nor is one *required* to assess an utterance from any particular context of assessment, in the sense that one would be making a mistake by assessing it from some different context of assessment. Rather, according to the relativist, a particular token utterance can be assessed for its truth-value at any number of different contexts of assessment. That is what makes the view distinctively relativist. Thus, while it is perfectly clear what it means to say that *the* context of utterance changes during Greg’s speech, it is not at all clear what it means to say that *the* context of assessment changes during Greg’s speech.

Perhaps what is meant is that the context of assessment that Greg himself occupies changes during his speech. When he says that he either ought to choose policy A or ought to choose policy B, he is projecting into the perspective of someone with full knowledge of the true norms of justice.²⁵ By contrast, when he says that he ought to choose policy C, he is speaking from his own context of normative uncertainty.

The first point to make in reply to this proposal is that, whatever it is supposed to achieve, there is still a crucial disanalogy between a contextualist appeal to a change in context of utterance, and a relativist appeal to a change in context of assessment. By appealing to a shift in the semantic value of ‘ought’, the contextualist explains how Greg does not actually express two inconsistent propositions. But whatever the appeal to a shift in context of assessment comes to, it will not deliver the result that Greg does not express inconsistent propositions. On the contrary, it is distinctive of the relativist view that utterances made by speakers occupying different contexts of assessment *do* express inconsistent propositions. In the case of Sarah and her professor, the relativist is at pains to stress that Sarah and her professor do express inconsistent propositions. Indeed, as noted in §2.3, relativists claim this as a data point in *favor* of their view, since it preserves the idea that there is a robust form of disagreement between the speakers. It is just that, from Sarah’s context of assessment, her utterance is true and the professor’s is false, whereas from the professor’s context of assessment, his utterance is true and Sarah’s is false.

So in a simple sense, the contextualist can deliver consistency to Greg’s speech in a way that the relativist cannot. What the relativist could deliver, at most, would be that each utterance Greg makes is true at the context of assessment he occupies when he makes it. That does not obviously deliver the *coherence* of the speech as a whole. After all, on the relativist account, from the context of assessment that Greg occupies when he says that he ought to choose policy A or he ought to choose policy B, the claim that he ought to choose policy C is false; whereas from the context of assessment that Greg occupies when he says that he ought to choose policy C, the claim that he ought to choose policy A or ought to choose policy B is false.

The relativist might initially be willing to bite that bullet. Perhaps, the relativist might say, Greg is just torn between two different perspectives here. Thus, the relativist might settle for the claim that each utterance Greg makes will seem true to him at the context of

²⁵ Although, clearly he cannot do this fully, since if he could, he would know *which* of policies A and B to pick. But we will let this pass.

assessment he occupies when he makes it. Unfortunately, however, the relativist cannot even deliver this result. Consider Greg's speech again. What is striking about the speech is that it forms a complete piece of *reasoning*. Greg's reason for ultimately concluding, 'I ought to choose policy C' is that he is sure that *either* utilitarianism is correct (i.e., he ought to choose policy A) or that lexical prioritarianism is correct (i.e., he ought to choose policy B), but uncertain as to which—and that policy C is a very close second on both metrics. His claim that he either ought to choose policy A or ought to choose policy B thus serves as a supporting *premise* in generating the claim that he ought to choose policy C. Yet, for the relativist, any context of assessment at which it is true that Greg ought to choose policy C must also be a context of assessment at which it is false that Greg either ought to choose policy A or ought to choose policy B. And if we interpret Greg as occupying a context of assessment relative to which it is false that he either ought to choose policy A or ought to choose policy B, then he loses his rationale for concluding that he ought to choose policy C. In other words, the relativist cannot identify a context of assessment from which Greg can reflectively endorse his claim that he ought to choose policy C. Consequently, his claim that he ought to choose policy C ceases to be rationally intelligible.

The contextualist has no analogous problem. On the contextualist's view, Greg uses the premise that he ought (on a more objective reading of 'ought', viz. what one ought to do given the true theory of justice) to choose policy A or he ought to choose B to generate the conclusion that he ought (on a more subjective reading of 'ought', viz. what one ought to do given one's normative uncertainty) to choose policy C. There is no problem with being able to simultaneously endorse both of these judgments from one perspective, because fundamentally they are consistent propositions that just happen to both use the same context-sensitive term, 'ought'. When you are speaking at one context of utterance, you have no problem seeing the truth of claims made at different contexts of utterance, even if the sentences used to express those claims would be false uttered at your context of utterance. For you can see that the underlying propositions that those sentences express at their actual contexts of utterance are true. Nothing about your occupying your present context of utterance interferes with that judgment. By contrast, it's part of what it is to *occupy* a perspective or context of *assessment* in the relativist's sense that you assess other utterances by those standards too.

At this point, the relativist might protest that she has her own way of resolving these problems.²⁶ She might say that Greg can occupy a more detached perspective from which he acknowledges, compatibly, that relative to one context of assessment—CA₁—he either ought to choose policy A or ought to choose policy B, but that relative to another context of assessment—CA₂—he ought to choose policy C. And his speech and reasoning reflect his acknowledgment of both of these things.

The problem with this proposal is not that it is implausible but rather that, on closer inspection and development, it turns out to concede the truth of contextualism. Greg does not say anything explicit about different contexts of assessment in his speech. So to develop the current proposal to explain the apparent coherence of Greg's speech, we need to say that Greg's claim that he ought to choose policy A or he ought to choose policy B is a way

²⁶ Thanks to the editors of *Mind* for sharpening this objection.

of *talking about* what is true relative to CA₁, and that his claim that he ought to choose policy C is a way of *talking about* what is true relative to CA₂. But now we can see why this proposal is in fact a particular form of contextualism. For it effectively posits an unarticulated, contextually-determined parameter for the semantic value of ‘ought’ that could be made explicit by saying ‘relative to context of assessment X...’. Even though the content of this parameter concerns a context of assessment, the value of this parameter is supplied, on the present proposal, by the context of *utterance*, and as different contexts of utterance supply different implicit relativizations, the semantic value of ‘Greg ought to Φ ’ will change.²⁷ That is a (particular kind of) contextualist view.

As we clarified at the start of this section, the relativist view as usually understood does not involve this contextualism: it is not claimed that there is a uniquely relevant context of assessment for every token utterance, or that claims of the form ‘you ought to Φ ’ *mean* ‘relative to context of assessment CA, you ought to Φ ’.²⁸ Rather, a single token utterance of the form ‘you ought to Φ ’ can be assessed as true or false at any number of contexts of assessment. If the relativist switches to a view on which context of assessment features as a semantic parameter, the value of which is fixed by context of utterance, then, she purchases a satisfactory handling of Greg’s case only in virtue of conceding the truth of metanormative contextualism. And metanormative contextualism is what we have set out to establish with Greg’s case.

As the above discussion helps to bring out, relativism and contextualism are not contradictory views, strictly speaking. Indeed, not only is relativism compatible with contextualism, but leading relativists such as Kolodny and MacFarlane acknowledge that the semantic value of ‘ought’ is context-sensitive in certain respects (2010, 131). For example, when someone says, after looking at a certain chess position, that Black ought to move her queen (in order to avoid a needless queen sacrifice), she does not contradict the father who, when confronting this same position in a game against his young son, says to himself, ‘I ought to leave my queen where she is, in order to give my son an opportunity to win the match’. Kolodny and MacFarlane would acknowledge that these two ‘ought’ claims are not inconsistent, since the contextually relevant ‘deontic selection function’ (to use their terminology) differs across the two utterances. In the first utterance, the relevant selection function evaluates options in light of the goal of placing white in checkmate, while in the second utterance the relevant selection function evaluates options according to some other criteria (e.g., moral criteria, or the criteria for being a good parent).

So Kolodny and MacFarlane do not oppose contextualism per se. What they oppose is a contextualist resolution to the puzzle presented by cases like Sarah the herbalist. One who seeks to explain the herbalist case in contextualist terms must endorse a version of contextualism that is more thoroughgoing than the restricted version admitted by Kolodny and MacFarlane. According to this more thoroughgoing contextualism, the semantic value

²⁷ Arguably, Harman’s (1975; 1996) version of ‘moral relativism’ is a disguised contextualism of this sort.

²⁸ Indeed, as we have already noted, one of the central arguments for relativism—that those speaking from different contexts nevertheless disagree with one another in a robust sense involving inconsistency—crucially turns on the *rejection* of such a view. If two speakers are each only talking about what is true at their own contexts of assessment, and their contexts of assessment differ, there is no longer any inconsistency in what they say.

of ‘ought’ is not fixed once it is determined that the speakers are concerned with the ‘ought’ of morality (or of chess, or prudence, or what have you). Even when restricting ourselves to the ‘ought’ of morality, for example, determining the semantic value will depend on further parameters (e.g., information) that vary with context. The contextualism required to explain Greg’s case goes much further than merely acknowledging the difference between the ‘ought’ of morality and the ‘ought’ of chess: the use of ‘ought’ to convey the normative implications of various states of uncertainty requires that the semantic value of ‘ought’ vary in a more systematic way that is along the lines of the contextualism posited to explain the case of Sarah the herbalist.²⁹

Kolodny and MacFarlane reject this thoroughgoing contextualism and argue that relativism can explain why Sarah and the professor both speak reasonably in the herbalist case. But as we’ve argued here, relativism cannot explain the reasonability and apparent coherence of Greg’s speech. This does not by itself show that relativism is false, but only shows that the relativist will either need to invoke a more thoroughgoing version of contextualism (thus conceding our argument) or appeal to one of the invariantist strategies. Moreover, to the extent that we should prefer a unified solution to the case of Greg and that of Sarah—cases that, notwithstanding the difference between normative and descriptive uncertainty, have some striking similarities—the case of Greg puts pressure on the relativist proposal about Sarah’s case.

We now argue that the invariantist strategies also fail to explain the apparent coherence of Greg’s speech, leaving metanormative contextualism as the only viable option.

3.3 *Hard invariantism*

Recall the two-pronged strategy that we sketched for the hard invariantist in §2.2 for dealing with cases of descriptive uncertainty: either say that the ‘ought’ under uncertainty takes wide-scope, or say that the ‘ought’ under certainty is simply false and a result of mistaken projection. Do either of these options help with examples like that of Greg?

As we suggested earlier, the wide-scoping option is at its most plausible when the uncertainty on the part of the agent is not in fact justified. According to the wide-scoping option, the most charitable way to interpret Greg’s beliefs involves ascribing to him (i) the belief that he ought to choose policy A or he ought to choose policy B, and (ii) the (wide-scope) belief that he ought, if he is uncertain about which of policy A or policy B he ought to choose, to choose policy C. If Greg believes that his own uncertainty about whether he ought to choose policy A or ought to choose policy B is unjustified (or, indeed, if he just fails to believe that his uncertainty is justified), then it is clear that he can coherently believe (i) and (ii) together. But just as in the case of Sarah, it is easy to imagine the case so that

²⁹ Admittedly, Kolodny and MacFarlane themselves only commit to a relativist treatment of the informational component of the semantics. So if we are determined to understand descriptive uncertainty as ‘informational’ and normative uncertainty as ‘non-informational’ (though this is a questionable bifurcation), it may look like they are not committed to a relativist treatment of normative uncertainty. Since they are silent on normative uncertainty, it is not entirely clear what their treatment of it would be. Perhaps they would concede a fully contextualist treatment of normative uncertainty. But the point remains that this would amount to a concession of a much more thoroughgoing contextualism than they have seen the need to admit on other grounds.

Greg believes that his normative uncertainty is justified, indeed even required (that is, that he ought to be uncertain). Adding this stipulation does not in any way undercut the impression that Greg's speech is reasonable and coherent, which suggests that the apparent coherence of Greg's speech does not derive from an implicit assumption that he believes that his uncertainty is unjustified.

Note that simply denying that normative uncertainty is ever justified, as some philosophers are wont to do, is not sufficient to defend the wide-scoping option, at least so long as Greg himself thinks that his normative uncertainty is justified.³⁰ Nor is it likely that when typical hearers perceive Greg's speech as coherent, they are tacitly assuming that his normative uncertainty is unjustified. If it is a fact that normative uncertainty is never justified, it is a surprising one and not one that is in the background when typical hearers interpret Greg's speech as coherent. A viable *semantic* thesis should help to explain why Greg and others *take* his speech to be a reasonable one.

Another suggestion aimed at rescuing the wide-scope variant of hard invariantism would be that Greg's beliefs are unproblematically coherent because *deontic dilemmas* are possible. For our purposes, an agent counts as being in a deontic dilemma just in case it is true from a single context of assessment, and without any variation in the semantic value of 'ought', that the states of affairs he ought to bring about are collectively incompatible. If we stipulate that Greg believes that his uncertainty is required, and if we accept what the hard invariantist who takes the wide-scoping option has to say about Greg's beliefs, then the only way for Greg's package of beliefs to all come out true would be if Greg is in a dilemma of this sort.

This is easy to see if deontic detachment (introduced in §2.2 above) is correct. In this case, the wide-scope claim, together with the claim that Greg ought to be uncertain about whether he ought to choose policy A or ought to choose policy B, would entail that Greg ought to choose policy C, an action that is incompatible both with his choosing policy A and with his choosing policy B. One might think that if we deny deontic detachment, the deontic dilemma disappears. But in fact this is not so.³¹ Acting so as to ensure the truth of the conditional that [if Greg is uncertain about which of policy A or B he ought to choose, then he chooses policy C] requires ensuring either that Greg chooses policy C (which is incompatible with choosing policy A and with choosing policy B) *or* that he is not uncertain about which of policy A or policy B he ought to choose (which is incompatible with the claim that he ought to be uncertain). Since on the wide-scoping option for hard invariantism Greg's package of beliefs would imply that he is in a deontic dilemma, a natural move in the effort to explain the apparent coherence of Greg's speech would be to appeal to the coherence of the claim that Greg's situation constitutes a deontic dilemma.³²

³⁰ For arguments that appear to support the surprising conclusion that normative uncertainty is never justified, see Kiesewetter (ms.) and Titelbaum (2015), though the latter restricts his thesis to beliefs about *rationality*.

³¹ Thanks to an anonymous referee here. Interestingly, though Broome denies deontic detachment (as noted earlier), he seems hostile to the possibility at least of *rational* dilemmas (Broome 2007: 364).

³² A hard invariantist might alternatively make the proposal that Greg is in a deontic dilemma quite independently of the wide-scoping strategy. Our objections here apply to either version of the proposal.

Is it plausible to maintain that the apparent coherence of Greg's speech is explained by our recognizing at some level the plausibility of deontic dilemmas? We don't think so. For starters, whether Greg's speech appears coherent does not seem to depend on whether one finds deontic dilemmas plausible. (We, at least, find Greg's speech to be eminently reasonable even though we are not inclined to endorse deontic dilemmas.) Indeed, again, it does not seem to be a presupposition of his speech that *he himself* (even implicitly) takes himself to be in a deontic dilemma, which suggests that this is not what explains his speech's apparent coherence. Moreover, if Greg's hearers implicitly understood him to be describing a dilemmatic situation, then they would likely be left wondering what he will in fact do. But in fact, after hearing Greg's speech, with its superficially incompatible 'ought' claims, it is clear that Greg has resolved to choose policy C. The speech exhibits reasoning that resolves the question of how to act rather than reasoning that shows why Greg is torn between incompatible requirements and condemned to moral failure. Finally, and relatedly, the two judgments—that he either ought to choose policy A or ought to choose policy B, and that he ought to choose policy C—do not appear to be putting any kind of pressure on each other, as conflicting 'ought'-judgments do in cases that are typically taken as candidates to be deontic dilemmas.³³ (Indeed, as noted in the discussion of relativism above, the former judgment actually serves as the *grounds* for the latter.) This strongly suggests that the two 'ought's are in fact taking different semantic values, rather than a single value that is giving rise to a deontic dilemma.

The discussion thus far gives us reason for thinking that the first hard invariantist option, appealing to a wide-scope normative belief, will not adequately explain the apparent coherence of Greg's speech. That leaves us with the second prong of the hard invariantist strategy, namely that the 'ought' under certainty is just a mistaken projection. Obviously, the notion of mistakenly projecting makes more sense when we have one speaker who projects his own epistemic position onto that of a different subject, as the professor does. But Greg cannot be doing this to himself. So the best the hard invariantist can do here is to bite the bullet and simply maintain that Greg's claim that either he ought to choose policy A or he ought to choose policy B is just false, and that it is guaranteed to be so by his claim that he ought to choose policy C.

But again, this suggests that Greg's speech should appear incoherent, for the same reason we gave in response to the relativist. If Greg's claim that he ought to pursue policy C entails that he speaks falsely in saying that he ought to choose policy A or he ought to choose policy B, then his reasoning for the former claim is self-undermining, since the latter claim is a premise in his argument for it. Consequently, this version of the hard invariantist view makes him blatantly incoherent.

³³ Of course, the advocate of deontic dilemmas holds that in such cases, ultimately the two conflicting 'ought'-claims are logically consistent. But they still put at least some *prima facie* epistemic pressure on each other. For example, in Sartre's (1946/2007) famous example, the judgment that one should look after one's mother puts at least *some* epistemic pressure on the judgment that one should fight for the free French, and vice versa. Each provides at least some epistemic reason to doubt the other, even if we ultimately conclude that both judgments are true.

Could the hard invariantist maintain instead that it is Greg's claim that he ought to choose policy C that is false?³⁴ So-called "objectivists" hold that what one ought to do (in any legitimate, genuine sense) is simply not sensitive to one's epistemic position (or one's uncertainty),³⁵ and so will maintain that there is no legitimate sense in which Greg ought to choose policy C. But on closer inspection, this first-order normative view does not provide any defense of hard invariantism over contextualism as competing *semantic* theses. The fact that Greg speaks as he does, and seems coherent in doing so, still suggests that he is using 'ought' in a way that presupposes that differences in epistemic position can make a difference to what one ought (in at least some sense) to do. If objectivism is true, then this presupposition is false, and so claims that make use of such uses of 'ought' are false (or, at least, suffer from presupposition failure). But that does not mean that they are not *using* 'ought' in the way we have described. The objectivist view can yield an error theory about non-objective usages of 'ought', rather than denying that there *are* such uses.³⁶

Of course, an objectivist *could* instead say that all usages of 'ought' do in fact express the objective, uncertainty-insensitive 'ought'. But this theory simply gives no explanation of the apparent coherence of Greg's speech (and, more broadly, the wide range of linguistic data suggesting that people do in fact use 'ought' in a (putatively) uncertainty-sensitive way). So we think the error-theoretic version of objectivism, on which it is married with contextualism at the semantic level, is more promising. Again, our claim was not that Greg's speech is true but only that it is (at least apparently) coherent. The error-theoretic objectivist, by accepting contextualism, can agree with us that Greg's speech is coherent, and that his two claims do not contradict each other. She just goes on to say that one of them, given the normative truth of objectivism, is false. While we ourselves do not accept objectivism, our argument does not turn on its rejection.

3.4 Soft invariantism

Finally, we turn to soft invariantism. Recall from §2.1 that attempts to appeal to ambiguity between different *normative* 'ought's—for example, between 'objective' and 'subjective' 'ought's—are subject to a severe regress problem. Since one can be uncertain about what one ought to do even in the more 'subjective' sense posited, the soft invariantist is forced to posit a third 'ought', which in turn can be attended by uncertainty, forcing a fourth 'ought', and so on. When the semantic values of 'ought' begin to multiply in this way, it is

³⁴ Thanks again to an anonymous referee for this suggestion.

³⁵ See e.g. Graham (2010).

³⁶ Compare more familiar error theories about the moral 'ought' of the sort defended by Mackie (1977) and Joyce (2001). Joyce is explicit (and here he is drawing out what is plausibly latent in Mackie also) that he thinks there are true claims that use 'ought', so long as such claims only use the *rational* 'ought' as opposed to the moral ought. But his strategy is not to claim that every usage of 'ought' actually expresses the rational 'ought'. On the contrary: what makes him an error theorist is precisely the fact that he thinks that much ordinary usage of 'ought' *does* clearly employ the moral 'ought', presupposing that there are objective, categorical moral standards of the sort that Joyce thinks there cannot be. Consequently, on Joyce's view, much ordinary normative talk turns out to be false. Similarly, the objectivist can recognize that there are subjective or information-sensitive usages of 'ought' but be an error theorist about such usages, holding that they result in claims that are false (or suffer from presupposition failure).

clear that these various semantic values possess an underlying unity that calls for a contextualist semantics rather than a proliferation of multiple senses.

This problem applies equally when it is normative uncertainty that is at issue. One can be uncertain, not just about what the true theory of normative ethics is, but about what the true theory of how one ought to act under normative uncertainty is, and about what the true theory of how one ought to act under uncertainty about the true theory of how one ought to act under normative uncertainty is, and so on.³⁷

It's worth noting that this problem also arises for another proposal that may occur to a soft invariantist in the case of normative uncertainty—namely to try to distinguish between a 'moral sense' of 'ought' and a 'rational sense' of 'ought'.³⁸ Suppose one says that morally, Greg ought to choose whichever policy is actually recommended by the true theory of distributive justice, but rationally, he ought to choose policy C. This is no better than the binary objective/subjective distinction. Again, Greg can be uncertain about how he ought to act under normative uncertainty, and that forces this particular kind of soft invariantist to admit a proliferation of senses of the rational 'ought', just as it forced our original soft invariantist to admit a proliferation of senses of the subjective 'ought'.³⁹

In the case of descriptive uncertainty, however, recall, there was a different kind of soft invariantism that avoided these problems. This was to maintain that the sense in which Sarah ought to have given herb B was some non-normative sense (unlike the sense in which she ought to have given herb C). So when the professor says that Sarah ought to have given herb B, this 'ought' claim does not express a normative view about what Sarah had (or has) most reason to do, but rather (for example) an *evaluative* view about the optimific action. So even if the soft-invariantist is forced to concede that 'ought' can be context-sensitive, she can seek to preserve an invariantist semantics for the *normative* 'ought' by insisting that any the context-sensitivity is limited to the non-normative senses of 'ought'.

While this sort of soft invariantism is, we think, relatively plausible in the case of descriptive uncertainty, it is much less plausible in the case of normative uncertainty. Normative uncertainty is, after all, uncertainty about the truth of some *normative* theory.

³⁷ See also Sepielli (2014).

³⁸ Of course, the contextualist can recognize a difference between moral uses of 'ought' and rational uses of 'ought'. What follows is not an argument against distinguishing the moral and rational 'ought' per se, but rather against trying to use this distinction to resist contextualism.

³⁹ A soft invariantist might try replying to this by claiming that the 'rational "ought"' is not actually a normative 'ought'. In this way, like the soft invariantist who claims that only the evaluative 'ought' is context-sensitive, she avoids admitting contextualism about the *normative* 'ought'. Inspiration for this proposal might be drawn from Kolodny (2005), who argues that the requirements of rationality are not normative. However, the 'requirements of rationality' that Kolodny denies to be normative are pure *coherence* requirements. Central to Kolodny's argument is the idea that coherence never enters into deliberation: one does not think 'I ought to Φ , because Φ -ing will make me rationally coherent'; rather, one thinks, 'I ought to Φ , because Φ -ing would be good, or right, or get me what I want' – and coherence comes about at most as a byproduct of this kind of good deliberation. Hence, any sense in which one ought to be coherent *as such* is (arguably) a non-normative sense. But it is clear that the same argument does not apply merely to 'ought's under normative uncertainty. On the contrary, for Greg, the 'ought' under normative uncertainty plays a *crucial* role in his deliberation. So if what is meant by 'rational "ought"' is "'ought" under (normative) uncertainty' rather than "'ought" of satisfaction of coherence requirements', there seems to be little scope to deny the normativity of this 'ought'.

In the case of Greg, for example, utilitarianism and lexical prioritarianism are opposing normative theories of how one ought to distribute goods, and so the ‘ought’ claims they issue must be *normative* ought claims. (By contrast, remember that Greg *denies* that weighted prioritarianism is the true theory of distributive justice.) So when Greg sees that utilitarianism implies that he ought to choose policy A and that lexical prioritarianism implies that he ought to choose policy B, what he is seeing (assuming he is not confused) are normative implications of these two theories. Thus, a reasonable interpretation of Greg’s speech requires that we understand his claim that either he ought to choose policy A or he ought to choose B as using ‘ought’ in a normative sense.

The soft invariantist might try to reply that, contrary to appearances, utilitarianism and lexical prioritarianism are evaluative theories of the optimistic state of world, and not of what one ought to do in the normative sense. We think this is *ad hoc*. Utilitarianism and lexical prioritarianism are commonly stated as normative theories about what actions one ought to perform. And even if the soft invariantist claims that these theories would be better understood as evaluative theories, it is still not clear why *Greg* need be thinking of them as evaluative when he is uncertain between them.

Moreover, a cousin of the regress problem that doomed other soft invariantist accounts also looms here. On the particular elaboration of the soft invariantist view that we are currently considering, it turns out that the theories of so-called ‘normative ethics’ are not actually normative but evaluative. The truly normative theory is that of how to act under uncertainty about which of these theories is correct. Call this a theory of ‘first-order’ uncertainty. But one can, in turn, be uncertain about which theory of first-order uncertainty is correct.⁴⁰ In such a case, we could talk about how one ought to act given the true theory of how to act under first-order uncertainty. But we could also talk about how one ought to act given one’s uncertainty about the true theory of how to act under (first-order) uncertainty. The latter would comprise a theory of second-order uncertainty.

Just as the soft invariantist blocked contextualism by saying that theories of normative ethics are not really normative, it now seems that she must block contextualism by saying that theories of first-order uncertainty are not really normative. After all, one can be uncertain about them, and there are theories about how one ought to act given this uncertainty. The truly normative theories, then, are the theories of how to act under ‘second-order’ uncertainty. But of course, one can also be uncertain about the correct theory of second-order uncertainty. And so the regress keeps on going. As long as it is possible to be uncertain about the correct theory of how to act under some level of n -th order uncertainty, we can cook up a case to show that the soft invariantist must claim that theories of n -th order uncertainty are not normative. There are, then, no stably normative theories for this kind of soft invariantist.

⁴⁰ For example, in addition to his uncertainty between utilitarianism and lexical prioritarianism, Greg might also be uncertain between two opposing theories of how one ought to choose when uncertain between competing theories of distributive justice: a ‘winner-takes-all’ view on which one ought to choose the policy that accords with the theory for which one has the highest credence; and a ‘compromise’ view on which (when possible) one ought to choose a policy like C that is ranked highly according to each theory (even if it is not the policy required by either theory).

Perhaps an interesting defense of this radical idea could be made. However, we think that it is probably not what soft invariantists had in mind from the start. Ultimately, metanormative contextualism probably saves more of what typical soft invariantists wanted than this radical view does. After all, metanormative contextualism and soft invariantism share the contention that ‘ought’ takes different semantic values on different occasions of utterance. But contextualism can unify these different semantic values under a single semantics rather than positing ambiguity.

4. The relevance of *normative uncertainty*

Observant readers will have noticed that there are really two important differences between the cases of Sarah and Greg. First, there is the fact that Greg’s uncertainty is normative, whereas Sarah’s is descriptive. But second, there is the fact that Greg’s case involves a *single speaker* making the two contrasting ‘ought’ claims, whereas Sarah’s involves two *different* speakers. One might naturally be suspicious that it is really the single-speaker feature of Greg’s case that is making all the difference in securing the case for metanormative contextualism. And clearly it does play some role. However, we also think that it is crucial that the case involves normative uncertainty. To show this, let us consider a single-speaker version of Sarah’s case, and briefly illustrate why it does not tell as decisively in favor of contextualism.

Suppose that as Sarah is preparing an herb C concoction to give to Sam, she says the following: ‘What I *really* ought to do is give you the herb that has no side effects. But I can’t remember whether that’s herb A or B, and giving you the wrong one would kill you. Given this uncertainty, I ought to give you herb C instead. It will save your life and leave you with only mild hypersensitivity to the sun’. Again, Sarah’s initial ‘ought’ claim (that what she really ought to do is give the herb with no side effects) superficially appears to be inconsistent with the ‘ought’ claim made two sentences later (that she ought to give herb C instead).

The crucial difference between this case and Greg’s, however, is that Sarah can and does reason to the conclusion that she ought to give herb C without relying on the premise that she ought to give herb A or B. She can reason to this conclusion purely from the descriptive facts that one of herb A and herb B will kill Sam, and that she can’t remember which. Greg, by contrast, reasons to the conclusion that he ought to perform policy C from the *normative* premise that either he ought to choose policy A or he ought to policy B, but doesn’t know which.

As we saw earlier, this difference is important in arguing against both the relativist and hard invariantist stances. Both relativists and hard invariantists portray Greg’s claim that he ought to choose policy A or he ought to choose policy B and his claim that he ought to choose policy C as inconsistent. But if they are inconsistent, then the former cannot serve as Greg’s rationale for the latter, since he cannot endorse both simultaneously, and so his judgment that he ought to perform policy C is not intelligible. By contrast, even if we cut out any claim on Sarah’s part that she ought to give the herb with no side effects, her claim that she ought to give herb C is fully intelligible. So it is easier to say that the two claims

are actually inconsistent while preserving the intelligibility of the second claim, since in Sarah's case the second claim does not depend for its justification on the first claim.

The significance of introducing normative uncertainty is further underscored when we consider the error theory offered by relativists Kolodny and MacFarlane (2010) as an explanation for why Sarah might mistakenly affirm that she ideally ought to give the herb without side effects. Even if their error theory offers a plausible explanation in the example of Sarah, the error theory fails to apply to Greg's case. According to Kolodny and MacFarlane, Sarah can appropriately make the following claims, each of which is true at her context of assessment:

- (i) If herb A is the optimal herb, then I ought to give herb A.
- (ii) If herb B is the optimal herb, then I ought to give herb B.
- (iii) Either herb A is the optimal herb or herb B is the optimal herb.

It would seem that these three claims logically entail a fourth claim:

- (iv) Either I ought to give herb A or I ought to give herb B.

On the supposition that (iv) is logically entailed by (i)-(iii), then we would have to affirm that (iv) is also true at Sarah's context of assessment, which is of course what Kolodny and MacFarlane aim to deny. But Kolodny and MacFarlane maintain that (i)-(iii) does *not* logically entail (iv), since according to their relativist semantics *modus ponens* is not logically valid for conditionals that (like (i) and (ii) above) have information-sensitive terms in the consequent. Because *modus ponens* is not valid, Sarah is unable to validly infer (iv) from (i)-(iii), and thus cannot validly arrive at a conclusion that is inconsistent with the claim that she ought to give herb C. But the reasons that account for the failure of *modus ponens* are quite subtle, which explains why Sarah would make the mistake of voicing (iv) in situations when she is entitled to assert (i)-(iii). We thus have an explanation for why we are inclined to endorse as appropriate Sarah's assertion of (iv), despite the falsity of (iv) at her context of assessment.

But when we turn to the case of Greg, the relativist cannot appeal to *modus ponens* failure in order to explain why Greg makes inconsistent claims. To see this, let's first see how Kolodny and MacFarlane's strategy would translate to Greg's case. Kolodny and MacFarlane affirm that Greg can appropriately make each of the following claims, each of which may be true at his context of assessment:

- (i*) If utilitarianism is true, I ought to choose policy A.
- (ii*) If lexical prioritarianism is true, I ought to choose policy B.
- (iii*) Either utilitarianism is true or lexical prioritarianism is true.

These three claims would seem to entail the following:

- (iv*) Either I ought to choose policy A or I ought to choose policy B.

Because (iv*) seems to logically follow from (i*)-(iii*), we have an explanation for why Greg is inclined to affirm it (and why we are inclined to judge that this affirmation is appropriate). But because Kolodny and MacFarlane hold that *modus ponens* is invalid in this setting, they would deny that (iv*) follows from (i*)-(iii*). This provides an explanation of why Greg (supposedly) cannot properly affirm (iv*), a claim that is inconsistent with his claim that he ought to choose policy C.

Unfortunately for the relativist, however, the claim that Greg cannot affirm (iv*) does not save Greg from inconsistency entirely. This is because (iii*), itself being a normative claim, is *already* inconsistent with Greg's claim that he ought to choose policy C. In Sarah's case, (i)-(iii) contain only *conditional* normative claims. *Modus ponens* is needed to extract an unconditional normative claim, (iv), that can conflict with Sarah's claim that she ought to give herb C. And since this particular instance of *modus ponens* is supposedly invalid, no conflict can be generated. In Greg's case, however, (iii*) is an unconditional normative claim (unlike the merely descriptive (iii)). Because (iii*) is already normative in character, we do not need to use *modus ponens* in order to extract a claim that is inconsistent with Greg's verdict that he ought to choose policy C. (iii*) itself is inconsistent with this verdict because (iii*) on its own entails (without recourse to *modus ponens*) that either one ought to maximize total well-being or one ought to do the action that is best for the least well off.

The soft invariantist strategy is also much better off in the case of descriptive uncertainty than that of normative uncertainty, even when both are taken in their single-speaker form. The soft invariantist can quite plausibly say that Sarah's claim that she ought to give herb A or B, like the professor's claim in the two-speaker case, is really an evaluative claim about what would be optimific. But, as we have already argued, this is far less plausible in Greg's case, where his claim that he ought to enact policy A or policy B is a result of his certainty that one of two *normative* theories of distributive justice is correct. So, again, the case of normative uncertainty is much stronger for the contextualist than the case of descriptive uncertainty is.

We have been suggesting that the phenomenon of normative uncertainty bears upon the debate between contextualist accounts of the normative 'ought' and their rivals. However, it may be also be that the viability of metanormative contextualism is significant for debates over the normative significance of normative uncertainty. Without a contextualist semantics for the normative 'ought,' we have argued, moral reasoning of the sort voiced by Greg would be positively incoherent. Yet those who affirm that moral uncertainty matters morally appear to affirm the same sort of reasoning that Greg does. After all, they take the view that it can happen that what you ought to do is *not* just whatever the correct moral theory says, but rather some action that takes account of your uncertainty between plausible moral theories. This contention can appear to amount to saying that you ought not do what you actually ought to do, which sounds incoherent.

Weatherson (2014, 146) has recently noted the apparent incoherence of this sort of response to moral uncertainty, and partly on this basis concludes that moral uncertainty has no bearing on what one ought morally to do. But if metanormative contextualism is correct, then this type of reasoning is not incoherent after all, at least not in any straightforward way. As we have been suggesting, such reasoning, at least in Greg's case, does not *appear* incoherent, and none of the rivals to contextualism can explain even this

apparent coherence. So we think it is better to conclude (with contextualists) that reasoning like it is not *in fact* incoherent, than to conclude (with Weatherston) that reasoning like Greg's must be incoherent and to deny the normative significance of normative uncertainty on these grounds. This of course does not establish that moral uncertainty has normative significance, but it does block one sort of incoherence argument that has been used to suggest that it does not.

5. Conclusion

Metanormative contextualism is the view that the semantic value of a given 'ought' claim depends on the value of one or more parameters that are fixed by the context of utterance. So, if there is even a single case that shows that 'ought' can have different semantic values in a way that cannot be explained by a soft invariantist appeal to 'multiple senses,' then metanormative contextualism follows. We have argued that a (single-speaker) case featuring normative (rather than descriptive) uncertainty is such a case. This type of case establishes the truth of metanormative contextualism even if in other cases explaining the data requires appealing to relativism or one of the invariantist strategies assessed here.

That said, there is reason to prefer a unified explanation of these many cases. So even if the relativist or invariantist can do just as well as the contextualist at explaining some particular cases, the clear superiority of contextualism in the kind of case developed here gives us a reason for preferring the contextualist explanation in cases where contextualism does not have such a clear explanatory advantage. So, while our argument does not provide a knock-down argument that relativism is false—for contextualism and relativism could both be true—it puts considerable pressure on the attempt to motivate relativism with recourse even to the 'stalemate' cases such as that of Sarah. For relativism is unable to explain the nearby and structurally similar case that we have offered.

The invariantist is in one way better off than the relativist and in one way worse off. She is worse off because if contextualism is true, it simply follows that invariantism is false, given how we defined 'invariantism'. But many of the phenomena that the invariantist used to explain away the existing contextualist data may yet be genuine ones. And in this respect, she is perhaps better off than the relativist, for there are considerable independent reasons to think both that, as the soft invariantist argues, 'ought' can have 'multiple senses' (normative, evaluative, and so on), and that, as the hard invariantist argues, it can sometimes take wide scope over conditionals (and express true claims in so doing). Nevertheless, we have argued that there is a particular kind of data that only contextualism can account for, and thus that contextualism is true.

There is a huge amount of work to be done in precisifying metanormative contextualism, and we have done virtually none of it here. Rather, we have tried to argue that some form of it must be correct, however the details are to be spelled out, and however the challenges are to be overcome. Since metanormative contextualism has the capacity to

shed light on many debates and problems in normative and metanormative philosophy, that should be welcome news.⁴¹

References

- Björnsson, Gunnar, and Stephen Finlay 2010. 'Metaethical Contextualism Defended'. *Ethics* 121, 7–36.
- Brogaard, Berit 2008. 'Moral Contextualism and Moral Relativism'. *The Philosophical Quarterly* 58, 385–409.
- Broome, John 1999. 'Normative Requirements'. *Ratio* 12, 398–419.
- 2004. 'Reasons'. In *Reason and Value: Themes From the Moral Philosophy of Joseph Raz*, edited by R. Jay Wallace, 2004–28. Oxford: Oxford University Press.
- 2007. 'Wide or Narrow Scope?' *Mind* 116, 359–70.
- 2013. *Rationality through Reasoning*. Chichester: Wiley Blackwell.
- Brunero, John 2010. 'The Scope of Rational Requirements'. *The Philosophical Quarterly* 60, 28–49.
- 2012. 'Instrumental Rationality, Symmetry and Scope'. *Philosophical Studies* 157, 125–40.
- Cariani, Fabrizio 2013. "Ought" and Resolution Semantics'. *Noûs* 47, 534–58.
- Cariani, Fabrizio, Magdalena Kaufmann, and Stefan Kaufmann 2013. 'Deliberative Modality under Epistemic Uncertainty'. *Linguistics and Philosophy* 36, 225–59.
- Dancy, Jonathan 1977. 'The Logical Conscience'. *Analysis* 37, 81.
- 2000. *Practical Reality*. New York: Oxford University Press.
- DeRose, Keith 2004. 'Single Scoreboard Semantics'. *Philosophical Studies* 119, 1–21.
- Dowell, Janice 2012. 'Contextualist Solutions to Three Puzzles about Practical Conditionals'. *Oxford Studies in Metaethics* 7, 271–303.
- 2013. 'Flexible Contextualism about Deontic Modals: A Puzzle about Information-Sensitivity'. *Inquiry* 56, 149–78.
- Dreier, James 1990. 'Internalism and Speaker Relativism'. *Ethics* 101, 6–26.
- Feldman, Richard 1988. 'Subjective and Objective Justification in Ethics and Epistemology in Justification'. *The Monist* 71, 405–19.
- Finlay, Stephen 2014. *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University Press.
- van Fraassen, Bas C. 1972. 'The Logic of Conditional Obligation'. *Journal of Philosophical Logic* 1, 417–38.
- Gibbard, Allan 2005. 'Truth and Correct Belief'. *Philosophical Issues* 15, 338–50.
- Graham, Peter A. 2010. "In Defense of Objectivism about Moral Obligation." *Ethics* 121 (1): 88–115.
- Greenspan, Patricia S. 1975. 'Conditional Oughts and Hypothetical Imperatives'. *The Journal of Philosophy* 72, 259–76.

⁴¹ The authors contributed equally to this manuscript. For helpful discussion and/or comments we are grateful to Gunnar Björnsson, John Broome, Keith DeRose, Dan Greco, Josh Knobe, Ralph Wedgwood, the editors at *Mind*, and several anonymous referees.

- Hannon, Michael 2015. 'Stabilizing Knowledge'. *Pacific Philosophical Quarterly* 96 (1): 116-139.
- Hansson, Bengt 1969. 'An Analysis of Some Deontic Logics'. *Noûs* 3, 373-98.
- Hare, Richard M. 1971. 'Wanting: Some Pitfalls', in his *Practical Inferences*, 44-58. London: Macmillan.
- Harman, Gilbert 1975. 'Moral Relativism Defended'. *The Philosophical Review* 84, 3-22.
- 1996. 'Moral Relativism'. In *Moral Relativism and Moral Objectivity*, by Gilbert Harman and Judith Jarvis Thomson. Cambridge, Mass.: Blackwell.
- Henning, Tim 2014. 'Normative Reasons Contextualism'. *Philosophy and Phenomenological Research* 88, 593-624.
- Jackson, Frank 1991. 'Decision-Theoretic Consequentialism and the Nearest and Dearest Objection'. *Ethics* 101, 461-82.
- Joyce, Richard 2001. *The Myth of Morality*. Cambridge, UK: Cambridge University Press.
- Kagan, Shelly ms. 'The Paradox of Methods'.
- Khoo, Justin and Knobe, Joshua ms. 'Moral disagreement and moral semantics'.
- Kiesewetter, Benjamin ms. 'You Ought to Φ Only If You May Believe That You Ought to Φ '.
- Kolodny, Niko 2005. 'Why Be Rational?' *Mind* 114, 509-63.
- Kolodny, Niko, and John MacFarlane ms. 'Ought: Between Objective and Subjective'. <http://phil.unlv.edu/ought-unlv.pdf>.
- 2010. 'Ifs and Oughts'. *The Journal of Philosophy* 107, 115-43.
- Kratzer, Angelika 1977. 'What "must" and "can" Must and Can Mean'. *Linguistics and Philosophy* 1, 337-55.
- 1981. 'The Notional Category of Modality'. In *Words, Worlds, and Contexts*, edited by H. J. Eikmeyer and H. Rieser, 38-74. Berlin: de Gruyter.
- 1991. 'Modality'. In *Semantics: An International Handbook of Contemporary Research*, edited by A. Stechow and D. Wunderlich, 639-50. Berlin: de Gruyter.
- Lewis, David 1973. *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Lord, Errol 2011. 'Violating Requirements, Exiting from Requirements, and the Scope of Rationality'. *The Philosophical Quarterly* 61, 392-99.
- 2014. 'The Real Symmetry Problem(s) for Wide-Scope Accounts of Rationality'. *Philosophical Studies* 170, 443-64.
- MacAskill, William ms. 'Meta Decision Theory and Newcomb's Problem'.
- MacFarlane, John 2014. *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.
- Mackie, J. L. 1977. *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- Parfit, Derek 2011. *On What Matters*. Vol. 1. Oxford: Oxford University Press.
- Regan, Donald 1980. *Utilitarianism and Co-Operation*. Oxford: Oxford University Press.
- Sartre, Jean-Paul 1946/2007. *Existentialism is a Humanism* (trans. Macomber). New Haven: Yale University Press.
- Schroeder, Mark 2004. 'The Scope of Instrumental Reason'. *Philosophical Perspectives* 18, 337-64.
- 2009. 'Means-end coherence, stringency, and subjective reasons,' *Philosophical Studies*, 143: 223-48.

- 2010. ‘Ought, Agents, and Actions’. *Philosophical Review* 120, 1–41.
- Sepielli, Andrew 2009. ‘What to Do When You Don’t Know What to Do’. *Oxford Studies in Metaethics* 4, 5–28.
- 2014. ‘What to Do When You Don’t Know What to Do When You Don’t Know What to Do...’. *Noûs* 48, 521–44.
- Setiya, Kieran 2004. ‘Against Internalism’. *Noûs* 38, 266–98.
- Silk, Alex 2014. ‘Why “Ought” Detaches: Or, Why You Ought to Get with My Friends (If You Want to Be My Lover)’. *Philosophers’ Imprint* 14.
- Snedegar, Justin 2013. ‘Contrastive Semantics for Deontic Modals’. In *Contrastivism in Philosophy*, edited by Martijn Blaauw, 116–33. New York: Routledge.
- Titelbaum, Michael G. 2015. ‘Rationality’s Fixed Point (Or: In Defense of Right Reason)’. *Oxford Studies in Epistemology* 5.
- Way, Jonathan 2010. ‘Defending the Wide-Scope Approach to Instrumental Reason’. *Philosophical Studies* 147, 213–33.
- 2011. ‘The Symmetry of Rational Requirements’. *Philosophical Studies* 155, 227–39.
- Weatherson, Brian 2014. ‘Running Risks Morally’. *Philosophical Studies* 167, 141–63.
- Wedgwood, Ralph forthcoming. ‘Objective and Subjective “Ought”’. In Nate Charlow and Matthew Chrisman, eds., *Deontic Modality*. Oxford: Oxford University Press.
- 2006. ‘The Meaning of “Ought”’. *Oxford Studies in Metaethics* 1, 127–60.
- Wells, Ian ms. ‘Known Divergence’.
- Worsnip, Alex 2015a. ‘Narrow-Scoping for Wide-Scopers’. *Synthese* Advance online access, 1–30.
- 2015b. ‘Rationality’s Demands on Belief’. Doctoral dissertation, Yale University.