

Making Sense of the Moral ‘Must’

Avner Baz

avner.baz[at]tufts.edu

Abstract

I offer a critique of the dominant representationalist understanding of the moral ‘must’ and argue for an alternative understanding that is *second-personal* and *performative*. The representationalist understanding, I argue, faces serious *theoretical* difficulties, having to do with the nature of the *necessity* that is supposedly referred to by the moral ‘must’; and it is also *morally* problematic, in that it encourages us to suppose that utterances of the form ‘N (morally) must φ ’ may be understood, and their truth assessed, altogether apart from such morally significant matters as the nature and history of the relationship between the speaker and her addressee(s) or the illocutionary force of the utterance. The alternative understanding dissolves the theoretical difficulties faced by those who have tried to vindicate a representationalist understanding of the moral ‘must’. It is also *morally* superior, in that it underscores the dependence of the sense of sentences of the form ‘N must φ ’ – as uttered in moral contexts, and first and foremost in the second person – on morally significant contextual features such as who addresses whom, and with what illocutionary force, and what puts the first person in a position to address the other with these words and in that way.

Introduction

The topic of this paper is the sense of utterances of the general form ‘N must φ ’, when made in moral contexts – in contexts, that is, in which ‘morally speaking’ is implied. I take the difference between the moral ‘must’ and ‘must’ as used in other sorts of contexts and senses – epistemic, instrumental, and so on – to be clear enough. And I focus on the moral ‘must’ (and ‘have to’), and not on other moral ‘modals’ such as ‘should’ or ‘ought to’, for ease of exposition. My argument – both the critique of extant accounts and the alternative account I’ll be proposing – may fairly straightforwardly be extended to those other moral modals, however, and may in fact be bolstered by such an extension.

In contemporary analytic, English-speaking philosophy, the dominant understanding of utterances of the form 'N must φ ' in moral contexts is *representationalist* (or descriptivist), in the sense that they are taken to have a 'semantic content', or to express a 'proposition', which is aptly assessable in terms of truth and falsity, and which is separable, at least in theory, from whatever is being done with the words, illocutionarily speaking. On that dominant, 'truth-conditional' understanding, it is *by virtue of* having some such theoretically separable and identifiable 'semantic content' that words of that general form may competently be used to *urge* someone, for example, or *advise* (encourage, implore, beseech...) her, to φ . The sense of 'must' in such utterances – the contribution *it* makes to (the determination of) their 'semantic content' – is supposed to be representational as well, on that understanding, and to be a matter of its referring to a moral 'requirement' or 'obligation' to φ (or a moral 'prohibition' on, or the moral 'impermissibility' of, not φ ing):¹ if such a moral 'requirement' or 'obligation' exists – under the circumstances and at least as far as N is concerned, or more generally – 'N must φ ' would be true, on that understanding. Otherwise, it would be false.

Though enjoying the status of near-dogma in contemporary analytic philosophy, the representationalist understanding of the moral 'must' – on any of its extant variants – faces significant theoretical difficulties, as we will see. It is also *morally* problematic, I will argue, in that it encourages us to suppose that utterances of the form 'N must φ ' may be understood, and their truth assessed, altogether apart from such morally significant and even crucial matters as the nature and history of the relationship between the speaker and her addressee(s) or the illocutionary force of the utterance.

After presenting the representationalist understanding and the theoretical and moral problems it faces, I will propose an alternative understanding, on which the second-person form – 'You must φ ' – is primary, and the first- and third-person forms should be understood as secondary, or derivative, and on which the sense of the moral 'must' is not separable from the illocutionary force of the utterance.² That alternative understanding, I will argue, is not just

¹ For a recent expression of the common notion that natural language sentences of the form 'N must (ought to, should, have to) ϕ ' are used to refer to, or express, the 'deontic category' of the 'required/obligatory/mandatory/compulsory', see Berker (2022: 24).

² It should become clear as I go on that my invocation of 'the second person' bears little substantive affinity to Stephen Darwall's (2006). If anything, the connection is this: whereas Darwall invokes the

theoretically superior, in that it dissolves the main difficulties faced by those who have tried to vindicate a representationalist understanding of the moral ‘must’; it is also *morally* superior, in that it underscores the dependence of the sense of sentences of the form ‘N must φ ’ – as uttered in moral contexts, and first and foremost in the second person – on morally significant contextual features such as who addresses whom, and with what illocutionary force, and what puts the first person in a position to address the other with these words and in that way.³

1. The Dominant, Representationalist Understanding of the Moral ‘Must’

The dominant understanding of the moral ‘must’ (and of other moral modals as well) derives from a dominant representationalist, ‘truth-conditional’ understanding of modals in general. Referring to that understanding as ‘the canon’, von Fintel and Gillies summarize it as follows:

[M]odals are context-dependent quantifiers over a domain of possibilities. Different flavors of modality correspond to quantification over different domains of possibilities. Logical modalities quantify over all the possibilities there are, physical modalities over possibilities compatible with the laws of physics, deontic modalities over possibilities compatible with what ought to be [...] and] epistemic modals [...] quantify over possibilities compatible with what is known (2011: 108).⁴

second-person in an attempt to substantiate and ground the sort of moral ‘obligation’ or ‘commitment’ that on the representationalist conception is supposed to legitimize the use of the moral ‘must’, my invocation of the second-person is meant to show that that sort of legitimization of the moral ‘must’ is not needed.

³ To be sure, when a theory or account of moral discourse is ‘morally’ problematic in the way I will try to bring out, it is *thereby* rendered theoretically problematic as well. I believe that the distinction I’m drawing between ‘theoretical’ and ‘moral’ problems with representationalist accounts of moral ‘modals’ is nonetheless useful and apt. Useful, because the theorists I’m critiquing would likely recognize the problems I call ‘theoretical’ as ones they, or their theory, ought (ultimately) to address, and they’d likely consider the problems I call ‘moral’ as falling outside the scope of the ‘semantics’ (and perhaps even ‘pragmatics’) of moral ‘modals’, and therefore outside the scope of their theorizing; so there’s a sense in which I’m playing along with a distinction they tacitly make and rely upon, with the aim of showing that properly addressing the ‘moral’ problem with their account would also dissolve the ‘theoretical’ problems it faces. And apt, because there is no obvious or straightforward way in which the shortcomings I call ‘theoretical’ could manifest themselves in morally problematic *practice*, whereas there is a pretty obvious and straightforward way in which the shortcomings I call ‘moral’ would manifest themselves in morally problematic practice: in a word, they’d manifest themselves in (what I will call) *moralizing*.

⁴ See Stalnaker 1999: 35-6, for a similar summary of the dominant understanding of modals.

The representationalist understanding summarized here largely follows Angelika Kratzer's influential theory, on which modals in general are context-dependent quantifiers over "possible-worlds". The sense of modals, on Kratzer's account, is "relative to a conversational background" (2012: 21; see also 31), which, in the case of "deontic" modals – to which moral modals belong – consists, in part, of "a body of laws or regulations" (2012: 37). That background of laws or regulations serves as "an ordering source" that effects an order among all "accessible" worlds – worlds (deemed) possible from the perspective of those engaged in the conversation, or of their world) – in accordance with how close they each come to some (moral) "ideal". And a proposition such as $N \varphi$ is "a necessity" – making ' N must (or has to) φ ' true (in a context) – "just in case it is true in all accessible worlds that come closest to the ideal determined by the ordering source" (in that context) (2012: 40; see also 38).

Now, it is clear why a philosopher who wishes, as Kratzer does, to hold on to the Lewisian idea that modals have truth-conditions that may be cashed out, or explicated, in terms of "possible worlds" over which they "quantify", would need to incorporate in their account of moral modals the added complication of (a mechanism for) ordering worlds according to how close they come to some moral ideal – or, as von Fintel and Gillies put it, to "what [morally] ought to be" – as a way of distinguishing the moral use of those modals from other uses (instrumental, legal, epistemic, and so on). But even if sense could be made of the idea, which Kratzer hardly even sketches, of such moral ordering of worlds, or possibilities, it seems clear that in most cases in which an utterance of the form ' N must φ ' was either made or assessed in an everyday context, there would be morally excellent (and accessible) worlds in which N did not φ ; so there is reason to suspect that the "possible worlds" apparatus – if it isn't altogether circular, as suggested by von Fintel and Gillies's appeal to "possibilities compatible with what ought to be" – is not doing any real work when it comes to our understanding and assessment of utterances of the form ' N must φ ', and that (sufficient) closeness to a moral ideal is here just a matter of whether N 's φ ing is (deemed) necessary for (the attainment of) some such ideal, which presumably is determined by the contextually-relevant "laws or regulations". Whatever one thinks of the general Lewisian strategy of explicating modality, or domesticating it, in terms of "quantification over possible worlds", there is reason to suspect that in the case of moral modals the original difficulty of accommodating modality within a broadly

representationalist-referentialist conception of linguistic sense has not been overcome, but rather has been transferred over to some rather obscure and under-theorized mechanism for the moral ordering of worlds.

And indeed, in Alex Silk's more recent account of moral modals – which holds on to the representationalist assumption that utterances of the form 'N must φ ' are "statements" that "assert something about the world" (2017: 217) and, as such, have "ordinary representational contents" (2017: 214) and a "truth-value" (2017: 209) – the "possible worlds" apparatus is dropped, and the truth-value of utterances of the form 'N must φ ' is taken to depend on whether N's φ ing "follows from" (2017:218 and 219), or is "required" or "implied" by (2017: 206-207, and 222) a "contextually relevant", or "contextually supplied" or "determined", "body of norms" (2017: 207, 210, 219, and 234). The moral *necessity* supposedly referred to or invoked in such utterances is, on Silk's account, nothing more nor less than "logical"; and "the truth-conditional contents of deontic modal sentences are propositions about logical relations (e.g. implication, compatibility) between propositions [e.g., N φ s] and premise sets [i.e., the contextually-relevant bodies of norms]" (2017: 226).

2. Theoretical Problems with the Dominant Understanding

The representationalist understanding of 'N must φ ' faces a series of theoretical problems. There is, first of all, the skepticism – going at least as far back as Elizabeth Anscombe's "Modern Moral Philosophy" (1958), and arguably as far back as Kant's attempt to validate morality in the face of precisely that sort of skepticism – about the very idea of specifically *moral* requirements, when those are not, or no longer, understood as grounded in God's commands or moral authority. When contemporary philosophers tell us that N must ϕ , because ϕ ing is morally "required", or that N must not ϕ , because ϕ ing is morally "forbidden" or "impermissible", they seem to overlook the fact that a question that would naturally and reasonably arise in other contexts in which requirements, permissions, prohibitions, and so on, were invoked – namely, 'By whom, and with what authority?' – has no plausible answer when it comes to morality.⁵

⁵ For Anscombe, that meant that "the emphatic, 'moral', ought" (1958: 13) – and she surely would have said the same thing about the emphatic, moral 'must' – is empty: carrying, as she puts it, a certain

Contemporary philosophers such as Kratzer and Silk essentially bypass that general skepticism, by invoking a background of moral “norms” (rules, requirements, regulations, imperatives...) from which φ ing is supposed to “follow” logically, and by leaving it to the utterance’s “context” to determine the identity of those “norms” (cf. Silk 2017: 226). That, however, leaves the *moral status* of those “norms” unaccounted for, and thereby leaves the skepticism intact: are the “norms” themselves supposed to be somehow required or obligatory? If so, then how is *that* requirement or obligation to be understood, and how would we avoid a vicious regress – of precisely the sort that Kant thought his moral theory managed to avoid (see Kant 2012: 4:444) – of norms-requiring-norms-requiring-norms (or laws-securing-the-bindingness-of-laws-securing-the-bindingness-of-laws)? But if the norms themselves are not required or obligatory, then what *is* their moral status? I will later propose that moral norms – and more broadly the cares and commitments against the background of which modals such as ‘must’ do their work in ordinary moral discourse – are *inherently contestable*, and in general neither more nor less secure than the particular moral utterances they might be thought to support. And I will further propose that, far from constituting a theoretical problem, that inherent contestability is part of what gives utterances featuring moral modals their point in ordinary moral discourse. But in order to be able

“mesmeric force”, or “strong psychological effect”, but no “meaning” (1958: 8). On Anscombe’s diagnosis, the emphatic moral ‘ought’ relies for its seeming sense on a concept of “obligation” that has lost its “root” in a religious form of life and adherence to a divine law (1958: 6; see also 1). Anscombe proposes that we (moderns) would do well to “jettison” or “drop” that problematic use of ‘ought’ altogether (1958: 1 and 8) and content ourselves with using the word “in a nonemphatic fashion” (1958: 15). The ordinary (non-emphatic) ‘ought’, on Anscombe’s understanding, where it isn’t simply instrumental, is broadly Aristotelian, in that its ultimate basis is (what is deemed partly constitutive of) human flourishing or well-being (cf. 1958: 5). She acknowledges, however, that we, in contrast with Plato and Aristotle, lack any clear and shared conception of human flourishing, and calls upon philosophers to seek to develop and defend one (1958: 18). Here she seems to me to have been morally and philosophically complacent, however; for I take it that, contrary to what Plato and Aristotle may have supposed, what constitutes human flourishing, or well-being, is *inherently contestable*, with philosophers enjoying no special authority on that matter. I will later propose that, rather than rendering problematic our ordinary and normal use of words such as ‘must’ and ‘ought to’ in moral contexts, the possibility – and indeed undeniable reality – of fundamental moral disagreements, together with the inherent contestability of our most basic morally significant cares and commitments, is part of the background against which those words make the particular sense they make in such contexts. But in order to bring *that* clearly into view, one would need to break with modern moral philosophy even more radically than Anscombe did, and give up the objectivist, representationalist-referentialist understanding of moral modals.

to appreciate that, we'd need to give up the representationalist understanding of such utterances.

This leads us to another, related problem for the representationalist understanding of moral modals, having to do with the identity of the contextually relevant norms: what determines *which* norms are “contextually relevant”, or pertinent for an understanding and assessment – in terms of truth and falsity – of some particular utterance of the form ‘N must φ ’? In most of the examples Kratzer uses to illustrate and motivate her account of “deontic” modals, there is either some clear communally-accepted norm or rule that specifically requires φ ing (cf. 2012: 4–5), or there is someone in a position of authority who issues, or may issue, the requirement to φ (or not to φ) (cf. 2012: 56). But that is precisely what we don’t normally have when ‘must’ is used in its moral sense; and Kratzer herself notes that “realistic utterance contexts rarely provide unique modal bases or ordering sources for modal constructions”, and that speakers therefore “need to be able to manage a high degree of context-dependency and vagueness” (2012: 68; see also 32). Silk, on the other hand, while assuming, as we saw, that the utterance’s context does “supply”, or “determine”, a body of norms which, in turn, either require or do not require N’s φ ing, notes that “there is a range of factors which may be relevant to fixing the value of the [contextually-relevant body of norms] – speaker intentions, previous utterances and discourse moves, information structure, features of the concrete conversational situation, substantive normative principles, etc.”, and adds that “how these factors, whatever they are, interact to determine the relevant value is plausibly highly complex” (2017: 234-5). And since “one cannot evaluate a deontic modal utterance as true or false without making assumptions about what value for the deontic premise set variable is determined by the concrete discourse context in which the utterance was produced” (2017: 234), and since even “[d]etailed descriptions of concrete discourse contexts will likely fail to specify all the contextual features that might be relevant to determining the value [of that premise set]” (2017: 235), it is only to be expected that even fully competent speakers and moral agents would disagree with each other in what they take the relevant body of norms to be in any given context in which an utterance of the form ‘N must φ ’ was made or assessed, and would therefore disagree with each other in the truth-value they assign to that utterance (Silk 2017: 235). “Nothing less than a stipulation of [the relevant body of norms] may suffice for delivering truth-value judgments that are stable across speakers and reflect genuinely

semantic competence with deontic modals”, Silk concludes (2017: 235). –But if the representationally-understood “semantic content” of deontic utterances, and therefore their truth-value, are genuinely indeterminate in most everyday cases, as Kratzer seems to propose, or at any rate are commonly impossible for competent speakers to determine with even a modicum of certainty, as Silk readily concedes, then what justifies the widespread assumption that those are nonetheless key to the understanding and assessment of such utterances? Why hold on to a representationalist understanding of moral modals from which it follows that their sense (“semantic content”) is beyond the ken – either in principle or at any rate in practice – of their competent employers, who can, at best, only “*assume* that the conversational situation determines a value for [a set of moral norms] that would make the utterance appropriate” (Silk 2017: 217, my emphasis)?

These questions do not so much as arise for the philosophers we’ve discussed. They proceed as though truth-conditional semantics for modals – with or without the “possible-worlds” apparatus – were the only game in town. von Fintel and Gillies refer to it as “the canon”, as we’ve seen. Silk refers to the formal semantics for modals his account presupposes as the “standard” (2017: 117-8) and describes it as “independently motivated” (2017: 207), but does not himself offer any independent motivation for it. And Kratzer, recounting what had led her to begin developing her account of modals (and conditionals), simply says that she “was interested in the meanings of modals and counterfactuals, and *hence* in truth-conditional semantics” (2012: 1, my emphasis). I’m going to argue that the dominant, truth-conditional understanding of moral modals has a viable alternative, and therefore is not, or should not be, the only game in town.⁶ But first I want to explain why I consider that dominant understanding to be, not just theoretically problematic, but morally problematic as well.

⁶ Within mainstream, English-speaking, analytic philosophy, one often encounters appeals to the “success” of the truth-conditional semantics research program, as a reason to stick with that program despite whatever difficulties it may face, and as justification/explanation for its institutional dominance. Thus, for example, we find Gabe Dupre recently claiming that truth-conditional semantics has “massively increased our understanding of the behaviour of linguistic expressions such as quantifiers, modals, attitude reports, etc.”, and therefore is worth defending against various “contextualist” challenges (2020: 921). Part of what I’m trying to show in this paper is that, at least in the case of moral modals, that recurrent claim of success does not withstand close scrutiny. In previous work (Baz 2012 and 2017), I have argued that the same was true of truth-conditional, and more broadly representationalist-referentialist, accounts of utterances featuring ‘know(s)’.

3. The Moral Problem with the Dominant Understanding

To see why the dominant understanding of moral (“deontic”) modals is not just theoretically problematic, but morally problematic as well, consider an example Silk uses to illustrate moral disagreement, and more specifically disagreement over what someone must morally do. It consists of the following exchange:

Alice: Sally must give 10 percent of her income to the poor.

Bert: No, Sally doesn’t have to give that much. She can give less. (Silk 2017, 212).

And here (with a few minor abbreviations and elucidatory additions in square brackets) is Silk’s ultimate gloss on this example:

Alice and Bert are discussing our moral obligations to the poor. They consider the case of Sally. Alice utters “Sally must give 10 percent of her income to the poor.” Upon hearing Alice’s semantically underspecified utterance, Bert might (tacitly) reason roughly as follows (where *g* is the proposition that Sally gives 10 percent of her income to the poor): “Alice is intending to say something about the possibility *g*. In order to do so, given the grammatical properties of modals, a set of premises must be contextually supplied. Since Alice wouldn’t intend to say something false, she must be assuming a premise set *P* [in the form of a body of norms] that implies *g*. The current question under discussion concerns the extent of Sally’s moral obligations to the poor. Since Alice is cooperative, her utterance [...] must be relevant and realize an intention to provide at least a partial answer to this question. Assuming [a body of norms that logically implies Sally’s giving 10 percent of her income to the poor] would do so by ensuring that the moral norms endorsed in the conversation require Sally to give 10 percent. So, Alice must be assuming [some such body of norms], and have meant that [it] implies *g*.” [...] The appropriateness of Alice’s linguistic act of uttering [‘Sally must give 10 percent of her income to the poor’] requires that the moral norms operative in the context imply *g*. Since it’s mutually presupposed that Alice is obeying the [Gricean] conversational maxims [...], in uttering [that sentence] Alice implicitly proposes that it become taken for granted that such norms be accepted in the conversation. In accepting an utterance one normally accepts what the speaker committed to in uttering it. So, since it’s common knowledge that Alice can expect Bert to undergo an abductive reasoning process like [the one described above], it’s also common knowledge that he will object if he has relevantly different moral views, given their common goal of settling on what moral norms to accept. So if Bert doesn’t object, this will confirm that the context is as the appropriateness of Alice’s act requires, and the discourse-level

moral norms parameter can be set to a value that implies *g*. However, since Bert has incompatible moral views, he objects [...] For reasons parallel to those above, his doing so is appropriate only if the discourse-level moral norms are compatible with Sally's giving less than 10 percent. As he expects, Alice goes through an analogous abductive reasoning process and infers that he must wish to take for granted that the discourse-level moral norms are that way. If Alice accommodates, and accepts Bert's justification for his denial, it can become taken for granted that the context is as their present actions mutually require. If she doesn't, further negotiation may ensue [...]. In using deontic modals Alice and Bert can [thus] exploit their mutual grammatical knowledge, along with general pragmatic reasoning, to manage their assumptions about the conversational situation itself. (Silk 2017: 221–3)

Setting aside the theoretical difficulties discussed in the previous section, having to do with the determination or identification of the contextually relevant body of norms from which Sally's giving 10 percent of her income to the poor is supposed to follow (or not to follow) logically, and having to do as well with the moral status of those norms, I do wish to point out one further theoretical oddity that Silk's example makes vivid: can it truly be believed that any realistic body of norms – or even, more broadly, normative background – would *logically imply* a moral requirement with *that* level of specificity? I find it hard to imagine some such realistic body of norms or normative background. Amending the example, by having Alice say instead that Sally must give *at least* or *about* 10 percent of her income to the poor, would only slightly reduce the oddity. We could lessen the oddity further, and (to my ear) make the imagined exchange more natural, by having Alice say, "I think Sally should give at least 10 percent of her income to the poor", and having Bert respond, "I don't think she should give that much"; but I do not know how Silk's account might be extended to "softer" moral modals such as 'should' or 'ought to', which presumably would require him to posit a relation *weaker* than logical implication between the normative background and the course of action under consideration; and once we've gone *that* far, we might anyway be readier to give up the representationalist, truth-conditional understanding of moral modals.⁷

⁷ I do not doubt that, with enough ingenuity and technical sophistication, someone committed to the framework of truth-conditional semantics *could* find any number of ways of accommodating the distinction between "weak" and "strong" moral modals within that framework. Given the fundamental problems discussed in this paper with extant attempts to account just for the "strong" moral modals within that framework, I doubt such ingenuity and sophistication would result in true insight into ordinary moral discourse and the functioning of moral modals therein.

None of that gets at what I find morally troubling about Silk's example, however. So let me ask: Why are Alice and Bert discussing how much of her income Sally must – or, for that matter, should – give to the poor? Has Sally asked for their advice? If not, do they think she would benefit from their advice anyway? Or is it perhaps that Alice is upset with Sally, for she feels Sally has not been sufficiently generous in her charitable contributions to the poor? If so, why focus on Sally, when presumably there are any number of other people who give even less? What, more generally, is the nature of Alice's and Bert's relationship with Sally? Why do they care about how much she gives to the poor? Do they otherwise care about her, and about what would benefit her, morally or otherwise? Or do they not care about her personally, and are merely using her, or her "case", as an example, in the course of a theoretical discussion of "our obligations to the poor"?

From Silk's perspective, and equally from that of Kratzer and others who have tried to explicate the sense of moral modals representationally and truth-conditionally, none of these questions is pertinent for our understanding of (the "semantic content" of) such utterances, and for our assessment of them in terms of truth and falsity. From their perspective, the sorts of questions I've just raised, while admittedly pertinent for a *moral* understanding and assessment of exchanges such as Alice and Bert's, may safely be set aside when considering the *semantics* of such exchanges, and how speakers "can use deontic modals as a way of testing one another's normative views, inviting them to object if they accept different norms" (Silk 2017: 223), and to thereby "manage an evolving body of norms" (Silk 201: 229).

I have no a priori argument for why attempts such as Silk's or Kratzer's to separate theoretically the semantics (and certain portions of the pragmatics) of moral modals from morally significant aspects of their employment are doomed to failure.⁸ Our competence with these words, which surely is acquired and practiced as one whole package, *might have been* modular in the way Silk, Kratzer, and many others in effect presuppose.⁹ I contend that we've not been given any reason to suppose that it is, however. Nor have we been given a sufficiently worked-out, plausible story about how it could be. In the next

⁸ After all, there *are* possible stretches of human discourse that might well be morally significant – e.g., 'I gave her nothing', or 'I will teach him a lesson' – where such a separation seems at least plausible.

⁹ Though I note once again that Silk and Kratzer have not shown us that it is, and that by their own lights, the presumably-theoretically-separable truth-conditional "semantic content" of utterances featuring moral modals is normally beyond the ken of their competent makers and addressees.

section, I will offer an alternative understanding of such utterances on which their content is not separable from their illocutionary force and from their moral significance more broadly, and which at the same time dissolves the theoretical difficulties faced by the representationalist understanding.

4. The Second-Personal, Performative Understanding of the Moral 'Must'

The alternative understanding I would like to propose of the moral 'must' (and of moral modals more generally) proceeds from the idea that our vocabulary has evolved *in* and *for use*, and that our moral vocabulary in particular has evolved to enable us to *position ourselves significantly in relation to others*, rather than merely to pronounce, truly or falsely, on what anyone and everyone morally must (should, ought to...) do, or not do (whatever merely pronouncing on *that* might be thought to mean). The understanding I propose is accordingly *second-personal and performative*. It is *second-personal*, in that it takes the second-person form – namely, 'You must φ ' – as primary and paradigmatic, and takes the third- and first-person forms as derivative, in the sense that what we say some third person morally must (should, ought to...) do is to be understood as meaning that that is what *we would say to them* they must do, and what we say *we* morally must (should, ought to...) do is to be understood as meaning that that is what we, or what we think someone we morally trust, would *say to us* we must do.¹⁰ And it is *performative* in that it does not take the content of the moral 'must' – how that word, as used in moral contexts, is to be understood; what contribution it makes to the overall sense of an utterance of 'You must φ ' – to be separable theoretically from the illocutionary force of the utterance.

¹⁰ Compare Fricker's proposal that the "basic" case of blame is second-personal, and that other forms of blaming (first- and third- person) may usefully be understood as derivative from it (2016: 167). The primarily second-person use of the moral 'must' I'm seeking to elucidate should be distinguished from that which refers to what Williams calls "moral incapacity" (1993; see also Winch 1987; Gaita 2004: 90; and Taylor 2012: 113-8), which expresses itself paradigmatically in *first-person* proclamations (as in Luther's "Here I stand, I can do no other"), has no natural second-person correlate, and in the third person can only be *reported*, or otherwise objectively represented ('Luther (found he) could do no other'). Or let me put it even more clearly and precisely: the use of 'You must ϕ ' I'm trying to elucidate is not the second-personal correlate of the 'I must (=can do no other than) ϕ ' that Williams and others have emphasized and discussed. Williams himself underscores the distinction, when he notes that there is also a moral use of 'can (not)' – as in 'You can't do this!' – that is not expressive of what he calls "moral incapacity", and belongs with the moral modals in what he refers to as their "normative" use (1993: 67). *That's* the use that Kratzer, Silk, and others have sought to understand representationally, and the one I'm seeking to elucidate in this section.

The force of the moral ‘must’, and of the other moral modals as well, is not, on my proposal, “logical”, as Kratzer and Silk make it out to be, but rather, precisely, illocutionary; and the typical illocutionary force of utterances in the second-person about what the addressee must (should, ought to...) do is not that of simply *telling* that person, or *informing* her of, what she morally must (should, ought to...) do – as if there were something intelligible to tell or inform another by means of these words – but rather that of *urging* or *calling upon* her to do this or that, or else (as is typically the case with ‘should’ or ‘ought to’ utterances) that of *advising* her to do this or that, or *recommending* or *suggesting* that she do it. And that means, crucially, that caring about the agent-addressee, and about what would morally benefit her, is essential:¹¹ it is part of what puts one in a position to felicitously perform any of these illocutionary acts, and what distinguishes a moral address from a moralizing one.¹²

The background against which utterances of ‘You must (should, ought to...) ϕ ’ have their sense in moral contexts is (what are taken to be) the agent-addressee’s relevant cares and commitments, which may well include, but need not be limited to, communally-held norms or universal (or arguably universal) moral principles. And let me also emphasize that, on the understanding here proposed, the agent-addressee has no overriding authority with respect to their cares and commitments: others may see more clearly, and help the agent-addressee come to see more clearly, what, at the end of the day – or anyway at the end of that day – those are. That the relevant cares and commitments must ultimately (reasonably be taken to) be the agent-addressee’s marks another crucial difference between my proposal and accounts such as Kratzer’s or Silk’s: while the three of us agree that moral modals are “inherently relational” (Kratzer 2012: 8), in the sense that they require, for their sense on particular occasions, a suitable “conversational background” (Kratzer 2012: 31), I do not take that background to be there objectively and impersonally – independently of the addressee’s cares and commitments as seen and understood by the speaker. If our ‘You must (should, ought to...) ϕ ’ does not make the right contact with those cares and commitments, if it does not resonate with them

¹¹ I take the contrast between what benefits someone *morally* and what benefits her in some other way (financially, medically, athletically, in terms of the advancement of her career, and so on) to be clear enough, notwithstanding the former’s undeniable and unavoidable vagueness. And of course, something *could*, under suitable circumstances, benefit someone morally by virtue of benefitting her in one of those other ways. I thank Noah Greenstein for pressing me on this issue.

¹² I elaborate on that distinction in Baz 2022.

in the right way, it is empty in just the way Anscombe says it is: it may still carry a certain “mesmeric force”, or “strong psychological effect” – which may suffice if all we mean is to moralize – but no “meaning” (Anscombe 1958: 8). By ‘moralizing’ I mean, precisely, using moral language in order to get the addressee to do, think, or feel something, or else just in order to assert moral authority or superiority and feel good about ourselves, but without respecting their cares and commitments – and so without respecting them – in our address. The moralizing use of moral language is thus parasitic on its non-moralizing use.¹³

The agent-addressee’s relevant cares and commitments *might* sometimes be taken as simply *given* – that is, as clear enough for present intents and purposes, and as beyond dispute;¹⁴ but another crucial aspect of my proposal is that they need not thus be taken. They are inherently contestable, as I’ve suggested earlier, and far more pervasively so than Silk’s account allows us to realize. For what’s contestable is not just what the other person’s cares and commitments (ultimately) are, but also what they each come to, how they stand relative to each other, what (taken together) they morally imply (in general and on *this* particular occasion), and whether they are beyond dispute – all that need not be taken, and regularly is not taken, as already given, or settleable by impersonal “logic”: it may be, and regularly is, part of what is under investigation and assessment in moral conversations, part of what we’re trying to figure out together.¹⁵

I’m proposing that ‘You must ϕ ’ may, and regularly does, mean something like, ‘Given your cares and commitments as I see and understand them, I see no reasonable alternative to your ϕ ing, or no alternative you could live with, and therefore am urging you to ϕ ’; and ‘You ought to ϕ ’ may, and regularly does, mean something like ‘I recognize that there are other things you may legitimately, or not unreasonably, do, but given your cares and commitments as I see and understand them, and out of concern for what would morally (or otherwise) benefit you, I’m advising you to ϕ ’.¹⁶ These proposed paraphrases

¹³ I elaborate on that idea in Baz 2022.

¹⁴ That might be the case when, for example, addressing ourselves to the firemen in front of a burning building, we call out to them, ‘You must do something!’, or ‘You must save them!’.

¹⁵ Here I’m drawing on Cavell’s account of moral conversation in Cavell 1979.

¹⁶ To the best of my knowledge, Sloman (1970) was the first to propose that ‘ought’ typically implies the existence of real, or not unreasonable, alternatives to the recommended course of action, whereas ‘must’

are no more than rough approximations of what we ordinarily and normally are saying, in effect, when we confront each other with what the other must (should, ought to...) do. Their aim is to bring out three important features of the ordinary use of moral modals. The first is that these words, as ordinarily used in moral contexts, function not representationally or descriptively, but as what might be called ‘illocutionary force indicators’, where the force in question could be made explicit by means of locutions such as ‘I’m urging you’, ‘I’m advising you’, ‘I’m (strongly) recommending that you...’, and so on. The second is that acknowledgement of, and respect for, the other’s cares and commitments – which means acknowledgement of and respect for that other person – are essential: they are, as I’ve said, what distinguishes genuine moral conversation from mere moralizing, and what is required for endowing the utterance – and specifically the moral modal employed in it – with sense (as opposed to mere psychological force).¹⁷ The third feature is the internal relation between the sense of an utterance of the form ‘You must (should, ought to...) ϕ ’ in moral contexts and the background of cares and commitments against which it has that sense: ‘You must (should, ought to...) ϕ ’ has no more content than may be given to it by that background and how it bears on that person’s ϕ ing, or not ϕ ing; and a person’s acceptance, or else rejection, of an utterance of that form contributes to the ongoing determination, and articulation, of that person’s cares and commitments.¹⁸

typically rules out such alternatives. Nothing of importance hangs, for the purposes of this paper, on whether that distinction is as clear, and clear-cut, in ordinary and normal practice, as Sloman suggests.

¹⁷ An anonymous reviewer suggested adding the qualification that “the cares and commitments in question are ones that the speaker him/herself can embrace, not just ones he/she attributes to the listener.” I find ‘embrace’ or even just ‘can embrace’ a little too strong. I want to allow for situations in which at least some of the other’s (relevant) cares and commitments are not, or even could not be, embraced by the speaker. ‘(Can) respect (and accept)’ seems to me just right; and we may well be able to respect and accept (some of) the other’s differences from us. Needless to say, if the other’s (contextually relevant) cares and commitments, as you see and understand them, are not ones you can respect, you could still try to get her to do or not to do something, but telling her that she must (should, ought to...) do, or not do, that something, would be semantically empty – though possibly still psychologically effective – in just the way Anscombe describes. Compare Gaita: “To say of Hitler [...] that he ought not to have murdered the Jews and others strikes me as fatuous. He is, and was at the time, beyond the epistemic reach of such a remark. It therefore idles in much the same way that thoughts or statements that a person ought to do something idle if that person lacks the relevant interests or desires” (2004: 94).

¹⁸ This may sound similar to Silk’s account of how speakers are “using deontic modals [...] to manage an evolving body of norms” (2017, 229); and at some level of abstraction, it is. But shifting the focus, as I have proposed, from impersonal truth-conditions, and truth value, to the illocutionary force of the

Insofar as the other's cares and commitments – including what they mean or come to, what they imply, and whether they are beyond dispute – are not taken as given, the 'You must (should, ought to...) ϕ ' may not reasonably or plausibly be understood as expressing the independently-intelligible conclusion of a practical syllogism that has these cares and commitments, or some subset of them, as its independently-standing and independently-intelligible major premise. Rather, it is used for presenting the connection we see between the other's cares and commitments (as we see and understand them), the situation she faces (as we see and understand it), and her ϕ ing, and for making explicit how we stand in relation to that perceived connection, and to that other person. Rather than situating the other's ϕ ing in an objectively existing space of moral requirements, permissions, and prohibitions, such utterances situate us in relation to her – within, and against the background of, a more or less (as we are liable to find out) shared "moral universe", as Stanley Cavell calls it (1979: 268). Their force, broadly speaking, is that of inviting the other to see things a certain way, not that of asserting how things morally stand. Pace Silk, my brother's calling our mother to apologize for something he's said, for example, need not, and normally would not, logically follow from, or be logically implied by, any set of "norms" or "premises" I might reasonably adduce in support of my saying to him that he must (or should) call her. But so long as I said it in the right way, and out of concern for what would morally benefit him, my saying it could very well be, not true, but perfectly intelligible and apt.

If all we're trying to achieve is just to get the other person to do or feel something – perhaps in part by appealing to or invoking their cares and commitments, but without really caring about and respecting those cares and commitments, and without caring about that other person and what would morally benefit them – then we are moralizing, and, if so, are not likely to be troubled by the possibility that our 'must' ('should', 'ought to'...) has only mesmeric force, but no (clear) content. But if, in confronting another person morally – and it's important to keep in mind that the confrontation may well be friendly – we do care about them and respect them, and therefore care about and respect what they care about and are committed to; and if we neither take

utterance, and more broadly to the nature of the evolving relationship between speaker and agent-addressee, makes all the difference, morally speaking, and dissolves the theoretical difficulties that beset Silk's account.

nor present ourselves as having the authority to simply *tell* them what they morally must (should, ought to...) do – as if there were any such intelligible thing to tell; then it is no problem that the other's cares and commitments, as well as our own, are part of what's at issue, and part of what's under investigation and assessment, in the conversation. That which has seemed fatefully problematic on the representationalist, truth-conditional understanding of moral modals, is precisely what gives their employment its actual, everyday point, and suffices for giving utterances of the form 'You must (should, ought to...) ϕ ' all the sense we need them to have.¹⁹

References

- Anscombe, G. E. M., 1958. "Modern Moral Philosophy". *Philosophy* 33, 1–19.
- Baz, A., 2012. *When Words are Called For: A Defense of Ordinary Language Philosophy*. Cambridge, MA: Harvard University Press.
- Baz, A., 2017. *The Crisis of Method in Contemporary Analytic Philosophy*. New York: Oxford University Press.
- Baz, A., 2022. "On the Moralization of Moral Theory". *Mind* 131, 549–573.
- Berker, S., 2022. "The Deontic, the Evaluative, and the Fitting". In Howard, C. and Rowland, R. A. (eds.), *Fittingness: Essays in the Philosophy of Normativity*. New York: Oxford University Press.
- Cavell, S., 1979. *The Claim of Reason*. New York: Oxford University Press.
- Darwall, S., 2006. *The Second Person Standpoint: Morality, Respect, and Accountability*. Cambridge, MA: Harvard University Press.
- Dupre, G., 2020. "Idealisation in semantics: truth-conditional semantics for radical contextualists". *Inquiry* 66, 917–946.
<https://doi.org/10.1080/0020174X.2020.1712225>
- Fricker, M., 2016. "What's the Point of Blame? A Paradigm Based Explanation". *Nous* 50, 165–183.
- Gaita, R., 2004. *Good and Evil: An Absolute Conception* (2nd Edition). New York: Routledge.
- Kratzer, A., 2012. *Modals and Conditionals*. New York: Oxford University Press.
- Silk, A., 2017. "Normative Language in Context". In Shafer-Landau, R. (ed.), *Oxford Studies in Metaethics* 12, 206–243.
- Sloman, A., 1970. "'Ought' and 'Better'". *Mind* 79, 385–394.
- Stalnaker, R., 1999. *Context and Content*. New York: Oxford University Press.
- Taylor, C., 2012. *Moralism: A Study of a Vice*. Ithaca, NY: McGill-Queen's University Press.

¹⁹ I thank two anonymous reviewers and the editors for several very helpful comments and suggestions.

- von Fintel, K. and Gillies, A., 2011. “‘Might’ Made Right”. In Egan, A. & Weatherson, B. (eds.), *Epistemic Modality*. New York: Oxford University Press. 108–130.
- Williams, B., 1993. “Moral Incapacity”. *Proceedings of the Aristotelian Society* 93: 59–70.
- Winch, P., 1987. “Who Is My Neighbor?”. In Winch, P. (ed.), *Trying to Make Sense*. Oxford: Basil Blackwell.

Biographical Note

Avner Baz is a Professor of Philosophy at Tufts University. He has done work in moral philosophy, aesthetics, epistemology, philosophy of perception, and philosophy of language, with a special focus on the question of philosophical method. He published two books – *When Words are Called For* (Harvard University Press, 2012), and *The Crisis of Method in Contemporary Analytic Philosophy* (Oxford University Press, 2017) – that aim to dispel the widespread belief that the insights and procedures of ordinary language philosophy may safely be ignored by current practitioners in mainstream analytic philosophy. He also published two books on the topic of aspect perception: *The Significance of Aspect Perception* (Springer, 2020), and *Wittgenstein on Aspect Perception* (Cambridge University Press, 2020).