CANADIAN JOURNAL OF PHILOSOPHY, 2018
VOL. 48, NOS. 3–4, 405–424
https://doi.org/10.1080/00455091.2018.1432400

Routledge
Taylor & Francis Group

Two nondescriptivist views of normative and evaluative statements

Matthew Chrisman

Department of Philosophy, University of Edinburgh, Edinburgh, Scotland

ABSTRACT

The dominant route to nondescriptivist views of normative and evaluative language is through the expressivist idea that normative terms have distinctive expressive roles in conveying our attitudes. This paper explores an alternative route based on two ideas. First, a core normative term 'ought' is a modal operator; and second, modal operators play a distinctive nonrepresentational role in generating meanings for the statements in which they figure. I argue that this provides for an attractive alternative to expressivist forms of nondescriptivism about normative language. In the final section of the paper, I explore ways it might be extended to evaluative language.

ARTICLE HISTORY Received 31 December 2017; Accepted 10 January 2018

KEYWORDS Metaethics; expressivism; inferentialism; representationalism; modals

1. Introduction

Some metaethicists are impressed by the idea that normative statements (e.g. saying what someone ought to do) differ fundamentally from statements describing reality. This idea has founded a fecund research program in twentieth century metaethics including the development of emotivist, prescriptivist, projectivist, and (most recently) expressivist views. The sorts of expressivism currently on the market are highly sophisticated, inspiring great ingenuity even amongst their critics in thinking about the function of normative/evaluative language and the relations between this and normative/evaluative thought.[1] In my view, however, there is a different and better way to get to the idea that normative statements do not describe reality.

This alternative is based on the dominant view in deontic logic and formal semantics that ought-statements should be treated as expressing a type of modal judgment.[2] More specifically, in this literature, ought-statements are commonly

CONTACT Matthew Chrisman matthew.chrisman@ed.ac.uk

© 2018 Canadian Journal of Philosophy

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

regarded as 'weak' necessity claims, where context determines the flavor (moral, prudential, teleological, epistemic) of weak necessity. A different and older tradition in the philosophy of mind and language was impressed by the idea that there is a fundamental difference between representing empirical facts and making modal judgments concerning various ways in which things might necessarily or possibly be. For example, considering the content of judgments about the world as it can be represented in our experiences, Kant wrote, 'The modality of judgments is a quite special function of them, which is distinctive in that it contributes nothing to the content of the judgment' ([1787] 1998, 209). And considering the propositional content of judgments conceived of as representations of reality, Frege wrote, 'By saying that a proposition is necessary I give a hint about the grounds for my judgment. But, since this does not affect the content of the judgment, the form of the apodictic judgment has no significance for us' ([1879] 1967, 13).

Developing this tradition's insight about modality suggests treating ought-statements as nondescriptive *because they are modal* (rather than because they are emotive, prescriptive, projective, or expressive). This is the kind of view of 'ought' I favor as a foundation for a nondescriptivist view about normative thought and discourse.[3] Because of its contrasting explanation of why normative language is nondescriptive, I have been reluctant in my own work in metaethics to adopt the label 'expressivist' despite agreeing with many expressivist arguments against descriptivist views in metaethics. Kant and Frege didn't think modal statements express anything like prescriptions, noncognitive attitudes, or planning states, but they do seem to have regarded them as nondescriptive because they performed some function in our thought and talk different from describing reality. To draw out this contrast with the expressivist tradition in more detail, there are two questions I want to explore here.

First (and mainly), even if we agree with Kant that the modality of a judgment contributes nothing to the content of that judgment, or we agree with Frege that saying a proposition is necessary does not affect the content of the judgment that it is true, it is obvious that we should not say that 'ought' contributes nothing to the *meaning* of statements in which it figures. (For what it's worth, I think it is most charitable to interpret Kant and Frege as making claims about the *descriptive* content of these judgments.) This raises the question: if 'ought' considered as a modal doesn't contribute to the descriptive content of the statements in which it figures, how does it affect their meaning? (Sections 3 and 4).

Second (and much more tentatively), the expressivist tradition offers various unified explanations of why normative *and* evaluative statements are nondescriptive, but a modal-operator account of why ought-statements are not descriptions of reality does not extend naturally to evaluative statements. Words such as 'good', 'bad', 'better', and 'worse' are not plausibly understood as modal operators in any usual sense. This puts pressure on a metaethics based on the modal-operator account of ought-statements to abandon the natural assumption that normative and evaluative statements are nondescriptive for

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

basically the same reason, raising the question: what should metaethicists who follow me in thinking normative statements are nondescriptive because they are modal say about evaluative statements? (Section 5).

2. Normative and evaluative statements as nondescriptive; some historical lessons

When I say that some metaethicists, including myself, are impressed by the idea that describing reality differs from saying what anyone ought to do or evaluating things as better and worse, I think it is important to recognize that most of these metaethicists would recognize a prominent descriptive use of words such as 'ought' and 'good'. This was already hinted at by Ayer, who distinguished between the use of such words to describe the moral sense of a particular community and their use as 'normative ethical symbols.' In fact, on a careful rereading of *Language, Truth and Logic*, one can get the sense that Ayer regarded most uses of words such as 'ought' and 'good' as descriptive (1936). On his view, many uses of these words are sociological descriptions of the morals of some group of people or descriptions of the verdicts of some normative system that conversational participants are assuming for the sake of conversation. It's only when we get to the business of discussing what one *really* ought to do or what is *really* good that his verificationism about meaning led him to deny that the words carry descriptive content.

We shouldn't dwell on Ayer's verificationist reasons for making this distinction, but I do think we should recognize a use of words such as 'ought' and 'good' where one is, in effect, describing the verdicts of some set of norms or values applied to some case. We should set aside these uses of 'ought' and 'good' in attempting to explain why normative and evaluative statements differ from straightforward descriptions of reality.

With that distinction in place, Ayer famously argued that the use of 'normative ethical symbols' is purely expressive, neither adding to nor subtracting from the descriptive content of the rest of the statement in which they figure (if it already had one). This has seemed implausible to many metaethicists. How could a word be meaningful and yet have no effect on the descriptive content of statements in which it figures? But we have pretty clear examples of this in ordinary language. The word 'fucking' in 'The fucking kids trashed the park,' is plausibly viewed as a purely expressive word. It's not a qualification of 'kids', used to pick out a specific group of kids; rather it's a way for the speaker to express a negative attitude towards the kids who trashed the park. This statement still carries the descriptive content *that the kids trashed the park*, it's just made using a word that adds an extra expression of negative attitude towards the kids. On a charitable interpretation of Ayer, this is like what 'normative ethical symbols' do in most statements containing them.

Even when charitably interpreted, this idea is not credible, for many well-known reasons. The logical properties of normative and evaluative statements

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

seem to turn on their use of 'ought' and 'better' in a way that statements involving the expressive use of 'fucking' do not.[4] Nevertheless, in moving beyond Ayer's emotivism, we shouldn't reject the idea that words can figure meaningfully in statements without contributing to the descriptive content of those statements; and a statement containing one of these words can still be descriptive, even if it also performs some other discourse role in conversational dynamics (because of the nondescriptive word in it). Ordinary uses of the sentence 'The fucking kids trashed the park' are still plausibly regarded as describing something the kids did.

Hare's (1952) prescriptivism offers a radically different model for the role of normative and evaluative words. He thought they functioned somewhat analogously to markers for imperatival mood. A common account of the meaning of imperatives treats them as containing two elements: a descriptive content and a 'make-true' operator on this content. The idea is that imperatives are linguistically suited to prescribe rather than describe because they involve a semantic operation on a piece of descriptive content turning it from something that can be put forward as true to something that can be put forward as to be made true. For example, 'Kids, clean up the park!' might be said to carry the descriptive content *that the kids clean up the park*, but the imperative is not usable to describe reality in this way; rather it is usable to prescribe to the kids a complex action which would make this descriptive content true. This is why imperatives are often said to have satisfaction conditions rather than truth conditions. The situation is a bit more complicated for normative statements, according to Hare, but his core thought seems to have been to treat 'ought' in statements such as 'The kids ought to clean up the park' as operations on embedded descriptive content that make it usable to prescribe action which would make the embedded content true.

Even when charitably interpreted, this idea is also not credible as a thesis about normative and evaluative statements. Sentences deploying 'ought' and 'good' are linguistically embeddable in propositional contexts (e.g. under 'believes' and 'might') that do not embed imperatives. However, we shouldn't reject the thought that words can have their meanings in part because of how they function as operators on descriptive content rather than contributors to descriptive content. With Hare, we might want to say that such statements 'carry' descriptive content but they do not describe reality as matching this content. I think this lesson is crucial for making sense of the idea that ought-statements are nondescriptive because they are modal.

3. 'Ought' as modal operator

Some metaethicists might think that 'ought' is a descriptive word, describing a relation between agents and actions. For example, an ordinary use of 'You ought to call your mother' might be said to describe you as being obligated to perform

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

the action of calling your mother. But, even if that looks halfway plausible in this case, it cannot be right as a general thesis about the word 'ought'. There are many 'flavors' of 'ought', not all of which have anything to do with agents and actions; and even for a more relaxed version of the relational view to work, the ought-relation would have to be multifarious to the point of gerrymandering. No agent is plausibly related to an action in 'As they left an hour ago, they ought to be home by now.' And even if we think there are agents responsible, the relation that would be described by 'There ought to be no childhood starvation' would have to be quite different from the relation putatively described by 'You ought to call your mother.' Moreover, we make general normative evaluations about how people feel, as in e.g. 'One ought to feel sympathy for the bereaved' where it's very hard to see how this could be plausibly construed as describing a relation that is similar to any normative relation between an agent and an action.

In response to these counterexamples to the relational view, some have suggested that 'ought' is ambiguous – maybe sometimes describing a relation between agents and actions, other times describing a relation between agents and their attitudes, and still other times describing outcomes as highly likely in light of some implicit body of evidence or some state of affairs as overall best.[5] But the dominant view in semantics is that 'ought' – like other modal words – is not ambiguous but context sensitive More specifically, it is context sensitive in the way that many intensional operators are. These are usually treated as linguistic devices that shift the circumstances relative to which various pieces of embedded semantic content are to be semantically evaluated. Context provides input to the kind of shift that takes place.

This is not the place to get into the general theory of intensional operators, but it might suffice to say that allowing that some bits of language are non-extensional is crucial for making progress towards the ideal of compositionality in our theoretical models the semantics of natural languages; and the existence of operators capable of shifting the circumstances under which some embedded piece of content is to be evaluated semantically as part of compositional processing of language is at the heart of all model theoretic semantic accounts that have any hope of achieving the ideal of compositionality.[6] This is the semantic role of intensional operators, and modal words are commonly thought to be paradigm examples of intensional operators.

In philosophy, this idea is perhaps most familiar from discussion of the epistemic possibility modal 'might'. Consider a standard use of 'Sally might be at home.' Assume, vagueness aside, it is the case either that Sally is at home or that Sally is not at home. In this case, it appears misguided to say that 'might' in this statement describes a relation Sally stands in to being at home. Rather, the standard view is that this word operates on the propositional content it embeds (*that Sally is at home*), evaluating this propositional content as true in at least one set of circumstances consistent with the relevant body of evidence.

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

If we appeal to possible worlds to model various circumstances under which a propositional content can be evaluated, we could generate a simple

$$\text{MODAL RULE FOR 'MIGHT':}[[\text{might } p]]^{e}$$
$$= \text{T iff } p \text{ is true at some of the } e\text{-compatible worlds}$$

where *e*-compatible worlds are the worlds consistent with the relevant body of evidence. There is considerable debate about how the relevant body of evidence for might-statements is negotiated, but the more general idea is clear enough. The word 'might' functions not to represent a relation but rather to shift semantic interpretation of the proposition it embeds, in effect directing interpreters to evaluate this proposition for truth not at the actual world but at all possible worlds compatible with some body of evidence. For instance, with typical uses of 'Sally might be at home' the modal rule for 'might' suggests one evaluate whether the proposition *that Sally is at home* is true at all circumstances compatible with some body of evidence.

There are complexities in extending this intensional semantics to 'ought' that I don't want to get into here.[7] However, it's enough for my purposes below to fixate on a simplified intensional operator view about 'ought' that treats it as a necessity modal rather than a possibility modal and allows a two-factored determination of the set of worlds relevant for modeling the circumstances at which its embedded content is to be evaluated.[8] If we do so, we get something like this simple:

$$\text{MODAL RULE FOR 'OUGHT':}[[\text{ought } p]]^{fg}$$
$$= \text{T iff } p \text{ is true at all of the } fg\text{-compatible worlds}$$

where *f* and *g* are contextually negotiated parameters determining, respectively, the background conditions and an identification of top ranked worlds relative some ranking of worlds, e.g. in terms of moral ideality, prudential betterness for some agent, probabilistic likelihood given some evidence, etc.

For example, take a use of 'Peter ought to live in southern Spain' considered as a claim about achieving best quality of life. We could (as a first pass anyway) model semantic interpretation of this statement as follows:

p: Peter lives in Southern Spain

f: restrict to possible worlds where Peter lives in Europe, has a portable job, etc.

g: identify the remaining worlds that are prudentially best for Peter, where this is a resultant of lifestyle, food, opportunities for meaningful relationships, etc.

Then the ought-statement is true iff the *fg*-compatible worlds are ones where *p* is true.

The attraction of the modal rule for 'ought' is its unity, flexibility, and similarity to semantic rules proposed for other modal words. By treating 'ought' as a context sensitive weak necessity modal, we can predict fairly plausible truth conditions for all of the different flavors of ought-statements mentioned above

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

without massive gerrymandering. The relatively simple rule doesn't restrict application to agents and actions; and it can make sense of the different ways background conditions and ways of ranking things affect the semantic processing of various ought-statements in context. Moreover, assimilating words like 'ought', 'should', and 'must' to universal quantification over possibilia and words like 'might', 'may', and 'could' to existential quantifications over possibilia allows for an attractive explanation of what these words have in common qua modal auxiliaries.

4. 'Ought' as a modal and descriptive content

Going forward, let's assume the modal rule for 'ought' is roughly correct. How does that affect the issue of whether ought-statements are descriptions of reality? Here are two apparently competing answers:

First: if the modal rule is roughly correct, ought-statements have truth conditions. So, like other statements with truth conditions, they can be believed true and they can be known. This means ought-statements describe a way reality could be. To be sure, since 'ought' is not a relational predicate, ought-statements shouldn't be treated as *ordinary* descriptions of a relation. We need an alternative account of what they describe; we should develop an account of the pieces of reality the description of which grounds 'ought's semantic contribution to the sentences in which it figures. If we are realists about possible worlds and the relations amongst them, then we can say ought-statement describes a region of modal space: what's true in all of some set of possible worlds. If we're skeptical of modal realism, we can still recognize various real relations of normative necessitation and view ought statements as describing these.

Second: if the modal rule is roughly correct, 'ought' is a modal operator. Modal operators do not function to describe things in reality. This means that 'ought' does not add descriptive content to the statements in which it figures. But surely it has meaning. Ayer's expressive and Hare's prscriptive accounts provide two early and inadequate accounts of its meaning. However, Hare was closer to right: As an intensional operator, 'ought' is more like markers for imperative mood than the expressive use of 'fucking'. It functions to shift the role of the descriptive content it embeds, meaning that ought-statements 'carry' descriptive content but do not put this forward as a description of reality. The idea that ought-statements sometimes *prescribe action rather than describe reality* seems plausible as part of the story, but this is too narrow to work as a general account. We need a more general alternative account of what ought-statements do; we should develop an account of the conceptual role of 'ought' that grounds its semantic contribution to the statements in which it figures.

Something like the first answer is fine for uses of 'ought' that Ayer would have regarded as describing the verdicts of some assumed system of norms. I think we should be skeptical of the reality of 'regions of modal space', but talk of

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

what's true in various possible worlds can be a useful way to model the content of statements describing real relations of necessitation, and verdicts following from some assumed system of norms might be viewed as a real relation of necessitation. Perhaps many or even most uses of 'ought' are like this. Certainly, in contexts where one could preface the use of 'ought' with phrases such as, 'According to the values of capitalism …' or 'On a Christian way of looking at this …' In these cases, it makes sense ask: 'But what ought I really to do, think, or feel?' suggesting that there's a difference between describing the verdicts of some system of values or norms and making a genuinely normative statement. Indeed, arguably, many uses of 'ought' are pro tanto, in the sense that one makes an ought statement intending to contribute to someone's reasoning about what to do, think, or feel, but one still allows that the verdict could be undercut, overridden, or erased by other truths about what one ought to do, think, or feel. In these cases, the speaker might be viewed as describing what follows from (or is normatively necessitated) by some system of values or norms, without yet endorsing those verdicts as winning in the end.

If we accept that descriptivist view for some uses of 'ought', why not accept it for all uses of 'ought'? After all, descriptivism offers a simple account of why the modal rule for 'ought' generates the correct truth conditions (insofar as it does): those truth conditions articulate what ought-statements describe (e.g. what is necessitated and how it is necessitated). As already suggested, however, many feel an intuitive difference between describing the verdicts of some system of norms and making a genuinely normative statement. Moreover, if one is skeptical of posits of real relations of 'all things considered' or 'just plain' normative necessitation as part of the fabric of reality, we'll want some view other than descriptivism for genuinely normative statements, on pain of viewing much normative discourse as being in error.

How could we accept the modal rule for 'ought', which predicts truth conditions, and not view all ought-statements as descriptions of reality? To answer this question, it may be helpful to take a step back and consider what we're doing when we develop semantic rules such as the modal rule for 'ought'. These are meant to be parts of a semantic model capable of predicting semantic contents for whole sentences that are compositional in the sense that the content of the whole sentence is a function of the content of the parts and the way these parts are put together. This is viewed by philosophers of language and linguistic semanticists as crucial for explaining the learnability and productivity of language. However, these models don't tell us how to interpret them (e.g. they deploy terms such as 'true' and 'refers' but they don't include a theory of truth or reference). For this reason, I would say that philosophy of language includes the project of explaining what it is in virtue of which statements and their parts have the semantic contents that the best compositional semantics says they have. Because it is not part of (compositional) semantics, I classify this project as 'metasemantic.' It's part of our overall theory of meaningfulness,

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

though it's probably going to appeal to elements of metaphysics, philosophy of mind, psycholinguistics, pragmatics, decision theory, cognitive ethology, evolutionary game theory, etc.

Sometimes I suspect that the uncommitted end up accepting a descriptivist account of all 'ought' statements for lack of an alternative to the representationalist idea that truth conditions tell us how reality has to be in order for the statement to be true. What we need to avoid this easy slide to descriptivism is an alternative metasemantic framework with the space to treat some uses of 'ought' as descriptive and others as nondescriptive.

One way to make this space is to interpret a semantic model's predictions of a statements' truth conditions as telling us *not* how reality has to be for the statement to be true *but rather* what thoughts a speaker has to have in order to be using this statement in accordance with the core communicative rules of the language. For example, if someone says 'Grass is green' most semantic models will predict that this statement is true iff grass is green – a prediction we might then interpret as an articulation of the thought one has to have in order to be using this statement in accordance with the core communicative rules of the language (e.g. sincerity rules and meaning rules).

To get a version of nondescriptivism about genuinely normative ought-statements, one needs then to add that some thoughts are not representations of reality. For, of course, most philosophers want to say that the thought that grass is green is representational of the way reality is. However, representationalism needn't be correct for all thoughts. For example, when it comes to the thought that Sally might be at home, the more natural view is that this though is not directly about reality but rather some sort of qualification of one's epistemic position with respect to Sally's location (which 'in reality' is either in the house or not).

A sophisticated expressivist way of combining these two ideas would be to argue that the thoughts a speaker has to have, in order to be using genuinely normative ought-statements in accordance with the core communicative rules of English, are nonrepresentational because these thoughts have a desire-like direction of fit, ultimately being more like conditional plans or universalized preferences than beliefs about how reality is.[9] Insofar as we are impressed by the intuition that there is an important difference between describing how the world is and saying what someone ought to do, this sophisticated expressivist metasemantics for ought-statements provides a way to endorse the modal rule's predictions about the truth conditions for ought-statements while denying that these statements are always descriptions of reality.[10]

Unlike many other metaethicists, I'm not very moved by the idea that thoughts can be divided into the belief-like and desire-like, with attendant different 'directions of fit' with reality and different roles in practical reasoning. It's not that I doubt the utility of the direction-of-fit metaphor, but I suspect many thoughts fall somewhere in between, having both or neither direction of fit, in

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

some sense and I think the categories *belief* and *desire* are not usefully aligned to this distinction.

Consider, for example, the epistemic 'ought'. For reasons already alluded to, the usual use of 'They left an hour ago, they ought to be there by now,' seems no more descriptive than an all-in normative use of 'You ought to call your mother.' However, with epistemic modal statements, I think it is considerably strained to say that these statements don't describe reality *because they express thoughts with a desire-like direction of fit with reality*. No, the reason they don't describe reality seems to me to have much more to do with the way language can be used in conversation to position ourselves epistemically with respect to the truth of some proposition. Moreover, I suspect that many of the same reasons we might want to be nondescriptivists about ought-statements apply to other modal statements (might-statements, must-statements, etc.). So I long for some more general explanation of why the thoughts expressed by genuinely normative ought-statements (the content of which the sophisticated expressivist takes to be articulated by predictions of the modal rule for 'ought') are not representations of the way reality is; and this explanation needs to be compatible with the idea that there are uses of 'ought' that are descriptive of assumed systems of values or norms.

So that's the main reason why I'm not inclined to go the expressivist route. In seeking an alternative, it can be instructive to ask what a sophisticated form of prescriptivism about 'ought' would look like in this context? The original prescriptivist position was that normative ought-statements don't describe reality but put forward some descriptive content as to be made true (for reasons that generalize). We might bring this in line with a different metasemantic interpretation of the predictions of the modal rule for 'ought' by claiming that the predictions a semantic model makes for truth conditions tell us not (i) what speakers have to think in order to use ought-statements in accordance with the core communicative rules of the language, but rather (ii) what speakers are committed to in virtue of using ought-statements to make assertions in ordinary discursive practice. Then, we could use (iii) to turn the original prescriptivist position into a metasemantic interpretation of the modal rule for 'ought' by suggesting that, while some statements are descriptive in committing speakers ontologically to a way reality is, other statements are prescriptive in committing speakers practically to acting and reacting in accordance with particular prescriptions and whatever they entail.

On this sophisticated form of prescriptivism, the predictions of the modal rule for 'ought' would be interpreted as articulations of the practical commitments carried by ought-statements in ordinary discursive practice. For example, the modal rule for 'ought' predicts that a normative use of 'You ought to call your mother,' is true iff you call your mother in all possible worlds consistent background conditions including your practical situation which are ranked highest by moral ideals. Our sophisticated prescriptivist would not interpret the reference

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

here to possible worlds and moral ideals representationally, rather she would interpret this prediction as telling us something like the following: someone who asserts this ought-statement in ordinary discursive practice is practically committed in a moral way to acting in accordance with the prescription *call your mother!* across some range of circumstances (whatever circumstances are consistent with the background conditions including the practical situation of the 'you' to whom he is speaking). If we want to include a reactive element to this, perhaps we could add that this includes the speaker being committed not only to calling his own mother if in relevantly similar circumstances but also to blaming his audience for not calling her mother in the circumstances in which she currently finds herself.

So far, however, this looks like it has the same problems as the sophisticated expressivist view. It doesn't work well for epistemic uses of 'ought', and it doesn't extend well to other modals such as 'might' and 'must'. Moreover, there are apparently normative uses of 'ought' that are not comfortably assimilated to the idea of prescribing action. What practical commitment is one undertaking, conceived as endorsing the legitimacy of particular prescriptions, when one says 'There ought to be less childhood death and disease'? It's not impossible to answer this question, but most answers strike me as a stretch.

I think these problems can be overcome (but I'm not sure!) by modifying the alternative metasemantic view to be one about commitments to *think and reason in particular ways* rather than commitments to act in accordance with particular prescriptions.[11] This idea generates an account of why the modal rule for 'ought' generates the correct truth conditions (insofar as it does), an account which is comparable in terms of simplicity to that provided by descriptivism: those conditions articulate what ought-statements commit a speaker to, how they have to think and reason in order to satisfy the implicit conceptual commitments affirmed by using 'ought' to make an assertion in normal discursive practice.

To develop this, consider might-statements first. When a speaker makes a might-statement, we can view her as implicitly affirming a commitment to think and reason in particular ways: roughly, to avoid thinking or reasoning in ways that would be inconsistent with the proposition embedded in the might-statement being true. Whether it is acceptable to commit in this way depends on what is ruled out by the body of evidence counting as 'relevant' on the conversational score. Whatever exactly this is, we could conceive of an articulation of the truth conditions of a might statement (i.e. a prediction made in terms of existential quantifications over possible worlds consistent with a body of evidence) as an attempt to spell out which ways of thinking and reasoning the speaker has implicitly affirmed commitment to in making the might-statement.

For example, when a speaker says 'Sally might be at home,' he is committed to not thinking that Sally is at work and to not reasoning in ways that presuppose that Sally is not at home. This commitment is of course defeasible. If he later

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

gets really good evidence that Sally is at work, the speaker is certainly allowed to think that Sally is at work. It's just that the commitment undertaken with the previous might-statement must then be abandoned (even if only implicitly).

So, on this view, the commitment articulated by a statement of the truth conditions for a might-statement is not a practical commitment in the traditional sense of being a commitment to act (moving our bodies towards some end) in in accord with some prescription. But it is also not an ontological commitment to reality being a particular way. Rather it is a commitment to thinking and reasoning in some way.

Could a similar view about 'ought' retain the virtues of the sophisticated expressivist and prescriptivist views while avoiding their problems? To do so, because 'ought' is more flexible than 'might' in the way it can take many flavors (moral, prudential, teleological, epistemic), we'd need to enrich our conception of the ways one can commit oneself to thinking and reasoning. What were previously conceived as practical commitments to act in accordance with some prescription could now be conceived as commitments to reason practically in certain ways, e.g. taking certain considerations to be decisive reasons for acting. This could help with genuinely normative statements such as 'You ought to call your mother.' However, we could also allow for commitments to thinking and reasoning with preferences, e.g. taking certain considerations as reasons for preferring things (even if these never connect to someone's ability to act). This could account for statements such as 'There ought to be less childhood death and disease.' Similarly, we could allow for commitments to thinking and reasoning with credences, e.g. taking certain circumstances to be reasons for assigning a high credence to particular propositions. This might help with statements such as 'They left an hour ago, they ought to be there by now.'

Initially this might look like a hodgepodge, but the general idea is unified: interpret the truth conditions predicted by the modal rule for 'ought' as articulations of the commitment implicitly affirmed by one who uses the statement to make an assertion in ordinary discursive practice, where this commitment is not conceived, in the first instance, as an ontological commitment to the way reality is but rather a commitment to think and reason in some way. Importantly, however, some commitments to think and reason in some way *are* ontological. They're commitments to think and reason as if reality is some particular way. So this metasemantic view is nonrepresentationalist in the sense that *not all* statements are treated as true just in case reality is some particular way, but it still allows that some (even many) statements are true just in case reality is some particular way. So if Ayer is right that some uses of 'ought' are best conceived as descriptions of the verdicts of the moral sense of some community or some assumed system of values, then this account can make sense of these statements as descriptive. But it does so while preserving theoretical space to make sense of other ought-statements as nondescriptive. What this means is that it has promise of being the general account of that in virtue of which sentences have

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

the meanings that our best compositional semantic account will predict them to have, while nonetheless avoiding commitment to global representationalism about truth-apt statements (something we already wanted to avoid because of might-statements).

5. 'Good' as a measurement of value

At the beginning of this paper, I advertised my view about the meaning of normative statements as more inspired by Kant and Frege on modality than emotivists, prescriptivists, and expressivists on morality. I also suggested that this alternative route – via the idea that modal concepts are nonrepresentational – to a nondescriptivist view in metaethics doesn't extend straightforwardly to evaluative statements. This is because evaluative terms are not modals, and it's not natural to treat them as intensional operators. Relatedly, evaluative terms rarely embed whole propositions, such that it might make sense to explain their conceptual role in terms of some kind of formal modification of the propositions they embed rather than in terms of ordinary contribution to the propositional content of the statements in which they figure.

Some philosophers will see this as a big lacuna in my metaethical view. If one's metaethical nondescriptivism applies to normative words such as 'ought' but not to evaluative words such as 'good', then hasn't one failed to defend a nondescriptivist view of the concepts targeted by metaethical inquiry? After all, Moore (1903, ch. 1) was originally focused on judgements about things being *good*, not about what someone *ought* to do, and this is what sparked the familiar debates between nonnaturalists, expressivists, error theorists, and naturalists. Why should we care about a nondescriptivist view about normative statements if that view cannot be extended to a similar view about evaluative statements?

I understand the worry in these questions, but I don't share it. I think it would still represent an interesting form of nondescriptivism in metaethics if normative statements were treated as nondescriptive, but evaluative statements were treated as descriptive. Such a view would respect the *is-ought gap*; it would continue to allow that ought-judgments play a distinctive role in practical deliberation; and it would arguably respect some of the intuitions behind open-question style arguments that originally moved Moore and Ayer. For these reasons, I am sometimes inclined to combine nondescriptivism about 'ought' with some sort of sophisticated relativism about 'good'. However, I also think the above discussion makes space for a weak form of nondescriptivism about 'good', which I shall explore in this section.

Semantically, statements about something's being good seem to be very similar to statements about something's being tall or cold. Words such as 'tall', 'cold', and 'good' are vague gradable adjectives with thresholds. A good start towards articulating the semantic rules governing them would include something like this:

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

SCALAR RULES
$$[[\text{x is tall}]]^c = \text{T iff } \mathit{tall}(\text{x}) > \mathit{threshold}_c(\text{tall})$$
$$[[\text{x is cold}]]^c = \text{T iff } \mathit{cold}(\text{x}) > \mathit{threshold}_c(\text{cold})$$
$$[[\text{x is good}]]^c = \text{T iff } \mathit{good}(\text{x}) > \mathit{threshold}_c(\text{good})$$

Where the function *tall*, *cold*, and *good* take items and returns a measurement of how tall, cold, or good they are in some scale for measuring height, temperature, or value (e.g. feet, degrees Fahrenheit, or weighted preference satisfaction) and the function $\mathit{threshold}_c$ takes measurement functions and returns a threshold degree on the corresponding scale (e.g. how many feet, degrees F, or weighted preferences an item can have and still not be truly said to be tall, cold, or good).[12]

As we know, scalar adjectives are generally context sensitive in the sense that how tall, cold, good, etc. something needs to be in order to be truly said to be, tall, cold, good, etc. will vary from conversational context to conversational context; and even within context vagueness can mean that there are borderline cases. Moreover, as Geach (1956) stressed, most if not all uses of these adjectives need implicit (if not explicit) determination of the category of thing being evaluated before we have any idea what is being said. So, before context can even shift the threshold, it needs to fix the category of items being evaluated as tall, cold, good, etc.[13]

So it is obviously wrong to think that saying something is tall, cold, or good is to provide a context-independent description of reality. However, once context does its work, it is very natural to see ordinary uses of at least some of these adjectives as straightforwardly descriptive. After all, saying that LeBron James is tall in a context where we're talking about all men or that Vancouver is cold in a context where we're talking about North American cities in January would normally be regarded in metaethics as paradigmatic descriptive statements. These are precisely the kinds of statements with which expressivists contrast normative and evaluative statements.

For this reason, representationalism would seem to provide the default interpretation of the scalar rules for 'tall' and 'cold'. The representationalist view, recall, is that the rules above are correct (assuming they are) because they articulate how reality must be in order for the relevant statements to be true. It's important to note, however, that the predicted truth conditions for statements about something's being tall or cold include reference to a measurement scale and threshold. This is provided by context of utterance, and so part of what these statements describe – assuming that they are descriptive – is how something relates in its degree of height or temperature to the threshold on a contextually determined scale.

Turning to the scalar rule for 'good', does the plausibility of the representationalist-cum-descriptivist view about 'tall' and 'cold' mean we should also embrace a form of metaethical descriptivism, according to which the scalar rule for 'good' is interpreted as articulating the *evaluative* way reality has to be in order for

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

the relevant statements to be true? Maybe, but there does seem to me to be a difference between describing the verdicts of some way of measuring value and making a genuinely evaluative statement.[14] One can make the former without endorsing the verdicts, whereas the latter seems to carry some conceptual connection to what there is reason to do, think, or feel. We can all recognize that, as a murder weapon, poison is good, without being forced to accept that there is any reason to use poison to murder or to admire those who do so. And someone who really thinks that saving for their retirement is good would seem to be irrational or unreasonable or deluded if they deny any preference or inclination to save for their retirement.

Above, I argued that a broader nonrepresentationalism could make sense of descriptive *and* nondescriptive uses of 'ought'. The idea was to interpret the truth conditions predicted by standard semantic rules articulating something other than the way reality has to be in order for the relevant statements to be true but also to allow that this 'something other' could sometimes but not always carry ontological commitment. Even though I'm inclined to think that 'good' is used descriptively more than 'ought', this theoretical structure can clearly be extended to cover both descriptive and nondescriptive uses of 'good'.

One way to do this would be to use a sophisticated form of expressivism. That is, we might interpret the scalar rule for 'good' as telling us what thoughts someone has to have, in order to be using good-statements in accordance with the core communicative rules of English. But we go on to argue that, while some thoughts that something is good are thoughts about the way reality is, not all thoughts that something is good need to be like that. More precisely, for those uses of 'good' that look to be descriptions of the verdicts of a contextually supplied system of values (e.g. 'As a murder weapon poison is good'), the sophisticated expressivist says that these thoughts are representational and the statements expressing them are ontologically committing. However, that leaves room for the expressivist to argue that other uses of good (e.g. 'Saving for retirement is good') express thoughts that are not representational. What makes them nonrepresentational? Maybe the expressivist can convince us that these thoughts play a distinctive functional role in our cognitive economies more like sophisticated preferences than representations of reality.[15]

This sophisticated form of expressivism is attractive for the way it can accept the semantic similarity between 'tall', 'cold', and 'good', and for the way it can use the scalar rule for 'good' to treat some good-statements as descriptive and other as nondescriptive. However, it does commit the expressivist to developing a particular view about the nature of the thoughts expressed by genuinely evaluative statements. A nondescriptivism about 'good' more in line with the nondescriptivism about 'ought' outlined above can be less committal about the psychology of evaluative thinking (even if it has to be more commital about the normative structure of thinking and reasoning). If we think that genuinely evaluative uses of 'good' are the ones that carry some conceptual connection

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

to what there is reason to do, think, or feel, then we could try to write that connection into the commitments we think are implicitly affirmed by someone making a good-statement.

Recall, from above, that my preferred alternative to sophisticated expressivism is a metasemantic view that interprets the truth conditions predicted by our best semantics as articulations of how someone making that statement in ordinary discursive practice is committed to thinking and reasoning. This applies to both descriptive and nondescriptive statements. For statements we think are descriptive, we add an account of how someone making the statement is committed to thinking and reasoning, where this *includes* thinking that reality is some particular way or reasoning as if there is something in reality corresponding to the statement. For statements we think are not descriptive, we add an account of how someone making the statement is committed to thinking and reasoning, where this *excludes* thinking and reasoning about the way reality is. For ought-statements, it was natural to focus on practical commitments, commitments to do certain things (though not all ought-statements are practical). For good-statements, it may be more natural to focus on attitudinal commitments, commitments to feel particular ways, in particular to prefer certain things in certain ways.

The resulting picture is one in which both evaluative and descriptive uses of 'good' are assigned the same semantic rule (where the scalar rule for 'good' above is a first approximation). But we don't interpret this rule as telling us what reality has to be like for the relevant statements to be true or what thoughts one who makes the statement has to have in order to conform to the core communicative rules of the language. Rather we interpret this rule as telling us how someone making that statement in ordinary discursive practice is committed to thinking and reasoning. When the good-statement is descriptive, they are committed to thinking and reasoning about reality in some particular way; when the good-statement is nondescriptive, they are committed to thinking and reasoning in other ways about what to do, think, and feel.

6. Conclusion

I began this paper by contrasting the route into a nondescriptivist view about 'ought' offered by Frege and Kant on modality with the more familiar route in metaethics offered by emotivists, prescriptivists, and expressivists. I suggested that the former is attractive for how it hews more closely to the standard intensional semantics for 'ought' and fits with broader metasemantic observations about the linguistic and conceptual role of other modal words such as 'might'. I suggested, however, that a sophisticated expressivist could also provide a metasemantic interpretation of the modal rule for 'ought' capable of funding a kind of nondescriptivism about genuinely normative ought-statements. Moreover, both of these views could make sense of Ayer's observation that

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

some uses of words such as 'ought' are descriptive of the verdicts of particular systems of values that are assumed for the purposes of some conversation, even if not endorsed by the speaker.

So we might conclude that, if one wants a form of nondescriptivism about genuinely normative ought-statements, there are two available routes. The Kant-Frege route may look unattractive for how it doesn't extend straightforwardly to evaluative uses of words such as 'good'. The traditional metaethical project was to explain normative *and* evaluative thought and talk; so providing an account of one of these that doesn't extend to the other is a lacuna in one's overall metaethical view. And since the sophisticated expressivist story about normative thought and talk can be extended to evaluative thought and talk, we might think that's a reason to prefer the Ayer-Hare route to a nondescriptivist view in metaethics.

In the final section of this paper, I sought to deflect this objection. I don't have a settled view about evaluative thought and talk. But I think there are ways of developing a non-expressivist form of nondescriptivism so that it covers 'good' as much as it covers 'ought'. So, in the end of the day, the deciding factor between these two ways of developing a nondescriptivist account of normative and evaluative thought and talk will be broader considerations having to do with what one thinks is foundational to a theory of meaningfulness and concept possession/use. Since I am skeptical of attempts to divide mental states into 'belief-like' and 'desire-like,' and more generally since language is a (mostly) public and observable phenomenon whereas mind is a (mostly) private and inferred phenomenon, I tend to like the non-expressivist route to a nondescriptivist account of normative and evaluative thought and talk. It is one which moves attractively from the inferential/conceptual structure of language to the psychological structure of thought, rather than the other way around.

Notes

1. See Sinclair (2009) and Chrisman (2011) for discussion of the history and some relatively recent developments.
2. Portner (2009, chs. 2–3) contains a great introductory discussion; see also Chrisman (2015) for an introduction.
3. It is the view I defended in Chrisman (2012a) and have developed in more detail in Chrisman (2016a) and (2016b).
4. For different versions of this argument see Dreier (1996), Unwin (1999), Schroeder (2008), and Blome-Tillman (2009).
5. See Humberstone (1971) and Schroeder (2011) for a view in this vein. See Chrisman (2012b) for critical discussion.
6. See Heim and Kratzer (1998, ch. 12) and Von Fintel and Heim (2007).
7. I address some of these in more detail in ch. 5 of Chrisman (2016a), where I argue that there is an agentive use of 'ought' embedding prescriptive content which is not propositional, and I explain a way to capture the relative weakness of 'ought' compared to 'must'.

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

8. This two factored account is the idea pioneered by Kratzer (1981). See Portner (2009, ch. 3) for general introduction and Chrisman (2015) for introduction to the case of deontic modals.
9. Chrisman (2016a, ch. 5, 2016b). See Ridge (2014) for a sophisticated development of this approach within a 'hybrid' expressivist framework that is capable of embracing the modal rule for 'ought' and treating some ought-statements as descriptive and others as nondescriptive.
10. What about embedded uses of 'ought'? On the one hand, this might be a question about the semantic content of complex statements, such as 'If you ought to call your mother, then you ought to charge up your phone.' If so, the answer is relatively easy: the modal rule predicts truth conditions for complete ought-sentences like a general truth-conditional semantic model attempts to do for any declarative sentence. These predictions must then be integrated with the model's treatment of sentential connectives, such as 'if-then' to predict truth conditions for the whole complex sentence in which the simpler sentences figure. To be sure, it is a matter of considerable controversy in compositional semantics how 'if-then' works, but as long as it takes truth-conditional complements, the modal rule for 'ought' will be able to integrate with a rule for 'if-then' to produce truth conditions for the complex sentence. On the other hand, however, the question about embedded uses of 'ought' could be a metasemantic question about what grounds our semantic model's predictions of the truth conditions for complex 'if-then' statements. This is a very difficult question, but one that is perhaps even more difficult for representationalists than nonrepresentationalists. It is sometimes claimed in compositional semantics that ordinary language conditionals are covert modals, in which case some of the same descriptivist and nondescriptivist ideas explored here about 'ought' might be explored in developing a metasemantic interpretation of 'if-then'. But there are also other possible metasemantic accounts of 'if-then'.
11. In Chrisman (2017), I explain how this might be incorporated into a conceptual role account of meaning for the sorts of expressions of interest to metaethicists.
12. Compare Kennedy (2007).
13. See Thomson (2008) for a worked out version of this view addressing many of the shortfalls of Geach's own suggestion but continuing in a similar spirit.
14. Also 'good' unlike 'tall' and 'cold' seems to be multidimensional. I will largely ignore this here, but it provides another path to a nondescriptivim about 'good': perhaps some contexts do not determine how the various dimensions of value are to be weighed in determining whether something is good. In such cases, the function of good-statements might not be to describe something's value but to set a standard for weighing competing values. See Plunkett and Sundell (2013) for further discussion of this metalinguistic use of vague and context-sensitive adjective and some of the implications it has for metaethics.
15. Compare Silk (2015) and Köhler (forthcoming).

Notes on contributor

Matthew Chrisman is a professor in Philosophy at the University of Edinburgh. His research is focused on ethical theory, political philosophy, the philosophy of language, and epistemology. He has published widely in these areas, including books with Oxford University Press and Routledge and articles in *Noûs, The Journal of Philosophy, The Australasian Journal of Philosophy, Oxford Studies in Metaethics, Philosophy and Phenomenological*

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

Research, Philosophers' Imprint and Philosophical Studies. Matthew earned his PhD and MA in Philosophy from the University of North Carolina at Chapel Hill and his BA from Rice University.

References

Ayer, A. J. 1936. *Language, Truth and Logic*. London: V. Gollancz.

Blome-Tillman, Michael. 2009. "Non-Cognitivism and the Grammar of Morality." *Proceedings of the Aristotelian Society* 103 (1): 279–309.

Chrisman, Matthew. 2011. "Ethical Expressivism." In *Continuum Companion to Ethics*, edited by Christian Miller, 29–54. New York: Continuum.

Chrisman, Matthew. 2012. "'Ought' and Control." *Australasian Journal of Philosophy* 90 (3): 433–451.

Chrisman, Matthew. 2012. "On the Meaning of 'Ought'." In *Oxford Studies in Metaethics*, vol. 7, edited by Russ Shafer-Landau, 304–332. Oxford: Oxford University Press.

Chrisman, Matthew. 2015. Deontic Modals. *Routledge Encyclopedia of Philosophy*.

Chrisman, Matthew. 2016a. "Metanormative Theory and the Meaning of Deontic Modals." In *Deontic Modality*, edited by Nate Charlow and Matthew Chrisman, 395–430. New York: Oxford University Press.

Chrisman, Matthew. 2016b. *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. New York: Oxford University Press.

Chrisman, Matthew. 2017. "Conceptual Role Accounts of Meaning in Metaethics." In *The Routledge Handbook of Metaethics*, edited by Tristram McPherson and David Plunkett, 260–274. New York: Routledge.

Dreier, James. 1996. "Expressivist Embeddings and Minimalist Truth." *Philosophical Studies* 83 (1): 29–51.

Frege, Gottlieb. (1879) 1967. "Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens (Halle a. S.: Louis Nebert)." [Concept Script, a formal language of pure thought modelled upon that of arithmetic.] In *From Frege to Gödel: A Source Book in Mathematical Logic*, translated by S. Bauer-Mengelberg, edited by J. van Heijenoort, 1879–1931. Cambridge, MA: Harvard University Press.

Geach, P. T. 1956. "Good and Evil." *Analysis* 17: 33–42.

Hare, R. M. 1952. *The Language of Morals*. Oxford: Clarendon Press.

Heim, Irene, and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Blackwell.

Humberstone, I. L. 1971. "Two Sorts of 'Ought's." *Analysis* 32 (1): 8–11.

Kant, Immanuel. (1787) 1998. *Critique of Pure Reason* (translated and edited by Paul Guyer & Allen W. Wood). Cambridge: Cambridge University Press.

Kennedy, Christopher. 2007. "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives." *Linguistics and Philosophy* 30: 1–45.

Köhler, Sebastian. forthcoming. Expressivism, Meaning, and All That. *Canadian Journal of Philosophy* ,

Kratzer, Angelika. 1981. "The Notional Category of Modality." In *Words, Worlds, and Contexts*, edited by Hans Eikmeyer and Hannes Reiser, 38–74. Berlin: de Gruyter.

Moore, G. E. 1903. *Principia Ethica*. Mineola, NY: Dover Publications.

Plunkett, David, and Tim Sundell. 2013. "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosophers' Imprint* 13 (23): 1–37.

Portner, Paul. 2009. *Modality*. Oxford: Oxford University Press.

Ridge, Michael. 2014. *Impassioned Belief*. Oxford: Oxford University Press.

Schroeder, Mark. 2008. "How Expressivists Can and Should Solve Their Problem with Negation." *Noûs* 42 (4): 573–599.

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press

Schroeder, M. 2011. "Ought, Agents, and Actions." *Philosophical Review* 120 (1): 1–41.
Silk, Alex. 2015. "How to Be an Ethical Expressivist." *Philosophy and Phenomenological Research* 91 (1): 47–81.
Sinclair, Neil. 2009. "Recent Work in Expressivism." *Analysis* 69 (1): 136–147.
Thomson, Judith Jarvis. 2008. *Normativity*. Chicago, IL: Open Court.
Unwin, Nicholas. 1999. "Quasi-Realism, Negation And the Frege-Geach Problem." *The Philosophical Quarterly* 49: 337–352.
Von Fintel, Kai, and Irene Heim. 2007. *Intensional Semantics*. Unpublished Lecture Notes, Cambridge, MA.

https://doi.org/10.1080/00455091.2018.1432400 Published online by Cambridge University Press