

Practical Commitment in Normative Discourse

Pekka Väyrynen
University of Leeds

1. Introduction

Evaluative and deontic judgments play a practical role in our thought. Judging that something is good, or right, or what we ought to do, ranks things in ways that guide us towards the higher-ranked things. When I judge that I ought to eat more iron, normally I reveal (or make) myself to be motivated or at least committed to eating more iron, feeling bad if I don't, and so on. When we judge that we have most reason to slow down climate change, this normally “stops the buck” in deliberation and decision about what to do. It is controversial whether such a role in practical reasoning is essential to the evaluative or deontic concepts deployed in these judgments. It is controversial to what extent “thick” concepts like TACTFUL, KIND, and CHASTE play such a practical role in our thought. And it is controversial just what the practical role of a given normative concept is. But it is widely agreed that, at minimum, descriptively “thin” normative concepts play some such practical role for rational agents with a normal grasp of those concepts.¹ Thin normative concepts are, roughly, normative concepts which encode little descriptive information. They are characteristically expressed by terms like *good*, *right*, and *ought* in English and their equivalents in other languages. Call these ‘thin normative terms’.²

This is a preprint of a paper forthcoming in *Journal of Ethics and Social Philosophy* (<http://jesp.org>). Please cite the official published version when available.

- 1 I'll use small capitals to denote concepts, single quotes to mention sentences, and double quotes for a variety of other uses of quote marks, such as mixture of use and mention, direct quotation, scare quotes, and introduction of theoretical notions. Italicized words denote lexical items. The use of italics for emphasis instead will be clear from the context. I use ‘normative’ to cover both the evaluative and the deontic. What I call a practical role is called “practical role” by Ridge (2014: 21), “normative role” by Eklund (2017: 9), and “guiding or regulative role” by Wedgwood (2017: 15). The term ‘normative role’ may be misleading insofar as it may refer merely to how a normative concept is associated with other normative concepts – not my topic here. Finally, to keep things manageable, I'll bracket epistemic judgments concerning what one ought to believe.
- 2 These terms have other uses as well. I assume we have a reasonably determinate grasp of which uses of these terms express thin normative concepts. What I say won't hang on having a precise delineation of that class.

This paper concerns how the practical role of normative thought is reflected in *language*. Normative language manifests a parallel phenomenon. I reveal (or make) myself to be motivated or committed to eating more iron no less when I assert ‘I ought to eat more iron’ than when I judge that I ought to eat more iron. It is natural to wonder whether the phenomenon is semantic or pragmatic, especially since it may not be cleanly delineable in natural language. If the phenomenon were specific to bare first-person assertions like ‘I ought to eat more iron’, there would be little initial pressure to think it semantic. Even if normative assertions in the first person are important to understanding normative thought or its relationship to action, it doesn’t follow that they are linguistically special. And yet the standard assumption in metaethics is that at least thin normative terms are associated with their practical role somehow as a matter of meaning.³ This assumption is usually implicit and rarely articulated precisely. But it faces a serious problem in any form. As we’ll see, terms like *ought* and *good* often figure in assertions where their customary practical role is absent. Such cases pose a challenge: either offer some plausible explanation of why the relevant practical upshots don’t show up in these cases despite featuring in our overall semantic theory for these expressions, or else don’t build them into such a theory. I’ll argue that existing accounts of normative language have no adequate explanation of how the association between thin normative terms and practical role can be broadly semantic and yet tolerate the exceptions I’ll highlight. In closing I’ll consider the prospects of a pragmatic account of the phenomenon. This is an option to take seriously because generalizations that don’t rise to the level of full generality are normally more apt to be explained in pragmatics.

My enterprise is hermeneutic: it concerns how the association between thin normative terms and their practical roles is in fact reflected in natural language. This gives the issues at stake and my arguments significant broader interest and importance. If the association between thin normative terms and practical role isn’t broadly semantic, this puts significant pressure on standard forms of metaethical expressivism and inferentialism. For expressivists typically account for the meaning of normative language in terms of practical states of mind, such as plans, desires, or sentiments. And inferentialists typically account for the meaning of normative language in terms of practical commitments to not only think but also act and feel in

3 Exceptions include Strandberg (2012) and Finlay (2014); see §6.1. My argument will be more general than theirs.

certain ways. Existing accounts of neither sort have an adequate solution to the challenge I'll raise. If the association between thin normative terms and practical role isn't broadly semantic, this will also suggest certain general lessons about how normative thought and language relate. For instance, our normative discourse can be in perfectly good working order without containing predicates that are associated with certain practical roles in any broadly semantic way. Even if there are distinctively practical OUGHT-concepts that are conceptually or constitutively linked to certain practical upshots, we have reason to doubt that this link will be encoded in the correct overall semantic theory for natural languages. Such concepts may be privileged in normative theory, but not in our overall semantic theory for normative language. We might then also wonder how crucial such concepts really are for normative thought. Philosophers of normativity will therefore need to think more carefully about the relation between normative thought and normative talk.

2. Semantic Views of the Practical Role of Normative Discourse

One of the few widely shared convictions among moral philosophers is that normative judgments play a distinctive role in our practical thinking. We may disagree about just what the practical role of a given normative concept is and be unable to state it in uncontroversial terms. But we have some grasp of the notion, since we routinely compare the practical roles of different normative concepts (Eklund 2017: 45-46). For instance, moral evaluation is different from aesthetic evaluation, and the practical role of CHASTE is different from that of RIGHT. Some normative concepts may even be individuated, at least in part, by their practical role. I won't deny that people can have such concepts.⁴ But don't words express concepts? So if there were distinctively practical concepts, wouldn't it then be part of the meaning of the words that express these concepts to have this connection to practical role? Sincere moral claims by rational agents, for instance, are often thought not to tolerate the absence of corresponding practical upshots. This line of thought might also help distinguish normative terms (and not just normative concepts) from non-normative terms. What could the difference be, if not that something about the meaning of normative terms connects them to a practical role?

4 Whether analyzing the nature of normative concepts in terms of their practical point would yield any substantive results is a further question. For a cautionary discussion, see Sharadin and van Someren Greve (2021).

This is all too quick. If words express concepts, it doesn't follow that those concepts directly determine the meaning of those words. In the normative case, the following is possible, for instance: (i) if a concept plays a role *R* for a thinker, then it picks out moral rightness but (ii) the word *right* isn't constrained by the rules of the language to stand *only* for such a concept. In other contexts, it might be used to express concepts that lack a practical role, such as what is right according to my parish priest. So even if some normative concepts are essentially associated with certain practical roles, it doesn't follow that this connection is reflected in the meanings of the words used to express those concepts. In general, linguistic meaning and mental content might not line up in any neat way (cf. Yalcin 2014; Glanzberg 2018). It remains similarly open whether what makes normative discourse different from non-normative discourse is how their meanings relate to some distinctive kind of practical import.

What might it be for thin normative terms to be associated with their practical roles as a matter of meaning? Our overall theory of meaningful linguistic communication is normally taken to have three main parts (whose boundaries may not be sharp): descriptive semantics, metasemantics, and pragmatics. When I ask whether the practical role of thin normative terms is a "broadly semantic" phenomenon that holds "as a matter of meaning" and is to figure in our "overall semantic theory", the question is whether the phenomenon is a descriptive-semantic or metasemantic rather than pragmatic phenomenon. A pragmatic account would take an account of the meaning of a normative sentence and say that the relevant practical upshots are some further kind of information which may be conveyed by its utterances, given its meaning and the context: conversational implicature, pragmatic presupposition, or the like.

On descriptive-semantic accounts, the association between thin normative terms and practical role would be a feature of assignments of semantic value to thin normative terms. The term 'semantic value' is a neutral term for the things which, according to a semantic theory, "provide the interpretations of simple expressions and are the arguments and values of the functions defined by the compositional rules that interpret the complex expressions" (Stalnaker 1997: 535). I'll take descriptive semantics to include an account of the lexical properties of simple expressions and other conventional aspects of their meaning in a language, since these will constrain

assignments of semantic value.⁵ (For context-sensitive expressions, context provides further constraints.) So, for instance, if the practical role of *ought* were a lexical property of the word, it would be a descriptive-semantic phenomenon. Assignments of semantic value to expressions by a theory are supposed to account for features like the entailments licensed by the expressions and the truth-conditions of the sentences they compose.⁶

On metasemantic accounts, the association between thin normative terms and practical role would instead be a feature of that in virtue of which thin normative terms have the semantic values they do. Metasemantics (sometimes called “foundational semantics”) is about “what the facts are that give expressions their semantic values, or more generally, about what makes it the case that the language spoken by a particular individual or community has a particular descriptive semantics” (Stalnaker 1997: 535). It involves explaining both why expressions have the context-invariant descriptive-semantic features they do and what it is about the situation, behavior, or the mental states of the speaker that makes it the case that a context-sensitive expression has the contextually determined semantic value it does. For instance, if the semantic value of *ought* were determined by its inferential role and its inferential role were essentially practical, then the practical role of *ought* would be a metasemantic phenomenon (see §5.4).

Canonical statements of semantic or metasemantic accounts of the practical role of normative discourse are hard to find. But there are suggestions in the ballpark. Hare thinks that the “primary meaning” of *ought* is “prescriptive” – one that, by Hare’s stipulation, entails an imperative (Hare 1952: 159, 164). Similarly, the primary meaning of *good* is an evaluative meaning that entails a recommendation. This, he thinks, is what anchors meaningful connections between the applications of these terms to very different sorts of things by different people. (Consider the example of the missionary and the cannibals in Hare 1952: 148-9.)

5 Conventional aspects of meaning are (very roughly) relatively arbitrary regularities of a population using a certain linguistic form (e.g. a word or a sentence) to mean a certain thing. (Conventions are arbitrary in that no general account than a historical one is available for them.) For a discussion of various complications, including that conventionality may have multiple dimensions and come in degrees, see Simons and Zollman (2019).

6 The dominant approach to semantics in this sense is truth-conditional semantics: the semantic value of a sentence just is its truth-conditions, and the semantic value of an expression is the contribution it makes to the truth conditions of the sentences in which it features. But of course there are alternatives. Expressions may also have various non-compositional conventional features. Here I’ll have to bracket views on which moral claims play a practical role via conventional implicature (Copp 2001) or semantic presupposition.

Dreier’s “speaker relativism” analyzes ‘*x* is good’ as meaning ‘*x* is approved of by *M*’, where *M* is a contextually specified moral system. He takes it to be built “into the very semantics of moral terms” that *M* is picked out in part by certain kinds of motivational and affective factors (Dreier 1990: 6, 23-4). Blackburn writes that “since moralizing and valuing are distinctive activities, the words we use to communicate our morals and our values will have their distinctive meanings” and that capturing the way that ethics is “essentially practical” requires “showing how [ethical] judgment has a content or truth-condition that is itself magnetic” (Blackburn 1998: 87, 115). Wedgwood proposes that the meaning of *good* and *ought* each is given by the basic rules of rationality governing its use, where these rules concern commitments to certain preferences or planning states (Wedgwood 2001; 2007: ch. 4). According to Gibbard, the meaning of *ought* is explained by its tie to planning states; “plan-laden concepts” have “much to do” with what ordinary normative terms like *ought* express (Gibbard 2003: 139). These proposals all affirm a connection between what thin normative terms mean and some kind of practical role. As we’ll see, they share a significant assumption. But they resist summary into a single hypothesis.

One respect in which these proposals differ concerns what kind of thing the practical role of normative language is. Some focus on practical *dispositions*: at least all else equal, judging that one ought to ϕ in circumstances *C* (etc.) tends to motivate one to ϕ in *C*, or to close deliberation about whether to ϕ , or the like. Perhaps one will also tend to feel guilt if one doesn’t ϕ or resent or blame others if they don’t ϕ in relevantly similar circumstances. Others focus on practical *commitments*: judging that one ought to ϕ in *C* (etc.) commits one to thinking, feeling, or acting in certain ways with respect to ϕ , such as perhaps forming the intention to ϕ in *C*, or making ϕ -ing in *C* part of one’s ideal plan about what to do. Perhaps it also commits one to disapproving, resenting, or blaming those who don’t ϕ in *C*, and so on.⁷ Of course, it might be that the practical role of some normative terms involves motivations while

7 Practical commitments differ from various species of theoretical commitment. They aren’t merely *intranormative*. Judging that I (morally, legally, etc.) ought to ϕ in *C* might commit me to thinking there is (moral, legal, etc.) reason for me to ϕ in *C*. Whether this commits me to acting, feeling, or thinking on the basis of those reasons is a further question. Nor can these commitments be merely *alethic* or *doxastic*. If I judge that a measure of voting intention is unreliable, I commit myself to the truth of that proposition, but not any intention or sentiment. The same goes for the commitment to believe *q* generated by judging that *p* and believing that $p \rightarrow q$. Practical commitments are more like the following: if I promise that I’ll be a better parent, I commit to being a better parent in the future, and the same goes if I endorse the thought that I shall be a better parent.

that of some others involves commitments, and the practical role of some terms might involve both. In what follows I'll mainly talk of practical commitments, since this is a logically weaker notion. Affirming a practical commitment requires less of an agent than having a matching disposition to comply.⁸

Another respect in which the proposals noted above differ concerns what kind of connection is supposed to obtain between thin normative terms and their practical role. Hare and Blackburn seem to make claims in descriptive semantics. Dreier, Gibbard, and Wedgwood seem to have in mind metasemantic views. But they all seem to imply that thin normative terms' association with their practical roles is some indispensable feature of their meaning which any competent user of these expressions must master.⁹ This is a useful starting point; it'll help to articulate a problem which any broadly semantic view must address somehow or other.

3. The Problem of Diverse Uses

I'll now argue that thin normative terms are often used sincerely and felicitously to state requirements or recommendations in ways that don't commit the speaker to thinking, feeling, or acting in accordance with those verdicts.¹⁰ If the practical role of these terms were a part of their conventional profile in a language, it shouldn't be subject to such exceptions but instead should be present in all literal uses in normal contexts.¹¹ This raises what I'll call the Problem of Diverse Uses: how do you reconcile the diversity of uses to which thin normative terms may be put with the claim that their association with their normative roles is broadly semantic? The problem prompts a challenge: either offer some plausible explanation of cases where the relevant practical upshots are absent which reconciles these claims, or else don't

8 Either kind of practical role is richer if some form of internalism about it is true, such as motivational internalism about normative judgment. Each is distinct from Klimczyk's (2020) idea of "authored practicality".

9 In due course I'll ask whether all broadly semantic views carry this implication in full. Some recent metasemantics for normative discourse might not. These include Ridge (2014), Chrisman (2015, 2018), Köhler (2018), Bedke (2019), and Tiefensee (2016, 2021). I'll discuss Ridge, Chrisman, and Tiefensee in §§5.3-5.4.

10 As the examples below make clear, I'm not talking about sentences like 'Martha says I ought to tell the truth', 'Either I ought to tell the truth or it is not the case that I ought to tell the truth', or the embedding of thin normative terms under negation or possibility modals like *might*, or in the antecedent of a conditional.

11 This holds irrespective of whether practical upshots are features of propositions speakers communicate when making claims using normative terms, or of acts of making such claims, such as their propriety conditions. An example of the latter kind of view would be the ethical neo-expressivism of Bar-On and Chrisman (2009).

build such upshots into our overall semantic theory for thin normative terms. I'll be driving us to the latter option.

One sort of evidence for the relevant kind of diversity of uses is that any suggestion of practical commitment is absent in cases like (1)-(2):

- (1) One ought to prioritize profit over fairness. But is that really the thing to do?
- (2) a. [Client:] What is my legal obligation, and what do you expect me to do?
b. [Lawyer:] You have to report your liability, but I don't know if you will; you may prefer to push the limits of the law and just conceal it. (Mandelkern 2021: 43.)

Claims like (1) can make perfectly good sense in conversations about capitalist worldviews. The corresponding interpretation of *ought* need only be implicit in the context for the overt utterance to sound coherent. Variants of (1)-(2) arise by varying the modal verbs: *ought* for *have to*, and so on.¹² The point of these examples is that practical commitments don't result conventionally from using *wrong* or *ought* to state requirements or recommendations. Conventional features of utterances aren't as easily defeasible as the absence of practical upshots in (1)-(2) would require.

One might claim that (1)-(2) are insufficient to show that the association between thin terms and practical role isn't semantic. Perhaps not all sincere assertions of thin normative claims in normal contexts must carry practical upshots for such association to be a broadly semantic matter. For instance, the phenomenon seems clearest in bare first-person uses, but (1)-(2) aren't such cases. However, there is no general reason why a difference in personal pronoun should make this particular kind of semantic difference. Moreover, practical upshots are absent also in some bare first-person assertions. To many people's ears, sincere assertions of (3)-(4) are indicative of a substantive normative mistake rather than linguistic incompetence:

- (3) I would be wrong to kill. But I'm okay with killing and don't feel bad about it.

¹² See e.g. Silk (2016: 40): 'Ernie ought to be home by 10. Aren't his parents stupid? I'd stay out if I were him.'

- (4) I ought to finish grading. I have absolutely no intention to do so, though.¹³

If a sincere assertion of ‘I ought to finish grading’ always committed you to planning to do the thing, then asserting it while also rejecting such plans, as in (4), should induce a strong feeling of incoherence. But it needn’t. Even if *wrong* and *ought* in many contexts express concepts whose application is practically committal in this way, instances of (3)-(4) needn’t express such concepts. The first half of (4) would often express something more like a role obligation.

So (3)-(4) are evidence that a practical commitment doesn’t result conventionally even from first-person assertive utterances of lexical items like *wrong* and *ought*. Similar evidence comes from psychologically realistic characters who take delight in evil for evil’s sake, are committed to sheer malignity under that description, and so on. When Iago sincerely identifies his conduct as demonic, he is specifying a feature of his conduct which he is committed to pursuing and strongly inclines him to engage in the conduct (cf. Kramer 2019). Further first-person cases where practical commitment is absent include (5)-(6):

- (5) I should do the shopping today (as far as I know).¹⁴
- (6) a. [CEO:] What is our legal obligation, and what we should we do?
b. [Head of Legal:] We have to report our liability, but could decide to push the limits of the law and just conceal it.

In (5) the speaker states a normative requirement that follows (as far as she knows) from her household’s shopping arrangements. Her commitments to responding accordingly are simply a further issue. (6) reminds us that not all first-person uses are singular. A special convention for bare singular first-person present-tense assertions is all the less likely, and wouldn’t in any case be a conventional aspect of the meaning of thin normative terms themselves.

13 Woods (2014) discusses these kinds of examples in the context of expressivism. Ruiz and Stojanovic (2019: 159-163) suggest that such examples need a contrasting connective like *but* and this is evidence that if a rational and sincere speaker says something positive about an object, she normally expresses a certain practical stance. I can agree, but would deny (in part for reasons I’ll mention in §6.1 below) that this tells us as much about the meaning of normative terms as they take it to do.

14 I adapt (5) from a second-person example in Kaufmann (2012: 58).

(1)-(6) illustrate that thin normative terms are often used non-committally, to state normative verdicts in ways that don't carry corresponding practical commitments.¹⁵ My sample theorists acknowledge some such cases.¹⁶ Hare says that *ought* and *good* can have "secondary" descriptive meanings. Dreier, Blackburn, and Gibbard allow that even sincere first-person *ought*-statements don't always imply that the speaker is in a corresponding conative state. But none is very clear about just how thin normative terms' association with their practical role can allow exceptions and yet be a broadly semantic phenomenon.¹⁷ (I'll extract some suggestions shortly.) Examples like (1)-(6) imply that such an association isn't the sort of conventional matter that these views take it to be.¹⁸ Similarly, insofar as assertive uses of thin normative terms carry directive or commissive illocutionary force, such force won't derive from their lexical meaning.

-
- 15 Some readers might be wondering about uses of normative terms to express imperatives. On the relationship between strong necessity modals like *must* and *have to* and imperatives, see e.g. Ninan (2005), Kaufmann (2012: §2.3.1), and Mandelkern (2021). Some of my examples suggest that these modals don't always (even in their deontic uses) express orders or commands. Although my focus isn't on orders or commands, it is worth noting that the performative element of an order (its inducement for the addressee to change the world in some way) might be a kind of pragmatic meaning. Kaufmann argues that imperatives ('Close the door!') are modalized declaratives ('You must/should close the door') which come with presuppositional ingredients that constrain them to usage in contexts in which the modalized declarative would achieve a non-assertoric, performative effect (Kaufmann 2012: 63, 132). The pressure to act which imperatives often induce on their addressees is located in certain pragmatic presuppositions that govern their felicitous use, not in an explanation of their meaning (Kaufmann 2012: 144-63). However, the debates about imperatives is ongoing.
- 16 In Wedgwood, this is implicit. His account is officially limited to a certain "deliberative" concept that *ought* can be used to express (2007: 97). He doesn't say how to generalize his analysis so that it would account for the meaning of *ought*, given the many other deontic concepts which the word can be used to express.
- 17 Hare, Blackburn, and Gibbard describe uses of *ought* for which they take motivational internalism to be a conceptual truth. But *ought* can be used in deontic contexts to express normative concepts for which motivational internalism doesn't hold. Dreier analyzes 'x is good' as meaning 'x is approved of by M', where M is a contextually specified moral system. According to him, it is built into the meaning of *good* that M is chosen from the context by balancing the content and subject matter of a set of rules and those rules' motivational and affective connections with the speaker (Dreier 1990: 24-5). But the meanings of these terms don't require us to interpret (1)-(6) by considering what rules have such connections. (Also, *good* has other than moral uses.)
- 18 My arguments allow that to understand the lexical meaning of a word, we may have to look at a group of semantically related words, such as *may*, *must*, and *have to* in the case of *ought*. How concepts relate in general to lexical meaning is a more complex issue to which I cannot do justice; Glanzberg (2018) is a rich discussion.

4. Diverse Uses and Descriptive Semantics

Any account of the diverse uses of thin normative terms which we have observed will take one of two forms. Practically committal and non-committal uses of thin normative terms either have a uniform descriptive semantics, or they don't. These are the options whether or not their practical role is a broadly semantic phenomenon. It'll be instructive to begin by considering why some simple broadly semantic views of each type fail to solve the Problem of Diverse Uses.

One proposal that posits semantic uniformity says that deontic uses of terms like *ought* affirm the practical commitments of some agent or other. The thought might be that committal uses convey the speaker's own commitments, and non-committal uses play the relevant practical role indirectly by committing the speaker to thinking or feeling in a certain way in such-and-such a contingency (cf. Gibbard 2003; Ridge 2014: 19). But this won't do. To interpret the deontic content of 'We ought to report our liability, but could decide to push the limits of the law and just conceal it' in (6), we needn't suppose that the Head of Legal is committed to reporting the liability in the case where the firm decides not to push the limits of the law. Or consider a value system which a science fiction author just invented. There is no pressure to interpret the *ought*-claims that describe that value system's verdicts as expressing the practical commitments of the members of this culture. In such cases, the relevant agent could only be specified as "someone who endorses the norms at issue". In general, we understand thin normative claims so long as we can identify the relevant standards: a capitalist value system in (1), the local legal code in (2) and (6), and so on. Stating the recommendations of a standard doesn't require assumptions regarding who if anyone is committed to those standards. And while such recommendations can always be represented by a preference ordering, this doesn't require preferences to encode motivations or commitments. They need only function as abstract ranking devices.

A correspondingly simple proposal without semantic uniformity is that although non-committal uses of thin normative terms don't directly convey practical commitments, these uses are ultimately to be explained by reference to committal uses. A proposal that is familiar from other contexts is to analyze non-committal uses as parasitic on committal uses. They might be claimed to report or otherwise allude to other people's committal normative judgments (cf. Hare 1952: 124-26, 164-65; Blackburn 1998: 59-68, 110). However, this solution undergeneralizes.

Normative language can be used make non-committed statements about what follows from some normative system which have never occurred to anyone before (Raz 1980: 235).

A different strategy that implies lack of semantic uniformity is to argue that cases where the relevant practical upshots are present and absent are distinguished by linguistic convention after all. One view is that thin normative terms have multiple distinct but related meanings. It is controversial whether modal expressions like *ought* are polysemous across different flavors of modality, such as epistemic, deontic, and dynamic.¹⁹ But examples (1)-(6) all concern *ought* in deontic contexts. It is much less common to think that the flexibility of modal expressions within a modal flavor is due to polysemy, rather than context-sensitivity (on which more in §5).²⁰ Features of a mere subset of uses of deontic expressions aren't evidence about their meaning.

A more sophisticated version of this strategy is to claim that deontic claims which don't play a practical role aren't genuinely normative claims, and so deontic terms aren't used normatively in (1)-(6). What 'normative' means in the talk of normative terms and claims is an infamously fraught issue (Finlay 2019). The notion of the "genuinely" normative also requires explication. But interpreting 'normative' to include only the genuinely normative is unduly narrow for the present purposes in any case. Suppose the sense in which a sentence like 'One ought to do that' is normative only applies to uses which essentially involve practical commitments. This class of "normative" sentences would exclude many uses of *ought* which also

19 The dominant tradition in the semantics of modals, developing the framework in (Kratzer 1981; 1991), aims to give a unified semantics of deontic, epistemic, and other flavors of modality. Viebahn and Vetter (2016) argue that each modal expression has distinct but systematically related meanings, for different modal flavors, while their flexibility within each modal flavor is due to context-sensitivity. This is compatible with my arguments. Even if the lexical entry for *ought* includes a variable whose different values give *ought* an epistemic, deontic, or goal-oriented flavor, the practical role of deontic *ought* doesn't have a comparable status. See §5.2.

20 See also Kaufmann (2012: 60-62), who uses polysemy avoidance to make a case for a uniform semantics for descriptive and performative modal verbs. There is a debate in metaethics about whether *ought* is polysemous between a "practical"/"deliberative" sense and an "evaluative" sense (Schroeder 2011; Finlay and Snedegar 2014; Chrisman 2015: 124-33). But the distinctions cited in that debate differ from those on the table here. Perhaps the deliberative *ought* is always indexed to a particular agent and time (Wedgwood 2007: 90). But *ought* can be so indexed when it is used to state what follows from some assumed standard for some other agent's situation at a given time. Or perhaps the deliberative *ought* "matters directly for advice" and "is the right kind of thing to close deliberation" (Schroeder 2011: 17). But an *ought*-claim describing what follows from a set of standards can constitute advice which is suitable for closing deliberation without the relevant practical commitments being in place; recall (2) and (6). (Cf. Bronfman and Dowell 2018: 107-8.)

are distinctive from typical uses of paradigmatically descriptive, non-normative words like *tree* or *tall*.²¹ The deontic claims in (1)-(6) are normative in the broader sense that they concern standards falling short of which opens one up to certain kinds of criticism – legal standards in the case of (2) and (6), and so on. Even if both ‘I ought to finish grading’ (as asserted in the context of (4)) and ‘Matt is tall’ are descriptive in *one* sense, the implications of falling short of the standards invoked by (4) are different from those of falling short of the contextually supplied standards of tallness. The lack of practical upshot may mean that (1)-(6) express requirements whose normativity is merely “formal” rather than “substantive” or “authoritative”. But these cases show that one function of these words in English is to express formally normative claims. So the sense of ‘normative’ that is relevant to our semantic theory for expressions like *good* and *ought* had better include also such formal normativity.

We could instead seek progress by characterizing the different uses to which normative terms may be put. We distinguish “engaged” or “committal” and “detached” or “non-committal” uses of normative language (Raz 1980: 234-35 and above). We might gloss this further by saying that normative terms have non-committal uses that state requirements or recommendations that follow from some assumed body of standards, and committal uses that somehow (semantically or pragmatically) also convey corresponding practical commitments. One proposed test is that a use is at least normally non-committal if the relevant utterance is reasonably interpreted as if it contained an implicit *according to*-type phrase.²² This is a fallible test. When the doorman says to me ‘According to the rules of our club, patrons may not wear sneakers’, this isn’t a merely descriptive use. The doorman is committed to not letting me in. But we recognize the distinction, however particular cases fall with respect to it. (In this example, the practical upshot of the doorman’s utterance is clearly a kind of pragmatic meaning.)

It is one thing to grant this distinction of uses of normative terms, quite another to claim that it is somehow marked by linguistic convention or otherwise

21 The more narrowly we understand ‘normative’, the less clear it is that there strictly speaking are normative words or sentences in English. There may just be words and sentences which can be used normatively (in the various senses of ‘normative’) or non-normatively without changing meaning. Alwood (2016: 15) makes this point about the notion of ‘normative’ in Ridge (2014). I’ll discuss Ridge’s view in §5.3.

22 Cf. Silk (2016: 130-32) and Chrisman (2018: 412). Ruiz and Stojanovic (2019) propose various other criteria for distinguishing “expressive” and “factual” uses of terms. Not all of these, however, seem likely to carry over.

drawn in our semantic theory for normative terms. I see little reason to suppose that the distinction between committal and non-committal uses of terms like *ought* and *good* has such a status. The distinction doesn't involve a difference in linguistic form, whereas linguistic conventions pair forms with meanings. So even if there are complex regularities regarding when an assertion of 'I ought to finish grading' is committal and when it is non-committal, it isn't at all clear how they could be candidates for linguistic convention.²³ Whether an assertive use of 'I ought to finish grading' is committal or non-committal is a function of features of the context of utterance which don't normally ground linguistic convention. The distinction looks more apt for speakers to track by means of their world knowledge and pragmatic reasoning. Nor are there clear analogues in other segments of natural language. For instance, to be a competent user of 'tummy' in English, I need to grasp not only that it denotes stomachs but also that it belongs to infant-directed speech; 'tummy' is unsuitable by virtue of its meaning for me to use at a gastroenterologist's (Ferguson 1964). The difference between committal and non-committal uses of normative language isn't this kind of conventionalized difference in register. Although I will be a fair target of criticism if I assert 'I ought to do that' in a seemingly committal way without being practically committed in the relevant ways, my *faux pas* (or, rather, lack of any *pas*) is just a special case of violating a general sincerity condition on speech acts, nothing specific to normative language.

It doesn't help the case to claim that committal uses express concepts that are crucial to a certain kind of practical thought. For instance, if non-committal uses of normative sentences just update our picture of what is to be done relative to some standard, their acceptance needn't involve more than recognizing those verdicts. Even if the acceptance of a normative sentence involves more in other cases, features specific to those uses aren't evidence about the meanings of normative terms (contrary to Gibbard 2003: 153). There is also no systematic link between whether *ought* is used committally or non-committally and what kind of normativity it is used to express. People can get very committed about what they should do by the formally normative standards of feline excitement or espresso excellence. In sum: to solve the Problem of Diverse Uses, it isn't enough just to distinguish committal and non-committal uses of normative terms.

23 For a useful exercise, compare how different this case is from whether indirect speech acts (such as using 'Can you pass the salt?' to issue a request) are conventional. See especially Simons and Zollman (2019: 13-22).

What I take my discussion so far most strongly to suggest is that thin normative terms' association with their practical role is unlikely to be a phenomenon at the level of descriptive semantics. In that case we should expect committal and non-committal uses of thin normative terms to be uniform in their descriptive semantics. This is dialectically significant. For if the relevant practical upshots are absent in many contexts, then semantic uniformity would be easy to explain if the presence of practical upshots in other contexts weren't a broadly semantic phenomenon. The pressure is therefore on for the alternative broadly semantic view that thin normative terms' association with their practical role is a metasemantic phenomenon instead.

5. Practical Role and the Metasemantics of *Ought*

To assess the metasemantic view, we require some idea of what kind of descriptive semantics thin normative terms have. I'll use as my sample the standard sort of descriptive semantics for deontic modal expressions in English. I'll argue that nothing in this formal semantics or plausible supplementary assumptions supports treating those uses of deontic terms which are associated with a practical role as semantically or metasemantically exceptional.

5.1. *The Standard Semantics for Deontic Modals*

The dominant view in linguistics and philosophy of language is that expressions such as *ought/must/should/have to/...* are context-sensitive modal operators. For instance, in *Ought(p)*, *ought* is a non-propositional device for shifting the circumstances relative to which *p* is to be semantically evaluated at a world of evaluation *w*. Relative to *w*, 'I eat more iron' might be false but 'I ought to eat more iron' true. On the account due to Kratzer which is the starting point for many further developments, *ought* determines the set of worlds relevant to modeling those circumstances along two contextual parameters (Kratzer 1981; 1991). Intuitively, the first is which worlds matter, the second is how we rank them. In the terminology of the theory, the first is a *modal base*: a function *f* from *w* to a set of worlds compatible with *w*. This is a restriction on *w* by whatever background conditions may be selected as relevant in context *c*.²⁴ The second is an *ordering source*: a

²⁴ The relevant notion of context here includes both linguistic elements (such as previous utterances and topics of conversation) and extralinguistic elements (such as the intentions of the speaker,

function g from w to the set of best worlds in the modal base. This yields an ordering over worlds in terms of whichever standards are selected in c . Then “ fg -compatible worlds” will be the set of best worlds determined by c and w .

The semantic value of $Ought(p)$ relative to a context-world pair can now be represented by the following modal rule:

[[$Ought(p)$]] ^{g} is true iff all of the fg -compatible worlds are p -worlds.²⁵

For instance, suppose that ‘You ought to abstain from premarital sex’ is uttered when the question is what it takes to live a Roman Catholic life. The modal rule for *ought* then tells us that the sentence is true just in case you abstain from premarital sex in all worlds that are compatible with something like the following two restrictions. First, a restriction to worlds in which you can engage in premarital sex (f). Second, a restriction of those worlds to ones that rank highly by Roman Catholic values (g). Interpreted relative to this context, the sentence is, plausibly, true. But now take a different context where the question is what it takes to live a life of high sensory pleasure. Interpreted relative to this context, the sentence is, plausibly, false. The modal rule allows *ought* to be interpreted relative to any ordering with certain formal properties: how well something conforms with the 1911 Boy Scouts Handbook; how excited my cat gets; what is better for a given agent; what is likely to happen anyway (normative fatalism is a thing for some); the standards endorsed by the speaker; some specific moral ideal; the most fundamental moral standards (whatever they may be); or whatever standard (objective or otherwise) is expressed in ‘I know I morally ought to ϕ , but ought I really to ϕ ?’²⁶ This captures our recognition that nearly anything can be felicitously called by terms like *ought*.

objects discernible in the environment, the social setting, and common presuppositions).

25 This formulation is from Chrisman (2018: 410). Double brackets stand for a function that maps a linguistic expression to its semantic value. My main points will go through even if *ought* needn’t order possible worlds (Ridge 2014: 33-35). I suspect they’ll also go through even if the Kratzer semantics is replaced by one that connects deontic modals to verdicts of a theory practical rationality, e.g. in ways proposed by Cariani (2016), Charlow (2016), or Lassiter (2016). For these accounts appeal just to formal decision-theoretic constraints on the ranking of options. Issues about deontic logic raised by Broome (2016: 11-13) may be more tricky. See also Bronfman and Dowell (2018) for responses to various challenges to Kratzer-style semantics for deontic modals.

26 As Worsnip (2019) puts it, the contextually selected standards needn’t be “parochial” but can be “aspirational”. As these examples suggest, standards needn’t be reified, they may lack handy natural language labels, and their contents may not be immediately obvious or transparent (cf. Evers 2014: 295-96).

5.2. Two Metasemantic Options

Suppose *ought* has roughly the above kind of descriptive semantics. What kind of metasemantic function might the practical role of *ought* play with respect to it? Its practical role won't explain any particular context-invariant features of its meaning, such as why its meaning includes parameters for which worlds matter and how we rank them. Two options seem to remain. One has to do with the metasemantics of the context-sensitivity of *ought*. Perhaps its practical role contributes to explaining its semantic value specifically in its committal uses. The other option is quite different. Perhaps the practical role of *ought* instead plays some significant role in explaining why a given formal semantics is an appropriate model for *ought* in the first place.

The first view is initially attractive. It promises to solve the Problem of Diverse Uses. Whether a use is committal or not is a difference in context. We might then think that when *ought* is used in a committal way, this can make a difference to the values of its contextual parameters. In this way, the practical role of *ought* might contribute to explaining its semantic value in some cases but not others. To assess this, it is important to see that the modal rule makes no semantic difference between committal and non-committal uses of *ought*. Nothing in the standard semantics *requires* the contextually supplied ordering sources or modal bases to be such that the results of feeding them into the modal rule for *ought* align with anyone's practical commitments.²⁷ In non-committal uses, the values of contextual parameters are fixed by factors other than practical commitments. Even in committal uses,

²⁷ One might claim that such a requirement can be motivated as an external constraint on the semantics. Things can be ranked in terms of badness rather than goodness. Formally speaking such a ranking is eligible to help determine the semantic value of *ought* in certain contexts. But could there really be a context in which a competent speaker who says that we ought to do something is thereby asserting that it is what we do in all the maximally *bad* worlds compatible with the modal base? If not, our overall semantic theory should posit a restriction on what rankings can supply the ordering source. One might then suggest that such a restriction has something to do with the relation between these rankings and the practical commitments of some contextually determined agent. In response, let's consider again Iago who is committed to sheer malignity for its own sake. Sincerely uttering 'You ought to kill, and it is better if you kill than if you study' is morally misguided and perhaps for that reason uncommon. But it needn't be linguistically incoherent if the ordering source is the same. (The sentence is true if killing is ranked highly by the contextually selected ordering source, and false if not.) By contrast, 'You ought to kill, but it is better if you study than if you kill' is incoherent if the ordering source is the same. Various semantic connections between *good* and *ought* constrain their interpretations (Lassiter 2017: ch. 8) independently of the practical roles of these terms. So explaining why *ought*-claims rarely assert what is done in maximally bad worlds doesn't require invoking the practical role of *ought* even in committal uses. It cannot require that in non-committal uses anyway.

appeal to practical role shouldn't single out any specific standard to be fed into our formal semantics. Often the selection of an ordering source is a substantive normative matter, at least beyond specifications such as 'the correct standards' or 'the most fundamental standards'. By the same token, however, nothing in the standard semantics *rules out* that contextually relevant motivations or practical commitments may contribute to fixing the semantic value of *ought* in committal contexts. So do they?

If you think that the practical role of normative language sets it apart from non-normative language in some broadly semantic way, you'll want its practical role to be a distinctive and comparatively significant feature of its meaning. Locating the difference in the metasemantics of context-sensitivity fails to fulfill this promise. The problem arises from perfectly general considerations. What is required for interpreting utterances featuring context-sensitive terms? A variety of contextual inputs must combine to set the values of contextual parameters and thereby transform non-propositional semantic values into compositional semantic values that are truth-apt. This may be fairly straightforward for simple demonstratives like *this*, or even conventionalized for indexicals like *I* and *now*. In general, however, just which contextual factors can contribute to determining the semantic value of an expression and how they may combine to do so defies simple generalization. Factors that are in general eligible are many. They include common presuppositions, social influences, the structure of the preceding discourse, salient objects, and more (Glanzberg 2007). The metasemantics of the ordering source parameter for *ought* is indirect in the same way: multiple factors influence how its value is set in context.²⁸ That must be so since in non-committal uses the semantic value of *ought* is determined by factors other than its practical role anyway. Suppose we have been discussing capitalism and I say 'We ought to prioritize profit over workers' interests'. It may well be clear from the preceding discourse, or the pins on my jacket, that I'm a committed critic of capitalism. I should then be construed as talking about what follows from capitalist values.

If factors other than the practical role of *ought* suffice to determine semantic value in non-committal contexts, we can expect them sometimes to override contextually relevant practical commitments even in committal contexts. In general,

28 On indirect metasemantics in general, see Glanzberg (2007, 2020). King (2014) defends a metasemantics on which the values of contextual parameters are fixed by speakers' intentions. But even on that view, it is a substantial issue whether or when the ordering source parameter is fixed by practical commitments.

different factors may pull in different directions and, thus, require balancing. For instance, social influences and speaker intentions may conflict. No particular factor enters such a calculation from a position of privilege, unless the conventional meaning of the expression in question so dictates. But we have seen that the practical role of *ought* isn't part of its conventional meaning, nor is its presence marked otherwise by linguistic convention. So when practical commitments help us to limit the set of contextually relevant ordering sources in some way, any other way of limiting them in that way would have done just as well. Thus, even when the practical role of *ought* plays an indirect metasemantic role, it plays no distinctive or comparatively significant role in determining the semantic value of *ought*.²⁹ So the claim that the practical role of thin normative terms is a particularly significant feature of their meaning cannot be vindicated in the metasemantics of their context-sensitivity.³⁰ My reasons for this claim may not be conclusive. The issues here are subtle and the metasemantics of context-sensitivity is an underexplored topic in general. But the reasons look strong enough to recommend looking elsewhere.

The second metasemantic view is that the practical role of *ought* is part of what explains why the dominant sort of formal models for modal language provide a good descriptive semantics for terms like *ought* in the first place. For instance, an inferentialist metasemantics might hold (roughly) that a given formal semantics is a good model for *ought* because it explicates a certain kind of inferential structure in which *ought* is nested (Chrisman 2015; Tiefensee 2016, 2021). Or, metasemantic expressivism might hold (again, roughly) that a given formal semantics is a good model for *ought* because it appropriately mirrors the structure of mental states which *ought* expresses (Ridge 2014; Köhler 2018).

If the association between thin normative terms and practical role is to be metasemantic in this way, a further claim is required. The explanation of why a certain formal semantics is a good model for *ought* must invoke specifically practical inferences or motivational states. But, on the face of it, an explanation that unifies the committal and non-committal uses of *ought* doesn't need this further claim.³¹ The standard semantics implies that committal and non-committal uses of

29 Practical upshots can also play an indirect metasemantic role even if they are heterogeneous across contexts in ways that wouldn't support explanatory generalizations.

30 In §6.1 I'll suggest that pragmatic accounts of the practical role of *ought* can accommodate the above kind of weak and indirect metasemantic role.

31 Silk's (2016) general framework for theorizing about normative language looks compatible with this alternative. Silk proposes that "normative uses" of *ought* present the speaker as endorsing the standards that justify the *ought*-claim in question (2016: 130-132, 137). Justification is an

any deontic sentence of the form *Ought(p)* have the same semantic value when the modal base and the ordering source are the same. Explaining non-committal uses only requires invoking theoretical commitments and cognitive states. By parity, that should suffice also for explaining committal uses. The standard semantics doesn't care about this distinction between uses. So, on the face of it, explaining why it is a good model for *ought* shouldn't require invoking practical role. If that is right, it would complete my case that nothing in the standard semantics for *ought* supports treating those uses which are associated with a practical role as semantically or metasemantically exceptional. I'll now argue that existing forms of metasemantic inferentialism and expressivism don't support assigning the practical role of *ought* a significant metasemantic function of this type. The Problem of Diverse Uses still pushes us not to build the practical role of thin normative terms into our overall semantic theory for these expressions.

5.3. Ideational Expressivism

First consider "ideational expressivism" due to Ridge. Ideationalism is a general metasemantic theory according to which "facts about the semantic contents of meaningful items in natural languages are constituted by facts about how those items are conventionally used to express states of mind" (Ridge 2014: 107). What Ridge's expressivism adds to this general framework is the claim that normative sentences have their meaning in virtue of expressing a certain kind of non-representational, motivational states while descriptive sentences have their meaning in virtue of expressing representational states (Ridge 2014: 110-11). My discussion won't touch ideational expressivism in its most general form, but only those specific forms that assign significant metasemantic work to the practical role of normative terms.

To illustrate, consider Ridge's broadly Kratzer-style descriptive semantics, on which any use of *ought/must* says roughly something of the following form: "Any standard of contextually specified kind S would, relative to a contextually specified set of background information or facts B, recommend/require X" (Ridge 2014: 28). Formally normative standards, such as legal standards, have no essential link to a practical role. But some standards do. If a sentence like 'You ought to give

intranormative notion, and endorsement (which Silk usually applies to endorsing a body of *information*) doesn't entail affirming specifically practical commitments. Nor is that entailed by his account of normative uses as *presupposing* a set of standards endorsed in the context, analyzed as a contextual parameter whose value speakers can manage by exploiting their mutual grammatical and world knowledge and general pragmatic reasoning skills (Silk 2016: 131, 145).

to charity' is used to make an "all things considered practically normative" claim, it can be paraphrased as "Any acceptable standard of practical reasoning would, given contextually specified circumstances, recommend that you give to charity" (Ridge 2014: 40). Normative contexts are to be understood narrowly as those in which the relevant sort of standard is 'any acceptable standard of practical reasoning'.³² In such contexts, Ridge claims, *ought* and *must* express a certain kind of motivational states (Ridge 2014: 19-21). In other contexts, *ought* and *must* are used non-normatively.

So Ridge acknowledges that terms like *ought* can be put to diverse uses. However, using ideational expressivism to explain why the standard semantics is a good model for *ought* requires no reference to the practical role of *ought*. No such reference is required for non-committal uses of *ought* to state recommendations. A further good question is whether one can use *ought* non-committally to state what would be recommended by any acceptable standard of practical reasoning.³³ (Why not?) To assess *ought*-sentences that count as normative under Ridge's narrow definition, consider his preferred account of expression:

Accountability Expression: A declarative sentence 'p' in sense S in a natural language N used with assertive force in a context of utterance C expresses a state of mind M if and only if conventions which partially constitute N dictate that someone who says 'p' in sense S in C with assertive force is thereby *liable* for being in state M. (Ridge 2014: 109.)

Whether an assertive use of *ought* states what would be recommended by any acceptable standard of practical reasoning is a function of features of the context of utterance which don't normally ground linguistic convention.³⁴ So the relevant

32 A standard of practical reasoning is a standard which is treated by the agent as a standard of practical reasoning and which can function in a certain action-guiding way (Ridge 2014: 40).

33 According to Ridge, to decide that a course of action is acceptable in a given set of circumstances is "in some sense to decide that the course of action is not ruled out for purposes of your deliberation – that it is still 'on the table'" (Ridge 2014: 41). Note that some non-committal uses of *ought* may satisfy this description. In (5), for instance, the speaker isn't (yet) engaged in deliberation but treats going shopping as being still on the table.

34 Ridge notes that we can use modifiers like *really* and *genuine* to flag that we are using *ought* or *must* in his narrowly normative sense. For instance, we can acknowledge that etiquette requires something but query whether that is what we *really* ought to do or have any *genuine* reason to do (Ridge 2014: 20). But this effect of *really/genuine* may be better explained in the pragmatics. It tells us nothing much about the meanings of assertive utterances of *ought*-sentences without such modifiers. Nor is a narrowly normative meaning of 'really ought to' a function of semantic

linguistic conventions would have to enter someplace else. Ridge does allude to “a theory of how the linguistic conventions provide a function from contexts of utterance to states of mind for any given sentence of the language” (2014: 130). But what are the linguistic conventions of English which dictate that assertive utterances of sentences of the form *Ought(p)* to state what would be recommended by any acceptable standard of practical reasoning make the speaker liable to be, specifically, in a motivational state? We saw in §5.2 that they won’t be conventions pertaining to how semantic value relative to context is computed from context as an operation on the modal rule for *ought*.³⁵ A more general point can be made even in the absence of the kind of general theory to which Ridge alludes. Whether ideational expressivism explains why the standard semantics is a good model for *ought* doesn’t depend on there being conventions that map the relevant uses of *ought* specifically onto motivational states. It is hard to see what conventions of English should rule out accounting for assertions concerning what would be recommended by acceptable standards of practical reasoning in terms of non-motivational states of mind, such as beliefs about what such standards recommend. In this vicinity lie also good but

composition. What *genuine* and *really* contribute to composition in these contexts is the same as when we ask whether refraining from noting a smudge on someone’s face is *genuinely* a polite thing to do, or whether going on a bar crawl against my parents’ wishes is *really* a fun thing to do.

- 35 Ridge’s own view may be more like the view that the practical role of *ought* does work in the metasemantics of its context-sensitivity. In the main text I discuss the prospects of recruiting his view to a different service, and so am not attributing all of the moves I discuss to Ridge himself. But ideational expressivism faces trouble also with showing that the practical role of *ought* plays a special role in explaining semantic value relative to context. Ridge individuates orderings finely: “A moral standard provides a different ordering from a legal standard, even in the case in which the law requires all and only the morally required actions” (Ridge 2014: 36). Consider two speakers who both assert the same *ought*-sentence, such as ‘One ought to report one’s liability’. The contextually specified standards may be isomorphic in the orderings they induce but such that one speaker treats the standard as an acceptable standard of practical reasoning whereas the other doesn’t. (Two speakers can differ in this way with respect to moral standards, legal standards, and more.) Ridge’s view implies that the contexts involve different standards, so the two utterances have distinct semantic values. That may be fine if the contexts are separate, though I myself find it odd if solely a difference in the motivational states which the speakers are liable to be in implies that they say different things. But suppose the speakers are involved in a conversation. It would be implausible to posit constant context shifts depending on whether the speaker at the time treats a specific standard as an acceptable standard of practical reasoning. Modeling such a conversation requires just a shared conception of what actions are recommended and a way of tracking what the speakers are planning to do about that. Moreover, it should be possible to report the conversation by saying ‘Timmy and Tammy both think that one ought to report one’s liability. Tammy is planning to report hers; Timmy isn’t.’ If their utterances said different things, such a report should be either false or odd. But it is an accurate report that would be felicitous in the described context.

underexplored questions about how philosophically committal we should want our metasemantic theories to be.

Given Accountability Expression, a sentence can express *M* without its being the case that the speaker actually is in *M*. One might then reply that *ought*-sentences have the semantic values they have in virtue of expressing the relevant kind of motivational attitudes in all contexts. What happens in non-committal uses is that further contextual information implies that the speaker's state of mind doesn't exemplify that structure. But this reply is problematic. First, Accountability Expression doesn't support it. In §4 we saw some reason to think that the liabilities incurred by non-committal uses of *ought* aren't plausibly parasitic on those incurred by its committal uses. We still have no evidence that a parasitic relation holds in such a way that non-committal speakers are liable to be in a motivational state by linguistic convention. So far as the conventions that constitute English go, deontic assertions can well uniformly express states representing what verdicts follow from contextually specified standards. Second, the standard semantics treats deontic talk as expressing what requirements follow from such standards. Nothing in it suggests that this semantic function of deontic talk is somehow parasitic on some more immediately practical function.³⁶ Explaining why it is a good model for *ought* thus doesn't require invoking the practical role of *ought*. I conclude that ideational expressivism doesn't help broadly semantic views to solve the Problem of Diverse Uses.

5.4. Inferentialist Metasemantics

Now turn to inferentialist metasemantics for normative language. Chrisman takes the truth-conditions predicted by the Kratzer-style modal rule to articulate how speakers “have to think and reason in order to satisfy the implicit conceptual commitments affirmed by using ‘ought’ to make an assertion in ordinary discursive practice” (Chrisman 2018: 415; cf. Chrisman 2015: ch. 5). My discussion won't touch metasemantic inferentialism in this general form, but only those specific forms that assign significant metasemantic work to the practical role of normative terms.

36 Ridge suggests that normative discourse might be like imperatival discourse, in that although imperatives can serve subsidiary communicative functions, these are “somehow parasitic on the more immediately practical function of the imperative form” (Ridge 2014: 21). Here I question the parallel. See also note 15 for the possibility that the practical function of the imperative form is a kind of pragmatic phenomenon.

Existing inferentialist metasemantics for normative language take this form. Chrisman acknowledges that some uses of *ought* merely commit one to concluding the verdicts of the contextually selected ranking. But he proposes that the commitments affirmed by “genuinely normative” uses could “be conceived as commitments to reason practically in certain ways” (Chrisman 2018: 416). The key idea is that the general metasemantic function of *ought* is still the same across these local differences. Tiefensee similarly proposes to understand evaluative terms such as *good* in terms of a general metaconceptual function of structuring and explaining the legitimacy of certain language exit transitions to intentions, actions, and so on (Tiefensee 2021).³⁷ Transposed to the deontic key of this paper, this is to analyze *ought* as a linguistic instrument for explicating certain commitment structures. Asserting ‘I ought to help my sister’, for instance, explains why a commitment to a premise like ‘Babysitting the boys will help my sister’ provides a reason to enter into the commitment that I shall babysit the boys, and likewise for premises concerning what one has promised, and so on (cf. Tiefensee 2021: 122, 125).

These views acknowledge that terms like *ought* can be put to diverse uses. But they, too, struggle with the Problem of Diverse Uses. Recall that on the standard semantics, an *ought*-sentence will have the same truth-conditions across committal and non-committal uses when contexts supply the same modal base and ordering source. ‘We ought to report our liability’, for instance, says that reporting our liability ranks highly on the contextually determined standard. On the face of it, explaining why the standard semantics is a good model for (deontic) *ought* should then be insensitive to differences in practical upshot. Metasemantic inferentialism would be more general and unified if it modeled *ought* as explicating structures whose language exit transitions consist in theoretical commitments. Examples would be commitments to form certain beliefs about what is required or recommended and to recognize reasons for action which correspond to the same ordering source.³⁸ This would account for one central function of deontic vocabulary of English: stating what requirements or recommendations follow from various sorts of standards to

37 For reasons of space, I simplify a great deal. I bracket Tiefensee’s (2021) view that we need to begin with words like *good* because *ought* presupposes evaluative orderings. As far as I can see myself, an ordering needn’t itself have evaluative content, and only a notion of ranking highly on an ordering (rather than, e.g., being among the *best* in any evaluatively laden sense) is required for identifying the set of recommended items.

38 Contrary to what Tiefensee (2016: 2444) seems to suggest, it isn’t sufficient for practical commitment in the relevant sense that a claim of the form ‘I ought to ϕ ’ gives a *pro tanto* justification or reason to ϕ . Recall note 7.

which the interlocutors needn't be practically committed. Its further function to convey practical commitments could well figure in some other part of our overall theory of meaningful communication, such as pragmatics.³⁹

For metasemantic inferentialism to solve the Problem of Diverse Uses, one of two different views must instead be true. One is the view that Chrisman gestures at: committal and non-committal uses of *ought* explicate different commitment structures; those involved in committal uses feature not only theoretical but also practical commitments. The other is the view that they explicate the same commitment structures and these are indispensably practical. Anyone making a deontic assertion using *ought* would normally be understood as undertaking the commitments which the sentence makes explicit, but further contextual information may suggest that the speaker's assertion only reports that structure. However, neither view explains why the standard semantics is a good model for *ought*. Nor does either view support the distinct view that the practical role of *ought* constrains our overall semantic theory in some other way.

If practical commitments were indispensable to the commitment structures which *ought* explicates or affirms, the opposite thesis should fall to a *reductio*. But it doesn't. Suppose that the language exit moves in the commitment structures which deontic *ought* makes explicit consist wholly in theoretical commitments. Deontic *ought*-claims would then make the same kind of commitment structure explicit in both committal and non-committal uses, and language exits would still be to states that are about actions. This variant of metasemantic inferentialism could equally well say that anyone making a deontic assertion using *ought* would normally be understood as undertaking the commitments which the sentence makes explicit, but further contextual information may suggest that the speaker's assertion only reports that structure. For it is perfectly compatible with all of the above that speakers who assert *ought*-sentences would normally be understood, on the basis of pragmatic inferences, as undertaking *also* some practical commitments. So the appeal to practical commitments doesn't seem indispensable in explaining why the standard semantics is a good model for *ought*.

39 Note a related issue about the metasemantic adequacy of versions of inferentialism which appeal to indispensably practical commitment structures. If modal expressions are polysemous across different "flavors" of modality, the polysemy doesn't seem to arise from differences in commitment structures. The arguments for modal polysemy in Viebahn and Vetter (2016) aren't like that, for instance.

This leaves the view that committal and non-committal uses of *ought* explicate different commitment structures. For this to solve the Problem of Diverse Uses, the distinction between committal and non-committal uses of normative language must be semantically significant. Chrisman suggests that “genuinely normative” uses of *ought* differ conceptually from the rest. In some sense that must be right, since there is a distinction between committal and non-committal uses to be drawn. But this minimal sense isn’t enough here. We have seen that the distinction isn’t lexically encoded or marked by linguistic convention. We have also seen that prospects are poor for showing that non-committal uses of normative language are semantically parasitic on some more immediate practical function of normative language. Nor have we been shown any other evidence that the distinction plays a role in explaining why the standard semantics is a good model for *ought*. As we have seen, if a word can be used in certain contexts to affirm or explicate certain implicit conceptual commitments, this isn’t automatically something to be reflected in our overall semantic theory. So even if inferentialism explains why the standard semantics provides a good model for terms like *ought*, we have yet to see why that explanation must feature not only theoretical but also practical commitments. Differences between committal and non-committal uses might instead be captured in some other part of our overall theory of meaningful communication.

Time to conclude. In this section I introduced the dominant sort of formal semantics for deontic modal expressions like *ought*. I then argued that explaining why deontic *ought* has this kind of descriptive semantics doesn’t require appeal to the practical role it plays in many contexts. My argument went through all the places I can think of where practical role might show up: the metaseantics of the context-invariant semantic features of *ought*; the metaseantics of its context-sensitivity; and explaining why a given kind of formal semantics is in general a good model for *ought*. Existing expressivist and inferentialist metaseantics for normative language don’t support the claim that the practical role of such language is a distinctive and particularly significant feature of its meaning. And yet metaseantics seemed to be the most promising place for vindicating broadly semantic accounts of the practical role of thin normative terms. I cannot claim to have conclusively ruled out that the practical role of normative language is a metaseantic phenomenon. Caution is due: the issues here are subtle and underexplored and my argument has many moving parts. But at minimum my argument sets a demanding bar to be met.

6. Broader Implications

I'll close by discussing the prospects of explaining the practical role of normative language in pragmatic rather than broadly semantic terms and highlighting some broader implications of my negative argument for metaethics and the philosophy of normativity.

6.1. Looking Ahead to Pragmatics

The challenges that I have raised against broadly semantic accounts of the practical role of thin normative terms bite because generalizations which don't rise to full generality are normally more apt to be explained in the pragmatics. Explanations of linguistic phenomena that appeal to general principles of communication instead of special theoretical posits are more unified and parsimonious. So if a phenomenon can be explained in the pragmatics, then (all else equal) it shouldn't be explained semantically. The obvious follow-up question is whether the practical role of thin normative terms can be adequately explained as a pragmatic phenomenon. If it couldn't, that would support introducing suitable posits into our overall semantic theory.

The general character of pragmatic reasoning concerns what can be inferred from the following sort of premise: the speaker believed that uttering a given particular sentence in the given context was best for promoting her conversational ends. Determining whether the practical role of normative speech can be derived through broadly pragmatic reasoning thus requires answering the following sorts of questions: In pursuit of what conversational ends may we engage in normative speech? By means of what kind of normative speech do we take those ends to be best achieved? Answering these questions is well beyond the scope of this paper. But there is reason to be optimistic about a pragmatic account. Everyone will need an account of what makes particular *uses* of words normative. A very wide range of non-normative terms can be contingently associated with practical upshots, given a suitable context. For instance, saying that a knife is sharp can convey a disposition or commitment to use it when it is common ground that a hard loaf of bread needs slicing. Any association between uttering 'This knife is sharp' and such practical upshots is pragmatic. Plenty theoretical space thus exists for a pragmatic account of the practical role of normative language.

One possibility is to adapt the pragmatic account from Finlay (2014: ch. 5). Finlay develops a certain kind of “end-relational” semantics of normative terms. He then argues that we can predict and explain the practicality of normative talk as a feature of “how we use normative language in context, pursuing our desired ends” by using “only maximally simple and conservative principles of pragmatics” (Finlay 2014: 116). The broad kind of pragmatic story might well be similar regardless of whether the relevant inputs from context are ends, rankings, or whatnot. Here is another possibility. Suppose we are particularly concerned to explain why we normally don’t need to work it out from scratch why the speaker of a bare first-person normative utterance is committed to acting or reacting accordingly. Here we might be able to adapt accounts which treat certain features of normative discourse as “generalized” pragmatic features. When a pragmatic feature is generalized, it arises from saying a certain thing in the absence of conversational moves or other special circumstances that would revoke the implication.⁴⁰ Strandberg (2012) argues that the motivational implications of moral assertions are generalized conversational implicatures. Väyrynen (2013: ch. 5) argues that the evaluative implications of “thick” terms, such as *chaste* and *generous*, are default but defeasible “not-at-issue contents” of their assertive uses in normal contexts. The practical upshots with which thin normative terms are associated might be susceptible to one of these types of explanation.

A fuller pragmatic account of the practical role of thin normative terms is beyond the scope of this paper. But another optimistic note worth sounding is that the practical role of thin normative terms can both be a pragmatic phenomenon and play the kind of indirect metasemantic role described in §5.2. Suppose someone makes a claim about what we ought to do but you don’t immediately get just what they are saying. One way you can try to understand them is by searching for practical commitments that would make sense of the utterance. For instance, information about the speaker’s or another salient agent’s practical commitments might help to identify standards which would recommend the action in question.⁴¹ (Other ways include a closer look at the preceding discourse, and so on. Recall the discussion of indirect metasemantics of context-sensitivity in §5.2.) Speakers often can expect their audience to interpret their utterances in this way by exploiting their

40 Here I generalize Grice’s notion of generalized conversational implicature (Grice 1989: 37-38).

41 See Silk (2016: 126-32) for one more detailed discussion of these kinds of aspects of normative discourse.

grammatical knowledge (that *ought* requires a ranking, and so on), world knowledge, and general pragmatic reasoning skills. The general point is that a context-sensitive expression with an indirect metasemantics may well be such that a factor can help fix the value of a contextual parameter in a context of its utterance and yet get communicated by the utterance as a matter of pragmatics. My arguments suggest that contextually relevant practical commitments may be one such thing.

6.2. *Implications for Metaethics and the Philosophy of Normativity*

If my negative arguments against broadly semantic accounts of the practical role of normative language are on the right track, they put significant pressure on a range of prominent accounts of normative language. As my discussions of Hare, Blackburn, Gibbard, and Ridge suggest, they put pressure on expressivist accounts of the meaning of normative language in terms of plans or other conative states.⁴² As my discussions of Wedgwood, Chrisman, and Tiefensee suggest, they put similar pressure on versions of conceptual or inferential role theories which aim to explain the meaning of normative language in terms of its practical role. These arguments also challenge the motivation for those moral error theories according to which it is a non-negotiable commitment of our moral discourse that such discourse plays a certain kind of practical role.

My negative arguments also show that we need to be much more careful in moving between claims about normative thought and claims about normative talk. To be clear, the arguments are compatible with there being normative concepts to which a practical role is essential. There are downstream questions, such as how we might manage to express such concepts if expressing them isn't a semantic or metasemantic feature of normative language. But the immediate lesson is just that even if there are OUGHT-concepts which are conceptually linked to motivation or practical commitment, this link is reflected neither in our descriptive semantics nor our metasemantics for the corresponding segment of English. Normative discourse can then be in perfectly good working order without predicates that are associated with certain practical roles in any broadly semantic way. This doesn't seem descriptively inadequate. For instance, my arguments raise no deep problems for explaining why *ought* is characteristically used to offer advice, *must* to issue more

42 Some recent work interprets expressivism more broadly as a view on which states of mind are fundamental in a theory of meaning (cf. Ridge 2014; Charlow 2015; Silk 2016). As emphasized in §5.3, my arguments touch such views only insofar as they appeal to essentially practically oriented states of mind.

insistent sort of requirements and orders, and so on. If their characteristic semantic function is to describe recommendations or requirements that follow from some assumed ranking, they will be naturally suited for directive conversational effects.

How my arguments bear on the philosophy of normativity is more complex. Some deny that discussions of how normative terms work in natural language are of much import to philosophers of normativity. For example, Broome argues that philosophers of normativity “do not have to be much bothered by the fluid and contextual nature of ‘ought’ in common English” (Broome 2016: 10). There is a central practical meaning of *ought*, specified by an “enkratic” rational requirement on intentions. The philosophy of normativity is principally concerned with the metaphysical nature of this ought, what determines when it obtains, and so on.

I won’t here discuss whether Broome is right that the enkratic OUGHT is at the centre of normativity. The main lesson of my paper is that even if a given OUGHT is privileged with respect to the philosophy of normativity, it won’t be linguistically privileged. For example, even when *ought* is used to express an authoritative OUGHT which resolves conflicts between explicitly relativized notions like “moral ought” and “prudential ought” and specifies what we ought to do without qualification, that is just a special case of the standard semantics. For it to be a well-behaved use of *ought*, it must induce a ranking of the things that matter in the context. (Broome 2016: 10 allows that there may be such a “final ordering”.) We might not have a dedicated name for such a ranking. But if there wasn’t one, this OUGHT wouldn’t be able to resolve conflicts between various explicitly relativized oughts. Regarding the enkratic OUGHT, it wouldn’t be clear why rationality should require that you intend to do what you judge you ought to do (in this sense) unless it was something that mattered highly in the context.

We probably won’t be able to say much at all about the content of the ranking induced by any such privileged OUGHT without doing substantive normative theory. But that is as it should be. We shouldn’t expect semantic theory to tell us much about normative theory. Nor, however, should normative theory do all that much to shape or drive semantic theory. So I counsel caution and reconciliation. Even if philosophers of normativity have identified a specific concept which lies at the centre of normativity and can be expressed by *ought*, this implies little regarding how its practical role is reflected in language. Even if a given OUGHT is individuated by a practical role that involves distinctive conceptual commitments, it doesn’t

follow that those commitments are to be reflected in our overall semantic theory. Just as semanticists should avoid exaggerated claims about normativity, philosophers of normativity should be sensitive to how the core features of their central concepts are reflected in language.⁴³

43 This paper was four years in the making, which feels long for something that began as a fun side project and remained a side project. My bumbling early runs at its core argument probably count as a distinct paper even by a fairly coarse-grained individuation of philosophy papers. That predecessor (“Evaluative Terms and Normative Role”) received valuable feedback from audiences at University of Leeds, University of Geneva, Uppsala University, Humboldt University of Berlin, and the New York Philosophy of Language Workshop. Something more like the present paper was helped along by discussions at University of Cambridge and workshops at NYU Abu Dhabi and Frankfurt School of Finance and Management. I’m grateful to Alex King as my commentator in Abu Dhabi and to Sarah Buss, Matti Eklund, Will Gamester, Camil Golub, Matthew Kramer, Nick Laskowski, Tristram McPherson, Eliot Michaelson, Christine Tiefensee, Robbie Williams, and Jack Woods for conversations and/or comments on various versions of the material. Comments from anonymous referees for *JESP* and many other journals along the way also led to significant improvements. Many thanks to you all.

References

- Alwood, Andrew (2016), "Should Expressivism Be a Theory at the Level of Metasemantics?", *Thought* 5: 13-22.
- Bar-On, Dorit and Chrisman, Matthew (2009), "Ethical Neo-Expressivism", *Oxford Studies in Metathics* 4: 133-165.
- Bedke, Matthew (2019), "Practical Oomph: A Case for Subjectivism", *Philosophical Quarterly* 69: 657-677.
- Blackburn, Simon (1998), *Ruling Passions* (Clarendon Press).
- Bronfman, Aaron and Dowell, J. L. (2018), "The Language of "Ought," and Reasons", in *The Oxford Handbook of Reasons and Normativity*, ed. D. Star (Oxford University Press), 85-112.
- Broome, John (2016), "A Linguistic Turn in the Philosophy of Normativity?", *Analytic Philosophy* 57: 1-14.
- Cariani, Fabrizio (2016), "Deontic Modals and Probabilities: One Theory to Rule Them All?", in Charlow and Chrisman (eds.), 2016: 11-46.
- Charlow, Nate (2015), "Prospects for an Expressivist Theory of Meaning", *Philosophers' Imprint* 15 (23): 1-43.
- Charlow, Nate (2016), "Decision Theory: Yes! Truth Conditions: No!", in Charlow and Chrisman (eds.), 2016: 47-81.
- Charlow, Nate and Chrisman, Matthew (eds.) (2016), *Deontic Modality* (Oxford University Press).
- Chrisman, Matthew (2015), *The Meaning of 'Ought'* (Oxford University Press).
- Chrisman, Matthew (2018), "Two Nondescriptivist Views of Normative and Evaluative Statements", *Canadian Journal of Philosophy* 48: 405-424.
- Copp, David (2001), "Realist-Expressivism: A Neglected Option for Moral Realism", *Social Philosophy and Policy* 18(2): 1-43.
- Dreier, James (1990), "Internalism and Speaker Relativism", *Ethics* 101: 6-26.
- Eklund, Matti (2017), *Choosing Normative Concepts* (Oxford University Press).
- Evers, Daan (2014), "Moral Contextualism and the Problem of Triviality", *Ethical Theory and Moral Practice* 17: 285-297.
- Ferguson, Charles A. (1964), "Baby Talk in Six Languages", *American Anthropologist* 66: 103-114.

- Finlay, Stephen (2014), *Confusion of Tongues: A Theory of Normative Language* (Oxford University Press).
- Finlay, Stephen (2019), "Defining Normativity", in *Dimensions of Normativity*, ed. D Plunkett, S. Shapiro, and K. Toh (Oxford University Press), 187-219.
- Finlay, Stephen and Snedegar, Justin (2014), "One Ought Too Many", *Philosophy and Phenomenological Research* 89: 102-124.
- Gibbard, Allan (2003), *Thinking How to Live* (Harvard University Press).
- Glanzberg, Michael (2007), "Context, Content, and Relativism", *Philosophical Studies* 136: 1-29.
- Glanzberg, Michael (2018), "Lexical Meaning, Concepts, and the Metasemantics of Predicates", in *The Science of Meaning*, ed. D. Ball and B. Rabern (Oxford University Press), 197-225.
- Glanzberg, Michael (2020), "Indirectness and Intentions in Metasemantics", in *The Architecture of Context and Context-Sensitivity*, ed. T. Ciecierski and P. Grabarczyk (Springer), 29-53.
- Grice, H. P. (1989), *Studies in the Way of Words* (Harvard University Press).
- Hare, R. M. (1952), *The Language of Morals* (Clarendon Press).
- Kaufmann, Magdalena (2012), *Interpreting Imperatives* (Springer).
- King, Jeffrey (2014), "The Metasemantics of Contextual Sensitivity", in *Metasemantics*, ed. A. Burgess and B. Sherman (Oxford University Press), 97-118.
- Klimczyk, Joanna (2020), "Compositional Semantics and Normative 'Ought'", *Axiomathes*. Advance access publication. doi: [10.1007/s10516-020-09500-z](https://doi.org/10.1007/s10516-020-09500-z)
- Köhler, Sebastian (2018), "Expressivism, Meaning, and All *That*", *Canadian Journal of Philosophy* 48: 337-356.
- Kramer, Matthew H. (2019), "Shakespeare, Moral Judgments, and Moral Realism", in *Routledge Companion to Shakespeare and Philosophy*, ed. C. Bourne and E. Bourne (Routledge), 235-245.
- Kratzer, Angelika (1981), "The Notional Category of Modality", in *Words, Worlds, and Contexts*, ed. H. J. Eikmeyer and H. Rieser (de Gruyter), 38-74.
- Kratzer, Angelika (1991), "Modality", in *Semantics*, ed. A. von Stechow & D. Wunderlich (Berlin: De Gruyter), 639-50.
- Lassiter, Daniel (2016), "Linguistic and Philosophical Considerations on Bayesian Semantics", in Charlow and Chrisman (eds.), 2016, 82-116.

- Lassiter, Daniel (2017), *Graded Modality* (Oxford University Press).
- Mandelkern, Matthew (2021), “Practical Moore Sentences”, *Noûs* 55: 39-61.
- Ninan, Dilip (2005), “Two Puzzles about Deontic Necessity”, *New Work on Modality, MIT Working Papers in Linguistics*, ed. J. Gajewski, V. Hacquard, B. Nickel, and S. Yalcin, 149–178.
- Raz, Joseph (1980), *The Concept of a Legal System*, 2nd edition (Clarendon Press).
- Ridge, Michael (2014), *Impassioned Belief* (Oxford University Press).
- Ruiz, Andrés Soria and Stojanovic, Isidora (2019), “On Linguistic Evidence for Expressivism”, *Royal Institute of Philosophy Supplement* 86: 155-180.
- Schroeder, Mark (2011), “Ought, Agents, and Actions”, *Philosophical Review* 120: 1-41.
- Sharadin, Nate and van Someren Greve, Rob (2021), “Is Deontic Evaluation Capable of Doing What it is For?”, *Journal of Ethics and Social Philosophy* 19: 203-229.
- Silk, Alex (2016), *Discourse Contextualism* (Oxford University Press).
- Simons, Mandy and Zollman, Kevin (2019), “Natural Conventions and Indirect Speech Acts”, *Philosophers’ Imprint* 19(9): 1-26.
- Stalnaker, Robert (1997), “Reference and Necessity”, in *A Companion to the Philosophy of Language*, ed. B. Hale and C. Wright (Blackwell), 534-554.
- Strandberg, Caj (2012), “A Dual Aspect Account of Moral Language”, *Philosophy and Phenomenological Research* 84: 87-122.
- Tiefensee, Christine (2016), “Inferentialist Metaethics, Bifurcations, and Ontological Commitment”, *Philosophical Studies* 173: 2437-2459.
- Tiefensee, Christine (2021), “Metasemantics for the Relaxed”, *Oxford Studies in Metaethics* 16: 108-133.
- Väyrynen, Pekka (2013), *The Lewd, the Rude and the Nasty: A Study of Thick Concepts in Ethics* (Oxford University Press).
- Viebahn, Emanuel and Vetter, Barbara (2016), “How Many Meanings for ‘May’? The Case for Modal Polysemy”, *Philosophers’ Imprint* 16 (10): 1-26.
- Wedgwood, Ralph (2001), “Conceptual Role Semantics for Moral Terms”, *Philosophical Review* 110: 1-30.
- Wedgwood, Ralph (2007), *The Nature of Normativity* (Oxford University Press).
- Wedgwood, Ralph (2017), *The Value of Rationality* (Oxford University Press).
- Woods, Jack (2014), “Expressivism and Moore’s Paradox”, *Philosophers’ Imprint*

14 (5): 1-12.

Worsnip, Alex (2019), “‘Ought’-Contextualism Beyond the Parochial”,
Philosophical Studies 176: 3099-3119.

Yalcin, Seth (2014), “Semantics and Metasemantics in the Context of Generative Grammar”, in *Metasemantics*, ed. A. Burgess and B. Sherman (Oxford University Press), 17-54.