9

METASEMANTIC QUANDARIES

Nate Charlow

Traditional interest theories hold that ethical statements are descriptive of the existing state of interests—that they simply give information about interests . . . It is this emphasis on description, on information, which leads to their incomplete relevance. Doubtless there is always some *element of description in ethical judgements, but this is by no means all. Their major use is not to indicate facts, but to* create an influence. *Instead of merely describing people's interests, they* change *or* intensify *them. They* recommend *an interest in an object, rather than state that the interest already exists.*

<div align="right">(Stevenson 1937)</div>

Conversation, then, is far more than a carrier of information. In talk we work out not only what to believe about things and events and people, but how to live. We work out how to feel about things in our lives, and in the lives of others.

<div align="right">(Gibbard 1990)</div>

1. Introduction

If an expression's semantic value or content is sometimes fixed in virtue of certain features of the context in which that expression is uttered, does it follow that that expression's semantic value, relative to the context of utterance, is *always* fixed in virtue of such features? If the semantic value of a sentence containing such an expression, relative to the context of utterance, can depend on the choice of a value for that expression, does it follow that the semantic value of such a sentence, relative to the context of utterance, *always* depends on such a choice? Although such questions are not often registered in the philosophy of language,

a great many philosophers of language—this paper will refer to them as "Referentialist" Metasemanticists—do seem to behave if the answer to both of these questions is a clear "yes."[1]

This paper proposes a different direction. It advocates a form of Expressivism as a strategy for resolving certain metasemantic puzzles about identifying the semantic value of a context-sensitive expression in context. On this version of Expressivism, while some utterances containing expressions in the target class do aim to proffer a proposition in the discourse in which they occur, such uses should be thought of as a kind of special interpretive case. Puzzles arising from the pressure to say what a putatively context-sensitive expression "refers to" in contexts that do not seem to specify a referent dissolve, once we appreciate that such attempts were ill-placed to begin with.

The version of Expressivism defended here will not deny that expressions in the target class *can* contribute a contextually determined semantic content to the computation of the semantic value of a larger syntactic constituent in which they occur. It instead says that expressions in the target class *need not* deliver a contextually determined semantic value as the input to semantic computation. Computation of a semantic content in context can avail itself of a variety of compositional mechanisms (e.g., λ-abstraction) in order to generate *nonpropositional* (more specifically, *prescription-type*) semantic contents, whose features are appropriate to realistic conversational/communicative aims of speakers.

Here is the plan. Section 2 reviews the basic metasemantic challenge raised by uses of context-sensitive expressions that do not appear to be specifically referential—that is, uses of such expressions in contexts in which they appear not to contribute a single entity of the relevant semantic type to semantic computation. Section 3 reviews a recent proposal for representing this phenomenon in the domain of gradable adjectives—the Expressivist account of MacFarlane (2016). It argues that there are *two* types of failures of specific referentiality that must be distinguished: those arising from semantic indeterminacy and those arising from broadly "nondescriptive" or "expressive" communicative aims (while also arguing that MacFarlane's account blurs these). While the former kind of failure can be modeled simply by allowing our theory of content-assignment to go indeterminate for certain kinds of context, the latter kind of failure cannot. Section 4 identifies some (at least prima facie) shortcomings with a metalinguistic treatment of the phenomenon of nondescriptive uses. Section 5 suggests an alternative model, on which nondescriptive uses are represented as semantically expressing a kind of *prescription*—a prescription that is semantically derivable, using familiar compositional machinery, from the standard semantic representation of,

---

1 Some linguists have been more circumspect (see, e.g., Dowty 1985, Jacobson 1999, and the literature on Direct Compositionality that grows out of this work). Theorists who have entertained the "no" answer to such questions (including Barker 2002; MacFarlane 2016) have tended to work with a (in my view) distorted perception of what the theoretical terrain will look like (and what the relevant modeling options are), once we've liberated ourselves of the demands of a Referentialist Metasemantics. Much more on this below.

for example, a gradable adjective in the positive form. Section 6 situates this proposal within the metasemantic framework of Gibbard (1990), while identifying one significant respect in which this proposal differs from Gibbard's.

2. Metasemantic Puzzles

According to our shared Kaplanian folklore, there are two ways a linguistic context can fix a semantic value for a context-sensitive expression. Some such expressions (e.g., the so-called "pure" indexicals, like "I") seem to come with lexically encoded rules (characters) that suffice to identify their semantic values in any context. Other such expressions (what King 2014 calls the *supplementives*) seem to be such that their characters do not suffice to fix their semantic values for any context whatever; in Kaplan-speak, such expressions require an associated "demonstration" (understood here to involve some type of indication of the speaker's referential intention that is supplementary to the utterance itself) in conjunction with the information provided by the utterance context and the expression's lexical character, in order to be assigned a semantic value in context. King argues that the class of natural language supplementives includes demonstratives, certain pronouns, modals, conditionals, gradable adjectives, and more besides (see also his 2017, 2018).

It is evident that the expressions that King understands to be supplementives exhibit *some form* of context-dependence (more precisely, that they *can receive* context-dependent semantic values). A speaker who says "Steph Curry is tall" can be interpreted as making a different claim depending on the relevant comparison class (e.g., American adults or NBA players). What the speaker says seems false if the speaker is comparing Steph Curry, who is just under 2 meter tall, to other NBA players, true if they are comparing him to American adults. A speaker who says "Bond might be in Zurich" can be interpreted as characterizing or describing a feature of their own knowledge, or what is known by a relevant group (what "we" know). What the speaker says seems false if the speaker is trying to characterize or describe what the group knows (and someone in the group knows Bond is not in Zurich). But it seems true if the speaker is trying to characterize what *they* know (and it is compatible with what they know that Bond is in Zurich).

None should dispute that supplementives exhibit context-dependence in this (quite thin) sense: for each supplementive, it is clear that *there are* contexts in which its semantic value is a function of that context. It does not, however, follow from this that supplementives exhibit context-dependence in a more theoretically robust sense. In particular, from the fact that there are contexts in which the semantic value of a supplementive is a function of that context, it does not follow that the semantic value of a supplementive is a function of context in *any* context.[2]

---

2 Throughout I bracket bound occurrences of supplementive expressions (e.g., pronouns).

Thank goodness for that! Suppose it were the case that, for *any* context *c*, the semantic value of a supplementive (tokened as part of a larger clause with a semantic value at *c*) was, in fact, a function of *c*; suppose, that is to say, that context always fixed the semantic value of a supplementive (when it occurred within a larger clause with a semantic value). Assigning a semantic value to the larger clause would *commit* the theorist to assigning a semantic value to the supplementive as a function of context. And this would in turn give rise to metasemantic questions about what features of the context the supplementive's semantic value was a function of.

Metasemantic questions like these have a tendency to become metasemantic quandaries. As King has stressed, supplementive expressions exhibit a strong degree of tolerance for **contextual nonspecificity**.[3] Even when no relevant facts about the context seem to settle which interpretation of the supplementive is "the" interpretation of the supplementive in that context, speakers still make liberal use of supplementive expressions (and addressees generally have no difficulty interpreting such utterances, assessing them for truth or falsity, affirming or denying the semantic content of those utterances, etc.). We might imagine, for example, that the speaker is[4] *indifferent* to (and therefore undecided about) whether, in saying "that is a beautiful car," they mean to be talking about the car type or token. An addressee who is, by stipulation, omniscient about the speaker's referential intentions will have no way of determining what "that" designates, but no difficulty in figuring out how to update on the speaker's utterance. The speaker appears to say something with content, even though the semantic value of the supplementive "that" is, almost by stipulation, not a function of the relevant context.

How should a theorist react to the phenomenon of nonspecificity? I can see three options:

1. One might think that there is pressure (from, e.g., compositionality) to hold that the semantic value of a supplementive at *c* (tokened as part of a larger clause with a semantic value at *c*) is *always* a function of *c* (if it is ever a function of *c*). Such pressures render metasemantic quandaries theoretically unavoidable.

2. One might also think that sentences with nonspecifically referential supplementives can receive *pragmatically workable interpretations* in context, but that such interpretations are not semantic values in the ordinary sense (e.g., because such interpretations arise as a kind of indirect speech act).

---

3 King's term is "underspecification" (see King 2017, 2018). "Nonspecificity" is clumsier, but it avoids the implication, which I will be denying here, that specificity of semantic value is the default/normal state of affairs for a supplementive in context.

4 Typically, I would think, speakers *are undecided* on whether they mean to be referring to the car type or car token. In most contexts, it just is unimportant, given the communicative aims of the speaker, for the speaker to be precise or decided in this fashion.

3. One might posit some sort of semantic distinction between two tokens of the same sentence type relative to a context c_1 in which the semantic value of the supplementive is specified and a context c_2 in which no semantic value for the supplementive is specified.

The next sections will examine these options, and identify some reasons for developing a theory in the general mold of (3). After that, I try to lay down a few semantic and pragmatic cornerstones for theorizing in this direction. The theory I will outline is in certain respects just a "generalization" of Gibbard's Expressivism to non-practical language. The driving idea is that Gibbard's theory of practical claims is a branch of a larger theory of "cognitive prescriptions", on which speakers express cognitive prescriptions – *ways* of thinking/representing on some at-issue matter, which they have *and* in some sense expect their interlocutors to share – by expressing properties of (contextually free) semantic parameters. (Practical claims are a special case, in which the contextually free semantic parameter is a "normative system" or "planning state".) This paper will ultimately register some (significant) differences with Gibbard's theory, while also taking issue with the, we might say, overly "literal" Gibbardianism of theorists like MacFarlane (2016). But the theory is, in its essence, an Expressivist theory – it presupposes the approach to thinking about thought and content that Gibbard pioneered.

3. Referential Metasemantics

It is evident that supplementives sometimes receive semantic values as a function of context (and that, when they do, their semantic values combine with the semantic values of sister nodes in the ordinary compositional fashion). In type-theoretic frameworks (like Heim and Kratzer 1998), the default mode of semantic composition is Functional Application.[5]

Functional Application (FA)
If α is a branching node whose daughters are β and γ, then, when defined, $[\![\alpha]\!]^c = [\![\beta]\!]^c([\![\gamma]\!]^c)$ or $[\![\alpha]\!]^c = [\![\gamma]\!]^c([\![\beta]\!]^c)$.

Suppose γ is a supplementive expression whose semantic value in c, $[\![\gamma]\!]^c$, is fixed as a function of c. If $[\![\beta]\!]^c$ is a function whose domain includes $[\![\gamma]\!]^c$, semantic composition proceeds in the usual fashion, and the semantic value of α at c, $[\![\alpha]\!]^c$, is simply $[\![\beta]\!]^c([\![\gamma]\!]^c)$.

This "familiar" state of affairs can come to have the air of obligation, if our only mode of semantic composition is FA: unless $[\![\gamma]\!]^c$ is defined (and of the right type to allow it to combine with $[\![\beta]\!]^c$), semantic composition (via FA) seems to break down. So, for any context c

---

5 If the reader is worried that the content of a complex constituent, relative to a context, is not compositional (see, e.g., Lewis 1980), they may read these claims as claims about the semantic value (or extension) of a complex constituent, relative to a context.

in which semantic composition appears not to break down for α—if, say, α is a sentence that appears to say something at c—there is pressure to hold that $[\![\gamma]\!]^c$ is a value of the requisite semantic type. And so there is compositional pressure to articulate what I'll term a **Referential Metasemantics** for supplementive expressions (by which I will mean an account of how, or in virtue of what, a supplementive expression receives *a semantic value of the ordinary type at any context* in which computation of semantic value for a syntactic constituent that contains it succeeds).

3.1 MacFarlane on Nonspecificity
Though Referential Metasemantics is the theoretical default for work on the metasemantics of supplementives, nonspecific uses of supplementive expressions present an immediate, and serious, challenge to the view.

Let's start with the following remark by MacFarlane:

> We have plenty of . . . flexible expressions, whose extensions are to a great extent up to the speaker to determine. The most obvious examples are bare demonstratives like "this" and "that." In principle, I can use "that" to refer to any object. But with this freedom comes great responsibility. I must provide my hearers with enough cues to enable them to associate my use of "this" with the same object I do, or communication will fail . . . [I]n every case, *we're obliged to do whatever is required to get our hearers to associate the same object with the demonstrative that we do.* If we fail to do this, it will be sheer luck if they understand us.
>
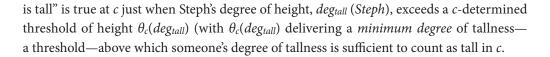> (MacFarlane 2016, 260–61, emphasis mine)

MacFarlane says, I think correctly, that successful communication in a context in which a speaker expresses a semantic value with a supplementive requires that speaker and addressee coordinate on that semantic value. Nonspecific uses of supplementive expressions present cases in which it appears *no* relevant facts about the context settle the expression's interpretation. In such contexts, a speaker *does not* (and, in many cases—e.g., in the case of specifying an exact comparison class for a gradable adjective—*cannot*) provide cues that allow her addressee to associate her use of a supplementive with the object, if any, she associates with the supplementive. If the supplementive contributed a semantic value of the ordinary type to the proposition expressed by the utterance, there would be no way for speaker and addressee to coordinate on the proposition expressed by the utterance. Since successful communication seems to demand such coordination, the prediction is that communication must *fail* in any such context.

We have seen that this is a false prediction. To avoid it, MacFarlane proposes denying, in at least some such cases, that a speaker expresses a proposition with her utterance that is a function of a contextually determined semantic value for the supplementive. MacFarlane's specific target is the contextually determined degree threshold invoked in the standard degree semantics for gradable adjectives (see a.o. Kennedy 2007). On that semantics, "Steph

is tall" is true at c just when Steph's degree of height, deg_{tall} (*Steph*), exceeds a c-determined threshold of height $\theta_c(deg_{tall})$ (with $\theta_c(deg_{tall})$ delivering a *minimum degree* of tallness— a threshold—above which someone's degree of tallness is sufficient to count as tall in c.

$$\llbracket Steph\ is\ tall \rrbracket^c = 1\ iff\ deg_{tall}\ (Steph) > \theta_c(deg_{tall})$$

MacFarlane takes this sort of semantics to be ruled out by the sorts of considerations described in the prior paragraph; coordinating on a precise degree threshold is not something that people are even ordinarily able to do in conversation.

Instead, MacFarlane claims, sentences like "Steph is tall" are semantically evaluated with respect to Gibbardian *hyperplans*, which are objects that, for any possible situation s, specify exactly which actions are forbidden/permitted in s (see esp. Gibbard 1990, 2003). A hyperplan is, inter alia, a plan for where to draw the line for the degree of height required to count as tall in c, for any context of utterance c. For MacFarlane, the semantic content of a sentence like "Steph is tall" at a context of utterance c is the set of hyperplans h such that Steph's height exceeds the degree of height required to count as tall in c, according to h—$\theta_h(deg_{tall})$:

$$\llbracket Steph\ is\ tall \rrbracket^c = \lambda h.deg_{tall}\ (Steph) > \theta_h(deg_{tall})$$

Equivalently, its content is a *property of planning states*—specifically, the property of planning to count anyone of at least Steph's height as tall. Contra Kennedy, it is not a contextually determined proposition. Proffering such a property in discourse amounts, *not* to proffering a proposition for addition to the Common Ground, instead to proffering to one's audience a (*practical*, rather than doxastic or epistemic) constraint on who *to count as* tall—namely, anyone of at least Steph's height.

3.2 Semantic Indeterminacy

Though I am broadly sympathetic to the conclusion—and will argue for a version of it later on—I believe the argument rests implicitly on a dubious contrast between bare demonstratives and gradable adjectives. MacFarlane draws the contrast as follows:

> While in using a bare demonstrative like "this" one must have a definite object in mind, and successful uptake requires recognizing what object that is, there are no analogous requirements for the use of "large." The speaker need not have in mind a particular delineation (even a "fuzzy" one), and the hearer need not associate the speaker's use with a particular delineation. What we get instead are constraints on delineations. In saying that apple C is large, I rule out certain ways of drawing a line between large and non-large apples, while leaving others open.
>
> (2016, 265)

Consider an utterance of "that's a beautiful car" in a context in which the speaker lacks a specific referential intention. Let us suppose further that it is common ground that a token of the car type is beautiful just when that car type is beautiful. Thinking this particular car beautiful is informationally equivalent, in this context, to thinking the car type beautiful. Communication seems to succeed here, even though the speaker does not provide the addressee the relevant cue about her referential intention (indeed, the speaker appears to lack any such intention to provide a cue *about*) (see again King 2017, 2018).

Now the fact that communication using a demonstrative can succeed, even absent specific referential intention, does *not* warrant the conclusion that the semantic value in context of a demonstrative-containing sentence like "that's a beautiful car" should be taken to be *anything other than* a proposition. Certainly, it does not warrant the (stronger and false) conclusion that the content of this sentence *is* a property of planning states (i.e., the property of planning to use "that" to refer to a beautiful car type or car token, but being indifferent between these[6]). One therefore wonders: what is supposed to differentiate the sort of semantic nonspecificity that warrants an alternative assignment of content (e.g., a property of hyperplans) from the sort of semantic nonspecificity that does not?[7]

Nor does this fact warrant even the weaker conclusion that the semantic value of "that's a beautiful car" is not a function of a contextually determined semantic value for the supplementive. "The" semantic value of this sentence could be *either* the proposition that cars of this type are beautiful or that this particular car is beautiful, in the sense that both propositions are semantically "eligible" in this context—there are contextually admissible resolutions of the demonstrative "that" that yield each of these propositions as the output of semantic computation. Either resolution of the demonstrative would yield a proposition with the right pragmatic profile (since, by assumption, updating on one yields the same informational change in this context as updating on the other).

We *could* ask which of these propositions is "the" semantic value of "that's a beautiful car" in context. But why? The speaker's utterance could express *either* of two propositions (and an interpreter may update on either of these two propositions, apparently without loss of information or understanding). There is no obvious cost to saying that the speaker expresses *both* semantic values in uttering this sentence in this way (compare King 2014, 106)—or, alternatively, that there is no determinate fact of the matter about which of these semantic values is the semantic value expressed by the utterance. Let us assume that, to realize their communicative aims without misunderstanding, speaker and addressee must

---

6   This seems to misconstrue the point of such an utterance. A speaker who says "that's a beautiful car" in the context we are imagining intends to express an aesthetic judgment about a car, not to express a constraint on one's plans for using the demonstrative "that."

7   MacFarlane writes that the standard picture of content as the proposition expressed by a sentence in context "assumes that speaker and hearer have shared knowledge of what it takes for the sentence to be true in the present context. When that assumption breaks down, truth-conditions lose their explanatory relevance" (MacFarlane 2016, 265).

be able to coordinate on a way of updating the context/their information. But this does not generally require that speaker and address be able to coordinate on a unique semantic value for the speaker's utterance. Semantic specificity beyond what is required for communicative aims is otiose (for the agents of a conversation, as well as for theorists trying to model their conversation). Why, then, would we posit it?

With a little effort, this strategy can be extended to MacFarlane's target, gradable adjectives. Someone who says "Steph Curry is tall" could, from the point of view of a referential metasemantics, express any of (continuum) many propositions—one for each way of drawing the threshold (compare Braun and Sider 2007; King 2014, 112). Do these propositions "carry" the same information in context? In one sense, obviously, no: the proposition that Steph's height exceeds θ_1 is, of course, distinct from the proposition that Steph's height exceeds θ_2; supposing θ_2 exceeds θ_1, the proposition that Steph's height exceeds θ_2 asymmetrically entails the proposition that his height exceeds θ_1. Still, if both θ_1 and θ_2 are candidate thresholds—both are contextually eligible ways of drawing the line between tall and non-tall—Steph's height must be assumed (by someone who says "Steph is tall") to exceed *both* θ_1 and θ_2: if Steph's height is not assumed to exceed θ_2, but θ_2 is regarded as an eligible threshold in the context, the utterance is marked.

> A: Steph is 1.9m. Do you have to be 2m to be tall?
> B: I'm not sure.
> A: Is Steph tall?
> B: ?? Yes, he is.

For any "contextually eligible" θ_1 and θ_2, someone who says "Steph is tall" will be taken to be committed to both the proposition that Steph's height exceeds θ_1 and the proposition that it exceeds θ_2. There is, again, no clear cost to saying that the speaker expresses *both* semantic values in saying "Steph Curry is tall"—or, alternatively, that there is no fact of the matter about which of the various compositionally eligible semantic values is *the* semantic value expressed by her utterance—since the speaker is taken as committed to any such value. Why, again, would a theorist demand semantic specificity, when the communicative aims of speakers and addressees *do not*?

Metasemantic quandaries dissolved? Expressivism circumvented? And this simply by introducing a plausible bit of indeterminacy into the semantico-pragmatic relation (*expression*) that relates speakers to the semantic values of their utterances?

3.3 Determinacy and Description

No, or so I will say. Semantic indeterminacy of the sort just described is a useful rubric for thinking about one kind of failure of specific referentiality (in the domain of gradable adjectives). But semantic indeterminacy does *not* help to account for another way in which speakers can use context-sensitive items nonspecifically.

A now common observation in the literature on gradable adjectives notes that they have two canonical "modes of use" (Barker 2002, 1ff.), one broadly *descriptive* in character (e.g., a speaker uses an utterance of a gradable adjective in the positive form to provide information about someone's height), another broadly *nondescriptive* in character. Barker contrasts these modes of use as follows:[8]

> Normally, (1) will be used in order to add to the common ground new information concerning Feynman's height:
>
> (1) Feynman is tall.
>
> But (1) has another mode of use. Imagine that we are at a party. Perhaps Feynman stands before us a short distance away, drinking punch and thinking about dancing; in any case, the exact degree to which Feynman is tall is common knowledge. You ask me what counts as tall in my country. "Well," I say, "around here, . . ." and I continue by uttering (1). This is not a descriptive use in the usual sense. I have not provided any new information about the world, or at least no new information about Feynman's height . . . All I have done is given you guidance concerning what the prevailing relevant standard for tallness happens to be in our community; in particular, that standard must be no greater than Feynman's maximal degree of height. (Barker 2002, 1–2)

It is *this*, apparently nondescriptive, use that MacFarlane proposes to model with hyperplan-type content (see, e.g., MacFarlane 2016, 256). That is to say, MacFarlane proposes to represent the speech act that Barker glosses as giving "guidance concerning . . . the prevailing relevant standard for tallness" as a speaker's proffering a (practical, not doxastic or epistemic) constraint on who *to count as* tall.

The question of whether an utterance expresses a specific semantic value in c is, however, distinct from the question of whether it functions to describe in c. That is to say, semantic determinacy is orthogonal to descriptiveness.

    a. An utterance may exhibit semantic determinacy, while expressing an assertion that updates the context with a determinate proposition. (The "Familiar" Case)

    b. No one proposition is the semantic value of the utterance, though the utterance is understood as expressing an assertion that updates the Common Ground with a determinate proposition (since the semantically eligible propositions are informationally equivalent in context). (most of King's Cases)

    c. An utterance may exhibit semantic determinacy, while the utterance's force is to proffer some kind of cognitive property or constraint.

---

8 Barker glosses the nondescriptive use as "metalinguistic." On the limits of this gloss, see section 4 below.

d. An utterance may exhibit semantic indeterminacy, while the utterance's force is to proffer some kind of cognitive property or constraint.

As a specific illustration, a speaker can use epistemic "might" in any of the following ways:[9]

a. To describe what is possible at c given a salient body of information at c. (This is the "ordinary" or "familiar" case.)
    [Context: A is gathering information about B's information, and this is common ground between A and B.]
    A: Where might Bond be [given your information]?
    B: He might be in Zurich.

b. To describe what is possible at c given the content of some or other body of information at c, while remaining undecided about which body of information. (King-type nonspecificity)
    [Context: A is gathering information about C and D's information, and this is common ground between A and B.]
    A: Where might Bond be?
    B: He might be in Zurich [given what C or D believe, it doesn't matter].

c. To determinately constrain someone's information so that, if the addressee accepts her utterance, she will regard a determinate proposition *as possible* (compare Moss 2013, 2015; Swanson 2006, 2016).[10]
    [Context: A and B are disagreeing about where Bond might be.]
    A: Bond has to be in London!
    B: No, Bond might be in Zurich.

d. To indeterminately constrain someone's information (in a context where satisfying one semantically eligible constraint is informationally equivalent to satisfying any semantically eligible constraint).
    [Context: A/B could be referring to either the car type or token.]
    A: That might be a Ferrari.
    B: No, but it might be a Maserati.

I think it is apparent that stating a theoretically adequate metasemantics for semantic indeterminacy—one liberated from the assumption that the relation that holds between

---

9 N.B. I don't claim all of these functions are attested for all supplementive expressions (although for some such expressions, like epistemic "might," it does appear that all are attested).

10 It is standardly held that such an update cannot be generally modeled as updating on a proposition (Russell and Hawthorne 2016; Veltman 1996; Yalcin 2011). I here treat disagreement about *whether to treat the proposition that Bond is in Zurich as possible* as indicative of nondescriptive disagreement, that is, disagreement that is not well-represented as disagreement about features of the actual world. This is not to say that other models of this sort of disagreement are ruled out (see MacFarlane 2011, 2014 for one).

a speaker and the content of what she says (i.e., *expression*) must be a *function*—does not free the theorist from the need to explain how certain context-sensitive expressions receive nondescriptive interpretations in context (or from the need to explain what such interpretations consist in). These are just different tasks. The problem of semantic indeterminacy is resolved—to a first pass, anyway—by getting comfortable with the notion that certain semantic facts are indeterminate (or that the semantic relation of expression is one-many).[11] It is not to be resolved with planning content (unless the theorist is willing to hold—contrary to apparent fact—that all utterances exhibiting semantic indeterminacy express nondescriptive planning content).

Nor is the problem of modeling nondescriptive, or constraint-type, interpretations resolved by making semantic facts indeterminate. It is, I will ultimately argue, resolved by recognizing a distinctive kind of *prescriptive content*—similar to, but also distinct in important ways, from the sort of planning content envisioned by MacFarlane.

4. The Metalinguistic Strategy

Nondescriptive uses of context-sensitive expressions raise metasemantic issues similar to those raised by nonspecific uses. Specifically, both are prima facie counterexamples to a Referentialist metasemantics for such expressions: if the expression contributed a semantic value of the ordinary type to semantic computation for the sentence, the output of semantic computation for the sentence would be a proposition involving that semantic value, and this would apparently fix a descriptive use.

This section will describe one strategy for resisting this argument—the metalinguistic strategy suggested in Barker (2002) and refined in Plunkett and Sundell (2013). On this strategy, sentences in the target class generally express propositions, but this need not fix a descriptive use for the sentence. I will argue that nondescriptive uses are not generally well understood as advancing proposals bearing on how to use language.

In contexts that do not provide a unique threshold for counting as tall, Barker (2002) represents the content of a claim like "Feynman is tall" as a constraint on the tallness-thresholds that are compatible with the context of utterance; the force of uttering such sentences is to

---

11 I say "to a first pass" because a general theory of how indeterminate semantic values determine illocutionary force in context will also need to be stated. This will be a difficult project, since there are a range of relations a set of candidate semantic values may bear to the semantic value on which agents update in a context (e.g., sometimes agents update on a disjunction of candidate semantic values, other times they update on the strongest semantic value; for examples of this variability, see King 2017). The project is then to say how interpreters derive the content on which they update in such contexts. (Note, by the way, that this is distinct from the problem of saying which proposition interpreters update on in such contexts: the information interpreters glean from the utterance is generally clear from the context, but it does not appear to be a function of which propositions are expressed by the utterance.) I will table this problem here.

eliminate those thresholds that are incompatible with the utterance from eligibility in the context of utterance.

Barker (2002) offers a substantive characterization of the *speech act* associated with proposing or expressing this sort of update as metalinguistic in nature or aim: "My purpose in uttering ['Feynman is tall'] . . . would be nothing more than to communicate something about how to use a certain word appropriately" (2). At first blush, however, Barker's metalinguistic gloss on this update looks undermotivated: proposing to exclude a candidate threshold for tallness from eligibility is not (obviously, anyway) the same thing as proposing to tell one's addressee how to use "tall." What reason is there to *identify* these speech acts? More generally, why think that modeling the nondescriptive interpretation of, for example, a sentence containing a gradable adjective *requires* a metalinguistic analysis?

In the Stalnakerian framework Barker takes on (Stalnaker 1978, 1984), proposals to update the context that are modeled as *eliminative* (e.g., as the intersection of the set of thresholds compatible with the utterance with the set of thresholds antecedently eligible in the context of utterance) are linked to a particular *functional role*. The function of updating a context with a set of possible worlds is straightforward to characterize: the antecedently eligible worlds—those in the pre-update context set—represent ways the world could be, given the present state of the conversation. The functional role of refining this set is to eliminate certain worlds as *candidates for actuality* in the conversation.

This philosophical context in view, it seems fairly clear that this functional story does not extend to the elimination of thresholds from contextual eligibility. An eligible threshold is not a *candidate for actuality*—there is, I will take it, no "actual" threshold that we are trying to figure out in conversation when we use gradable adjectives (MacFarlane 2016 agrees and offers considerations in support of this claim).
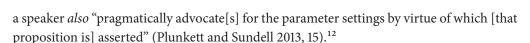
The metalinguistic strategy posits a *novel functional role* for the act of eliminating a threshold from contextual eligibility. The idea is that a threshold θ for a gradable adjective α is *c-eligible* when it constrains appropriate use of α in c: an utterance containing α must be compatible with some c-eligible threshold if the utterance is appropriate in c. (Compare: in the standard Stalnakerian framework, the possible worlds proposition one asserts must be compatible with the context set for the assertion to be appropriate.) To eliminate a threshold from contextual eligibility for α is therefore to make a metalinguistic proposal: to constrain how one's interlocutors use α in the conversation.

If the metalinguistic strategy is the right one for modeling nondescriptive interpretations of, for example, sentences containing gradable adjectives, these sentences might semantically express propositions—as Plunkett and Sundell (2013) are happy to take them to do—even while their functional role is not to refine our mutual representation of ways the world could be. Though Plunkett and Sundell (2013) do not use this terminology, on such an account, nondescriptive interpretations will arise as *indirect speech acts* (see, e.g., Asher and Lascarides 2001; Searle 1975): by asserting the (typically true) proposition that is determined by *their* preferred way of resolving, for example, the degree threshold for a gradable adjective,

a speaker *also* "pragmatically advocate[s] for the parameter settings by virtue of which [that proposition is] asserted" (Plunkett and Sundell 2013, 15).[12]

One difficulty is that, in cases where a contextual "parameter setting" is semantically indeterminate, this sort of account will struggle to say which proposition is asserted. Since the account takes the form of an indirect speech act account, it will struggle to say precisely which indirect speech act is performed by the speaker in the context of utterance. This isn't to say that this challenge couldn't be somehow met. It is just to note that it is not met by the sort of account that is described in Plunkett and Sundell (2013).

Another, more empirical difficulty—this one arising from the attempt to place nondescriptive interpretations under the rubric of indirect speech acts—is that, with bona fide indirect speech acts, an utterance's literal semantic content remains accessible to downstream "relational speech acts"—speech acts that are, in an intuitive sense, *anaphoric* to other speech acts in the discourse. For example:

    A: Can you pass the salt?
    B: Yes, I can!/No, I can't!

B's reply makes no sense if interpreted as relating to A's request.

    A: Please pass the salt.
    B: #Yes, I can!/#No, I can't!

B's reply in the first dialogue is licensed by the fact that A *semantically expresses a question about B's abilities*, albeit in service of a further communicative aim: requesting that B pass the salt (for discussions of this phenomenon, see Asher and Lascarides 2001; Charlow 2011).

By contrast, what the metalinguistic analysis treats as an utterance's semantic content in cases of nondescriptive interpretations does not seem to be accessible to the expected array of relational speech acts. It is surprisingly difficult to target the claimed propositional content of a nondescriptive interpretation with, for example, relational affirmation or denial.

    [Context: A and B agree that Feynman is 6 feet tall, but disagree about whether that makes Feynman tall.]
    A: Feynman is tall.
    B: ?? No, he is below the threshold.

---

12 To be fair, Plunkett and Sundell (2013) do their best to prescind from semantic debates: whatever one's semantic commitments, they argue, one will require a treatment of metalinguistic negotiation-type uses (see esp. their Section 6.1). My target here is the particular model of metalinguistic negotiation-type uses they use to illustrate their account.

[Context: A and B agree that A's information is compatible with Bond being in Zurich and B's information is not]

   A: Bond might be in Zurich.

   B: ?? No, the information rules out Bond being in Zurich.

The infelicity in both cases seems due to misunderstanding: when A says "Feynman is tall," A does not (ordinarily) mean to be interpreted as making a comparison between Feynman's height and "the" relevant threshold. When A says "Bond might be in Zurich," A does not (ordinarily) mean to be interpreted as making a claim about the properties of the relevant information.[13] (Intuitively, in the first case, A and B are not disagreeing about the truth of any proposition of the form *Feynman's height exceeds θ*: their disagreement appears to consist in the fact that A *regards* or *considers* Feynman's height as sufficient for being tall, while B does not. Intuitively, there is a disagreement in attitude, unaccompanied by any evident disagreement in fact.)

Until we see an explanation of why speech acts that are anaphoric to the claimed propositional content of A's utterances do not appear to be licensed in such cases, it will be unclear whether the alleged propositional content of "Feynman is tall" or "Bond might be in Zurich" has *any explanatory role to play* in accounting for nondescriptive interpretations of these sentences. If we can account for such interpretations, without appeal to such alleged contents, this is at least some reason to think that we should.
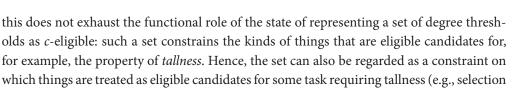
A final observation about the metalinguistic account: the action of eliminating, for example, a degree threshold from eligibility *need not be* understood metalinguistically (i.e., as targeting or aiming to constrain the *linguistic behavior* of the audience).[14] I would suggest that we do better to understand this action as an attempt to constrain the downstream *sortal attitudes* of the audience—as a kind of prescription (bearing on who to regard as tall, and who to regard as not-tall). Such a prescription no doubt bears indirectly on the use of words: to regard a set of degree thresholds as *c*-eligible might commit one to regarding *utterances* whose meanings are incompatible with at least one *c*-eligible threshold as inappropriate. But

---

13 Unsurprisingly, these dialogues sound a lot better when the context settles the referent of quasi-technical phrases like "the relevant threshold" and "the relevant information." If it is clear that *A is comparing Feynman's height to an explicit standard* (e.g., a line drawn on a wall), B can certainly reply to A's utterance "Feynman is tall" by saying "that's wrong, Feynman is below that line." If it is clear that B is seeking information from A about *where A thinks Bond might be*, B can certainly reply to A's utterance "Bond might be in Zurich" by saying "that is a damned falsehood, I know very well that your information rules this out."

14 This observation is distinct from a claim that Plunkett and Sundell (2013) are careful to rebut, namely that the action of eliminating, for example, a degree threshold from contextual eligibility *cannot* be understood metalinguistically. On this objection, metalinguistic moves in a conversation tend to be pointless—akin to merely verbal disagreement; we misconstrue the aims of speakers if we understand them as participating in such moves. But, as Plunkett and Sundell (2013) note, how we use words matters: whether or not Steph is regarded as "tall" will license (or not) a host of downstream actions (e.g., will we trade for Steph?). Metalinguistic judgments typically have nonmetalinguistic motivations and effects.

this does not exhaust the functional role of the state of representing a set of degree thresholds as *c*-eligible: such a set constrains the kinds of things that are eligible candidates for, for example, the property of *tallness*. Hence, the set can also be regarded as a constraint on which things are treated as eligible candidates for some task requiring tallness (e.g., selection for a game of pickup basketball, if that is the *c*-relevant task).[15]

I am thus inclined to think that the constraints on linguistic behavior that arise from representing a set of degree thresholds as *c*-eligible are better understood as a *consequence* of a more fundamental feature of this sort of representational state: representing a set of degree thresholds as eligible constrains who one *regards* or *considers as tall* in *c*, which subsequently constrains who one is able to appropriately *call tall* in *c*. On this model, the metalinguistic guidance that is provided by the elimination of a degree threshold from eligibility is not intrinsic to this kind of update; it is rather a kind of natural effect of adopting (in the sense of coming to satisfy) a constraint on who to regard as tall.

5. Prescription-Type Meanings

I will instead suggest that we recognize a distinction in *semantic type* between descriptive and nondescriptive interpretations. Descriptive interpretations have the usual propositional interpretation (with provisions for semantic indeterminacy). The nondescriptive interpretations in which I am interested here are *prescriptive* in nature (also with provisions for semantic indeterminacy). Drawing on the proposal for the semantics of imperatives—a dedicated clause-type for expressing prescriptions in natural language—developed in my earlier work (see, e.g., 2011, 2014, 2018), I will propose that these interpretations are semantically distinct (but semantically related). More concretely, when an expression of semantic type T is used nonspecifically (in the simple sense that it lacks a semantic value in context) its argument place can be optionally bound (by, e.g., λ-abstraction) to yield a characteristic function of objects of type T—that is to say, a *property* of objects of type T. Properties like this are well-suited to account for nondescriptive interpretations, I will here argue.

5.1 Prescriptions and Propositions
Imperatives encode/express constraints on states that have an *action-guiding* or *motivating* functional role (e.g., plans or preferences); in particular, they tell someone who accepts or updates on the imperative what their plans or preferences must be like (see Harris 2017a;

---

15 As noted above, Plunkett and Sundell (2013) utilize the fact that predicational questions (e.g., do we predicate "spicy" of this chili?) are linked to nonlinguistic questions (do we add more spice to the chili?) to argue that predicational questions are not merely verbal in import. While correct, this does not establish that the illocutionary point of saying "this chili is spicy" is to answer a predicational question (do we predicate "spicy" of this chili?). My eventual suggestion here will be that we do better to think of the resolution of such predicational questions as a natural effect of the resolution of nonlinguistic *normative questions* (e.g., *should we regard this chili as spicy*?).

Portner 2004, 2007, 2018; Roberts 2015, 2018; Starr 2020). Theorists differ about whether or not imperatives do this as a matter of their semantics: for Portner and Roberts, they do not, for Charlow, Starr, and Harris, they do. Here I will be assuming that the latter position is broadly correct.[16]

In my own account, an imperative like "Confess!" is *compositionally related* to a corresponding modal claim "you must confess" (it will take me a page or so to explain how). Consider the following (schematic, extensional) representation of a modal logical form:[17]

$$Modal_{f,g}(Restrictor)(Scope)$$

This representation is what we find in Kratzer (1977, 1981, 1991): modals are generalized quantifiers expressing a quantificational relation between (1) a domain of quantification jointly characterized by the Modal Base f, the Ordering Source g, and a (explicitly or implicitly provided) Restrictor, and (2) a set of possibilities characterized by the Scope.

Like other context-sensitive expressions, prioritizing (e.g., deontic) modals admit of descriptive and nondescriptive readings: "you must confess" can be used to describe what is required at c, given some or other body of norms or priorities made salient at c, but it can also be used to *tell someone to confess*. In such uses, the modal takes on a meaning that seems to be prescriptive, exhibiting many of the same features as the meaning of the corresponding imperative (e.g., infelicity when joined with denials that the relevant obligation will be discharged) (Ninan 2005; Portner 2007).

#You must go to confession, but you're not going to.
#Go to confession! You're not going to go to confession.

About the prescriptive meaning, Portner writes:

Since it seems that *must* has an obligation-imposing function, in addition to a traditional truth-conditional semantics, as part of its conventional meaning, the next question is what the nature of this obligation-imposing reading is. Ninan (2005) proposes to model it in terms of the notion of To-Do List (Portner 2004). Thus, he explains ["Confess!"] as follows: uttering the sentence places the property of going to confession on [the addressee's] To-Do List. But one cannot place a requirement on someone's To-Do List while at the same time asserting that it will not be met, and therefore the sentence is anomalous. What's important here is that the

---

16 All that will turn on this is whether the sort of contents I describe below for constraint-type interpretations are assigned in the semantics or by some "Dynamic Pragmatic" theory equipped with a mechanism like abstraction. On the Dynamic Pragmatic program, see especially Portner (2018) and Roberts (2018). For arguments against, see Charlow (2018).

17 The notion of "logical form" invoked here is syntactically neutral: I do *not* assume that representing a variable in logical form means representing a variable in morphosyntax.

ordinary, truth-conditional semantics for the modal does not play a role in explaining the patterns [above]. Rather, the independent imperative-like meaning does the job. (2007, 365–66)

Portner here assumes that any obligation-imposing (prescriptive) meaning carried by the modal would be "*independent*" of its "truth-conditional" (i.e., quantificational) semantics. Here, however, are two empirically viable possibilities for modeling the prescriptive meaning carried by prescriptive "must"; the truth-conditional dimension of the modal's meaning plays an essential role in both.

- **Performative**: Prescriptive interpretations consist in proposals to adjust a salient body of norms (or more generally a salient state with an action-guiding functional role) so that the modal proposition semantically expressed by the sentence is true relative to the adjusted body of norms.[18]

Like the metalinguistic account, a performative account makes central explanatory use of the proposition allegedly expressed by a sentence like "Confess!"; I will set such accounts to the side here (for arguments against performative accounts, see Charlow 2018).

- **Modally Derived**: Prescriptive interpretations of imperatives consist in proposals to adjust a salient body of norms (or more generally a salient state with an action-guiding functional role) so that it comes to satisfy a modally characterized property (see Charlow 2011, 2014, 2018).

The suggestion here is to represent the semantic content of both "Confess!" and "You must confess!" (on its prescriptive interpretation) at a context c with the same modally characterized property, namely:

$$\lambda g[must_{f_c,g}(addr_c \; goes \; to \; confession)]$$

This is the property an ordering source g has, iff all the g-best possibilities compatible with the c-relevant information f_c are possibilities where $addr_c$ goes to confession. Less technically, it is the property g has when g induces a *ranking* (more precisely, an ordering) on the possibilities compatible with f_c, according to which possibilities where the addressee of c goes to confession are *highest-ranked*. Less technically still, it is the property g has when, according to g, it is preferred/planned/required/ . . . that the addressee of c goes to confession.

---

18 More precisely, one proposes to alter a body of norms that characterizes a domain of quantification for the modal "must" so that all worlds compatible with the domain of quantification are worlds in which the prejacent proposition (e.g., that the addressee goes to confession) holds. For accounts of imperatives in this vein, see especially Lewis (1979) and Kaufmann (2012).

What links property-type contents of this sort to illocutionary forces or discourse moves? Such contents are properties that can, in a sense, be *instantiated* by motivating psychological states (e.g., an agent's plans).[19] It is natural to say that the force of an imperative is to propose a selected or salient motivating state—typically that of the addressee—come to instantiate this property.[20] The following Force Assignment principle gives the rough idea.

Force Assignment

The force of expressing a property of an ordering source at *c* is to propose that the action-guiding/ motivational state(s) of one's addressee(s) at *c* satisfy this property.

"Satisfaction," in the relevant sense, amounts to *representability*: an agent's action-guiding or motivating state satisfies $\lambda g[must_{f_c,g}(R)(S)]$ just when that state is representable with some g such that $must_{f_c,g}(R)(S)$ is true. So, for example, the force of "Confess!" at *c* is to propose that the motivational state of one's addressee come to be representable with priorities that require going to confession. Imperatives conventionally express directives, in the following sense: imperatives are conventionally associated with attempts to get their addressee(s) to be motivated or adopt plans with specific modally characterized properties.

5.2 Extending the Account

The strategy described in the last section is easily extended beyond imperatives. In general, when a speaker means to *express a property of some parameter* (for the sake of *adoption* by some agent in the context), rather than to describe a feature of this parameter, we will say:[21]

- The **semantic value** expressed by the speaker is a λ-abstract over a free (contextually unbound) occurrence of that parameter.[22]

---

19 Strictly speaking, I do not want to assume that there is any kind of tight/conventional link between expressing a property of *an ordering source* and the imperatival speech act of direction. In contexts in which an ordering source represents a body of *expectations* (as opposed to a body of plans or intentions), expressing a property of an ordering source means proposing that the expectations of one's addressee satisfy this property. Ordering sources come in different "flavors" depending on what information they are used to represent, and the force of expressing a property of an ordering source will depend on what type of information the ordering source is used to represent. (This sketchy suggestion is meant to build on Kratzerian (1981) orthodoxy about "modal flavor" (i.e., polysemy). That said I am doubtful that we in fact have a good understanding of what modal flavor is/how it is determined relative to a context of utterance.)

20 For a more complete statement of the account, see Charlow (2018).

21 For a related view of semantic composition and its metasemantic implications, see discussion (and references) in Harris (2017b).

22 There are different compositional routes to this semantic value. One familiar possibility is to treat the parameter as a free variable and apply the Predicate Abstraction rule of Heim and Kratzer (1998, 114). A less familiar, but also attractive, possibility eliminates variables (and variable assignments) from the metalanguage, instead treating semantically underspecified expressions as semantic placeholders, whose values are resolved "post-semantically" (see especially the "Variable-Free" system of Jacobson 1999).

- The **illocutionary force** of expressing such a meaning is proffering this property for adoption by the addressee in the context.

Proffering a property for adoption by another agent is an *inherently prescriptive* act: it amounts to offering a kind of *cognitive advice* (adapting a phrase of Swanson 2006, 39). It is therefore important to distinguish this general notion of prescriptive force from the more particular notion of prescriptive force that is relevant to the analysis of imperatives. Imperatival prescriptive force is *practical* or *directive* in nature: it bears on an agent's motivating or action-guiding psychological states; it tells the agent how she should plan (and therefore how she should act). Prescriptive force need not be practical or directive, in this more particular sense: when a speaker proffers a property of a nonmotivating (e.g., doxastic) state for adoption, she is not telling the addressee what to do; she is telling the addressee what, say, her beliefs must be like.

MacFarlane, recall, treats gradable adjectives as encoding practical information: information that constrains an agent's plans, rather than some other (e.g., nonmotivating) state of mind. This is *optional* on the framework I am proposing: while MacFarlane and I agree that gradable adjectives (in cases of nondescriptive uses) proffer prescriptions, I have here distinguished two senses of prescription: practical and more broadly cognitive prescriptions. The question of whether nondescriptive interpretations of gradable adjectives are best represented with practical or cognitive prescriptions remains open.

Consider again "Steph is tall" as uttered a context in which the sentence receives a nondescriptive interpretation—that is, roughly, a context in which the speaker is encouraging her addressee to *think of Steph as tall* (while not attempting to offer any descriptive information about Steph's height). Is the speaker well-understood as suggesting that the addressee *adopt a plan for* thinking Steph tall?
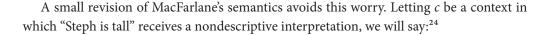
I do not think so. For one thing, thinking Steph tall—in the stative, rather than eventive, sense—isn't under the addressee's voluntary control, and typically isn't the sort of thing that you can plan to do. (This is why an imperatival expression of this kind of planning content, like "regard Steph as tall!," is generally heard as marked.[23]) A theory should not blur the distinction between the type of prescriptive force relevant to the analysis of imperatives (which is well-modeled with planning content) and the type of prescriptive force associated with proffering a property for cognitive adoption.

---

23 MacFarlane tends to use eventive language (e.g., where to "*draw the line*" between tall and non-tall) to describe the kind of cognitive question that a gradable adjective in the positive form is meant to resolve. There is nothing grammatically marked about telling one's addressee *to draw the line* between tall and non-tall in such a way that Steph counts as tall. Still, how I draw this line isn't really up to me: how I draw the line determines who I think tall, and who I think tall isn't really up to me. It therefore seems like a distortion (prima facie) to represent the cognitive question that a gradable adjective in the positive form is meant to resolve as a question about where to plan to draw the line between tall and not.

A small revision of MacFarlane's semantics avoids this worry. Letting *c* be a context in which "Steph is tall" receives a nondescriptive interpretation, we will say:[24]

$$[\![\textit{Steph is tall}]\!]^c = \lambda\theta.\textit{deg}_{\textit{tall}}(\textit{Steph}) > \theta(\textit{deg}_{\textit{tall}})$$

This is the property a higher-order threshold function has if it maps the degree measure $\textit{deg}_{\textit{tall}}$ into a value that is lower than Steph's degree of tallness. This property is the sort of property an addressee might be encouraged to "cognitively instantiate" by a speaker (whose communicative aim is to get the addressee to share their appraisal of Steph as tall).

It is already common ground in this debate that agents cognitively represent entities of type θ—such entities are, after all, represented in the dominant theory of what a speaker's *semantic competence* with respect to gradable adjectives consists in.[25] What will be controversial between a proponent of the Kennedy (2007) view and myself is the *functional role* of this sort of representation.

According to Kennedy's own interpretation of his semantics for gradable adjectives (and any form of Contextualism about gradable adjectives with which I am familiar), representing some $\theta_c(\textit{deg}_{\textit{tall}})$ as eligible in a context *c* is to represent *c* as being a certain way—for example, as being such that $\theta_c(\textit{deg}_{\textit{tall}})$ is the *c*-relevant threshold for the minimum degree of tallness needed to count as tall (see Kennedy 2007, 17ff.). Given this account—which is

---

24 The context does not bind anything to its right in this semantic proposal; this is a simplification. Context can fix the values of various parameters relevant to evaluating a tallness claim (e.g., a relevant comparison class, a discourse task that would be resolved with a way of sorting individuals in the comparison class into tall and not-tall, etc.). According to the proposal being advanced here, we would not be surprised to observe nondescriptive uses targeting these contextual parameters (e.g., expressions of *properties of* comparison classes or *properties of* discourse tasks). My view is that such interpretations are likely attested. This is *not* to say that a speaker can freely express a property of any contextual parameter—in fact, speakers' freedom in this domain appears to be tightly constrained: when an expression of semantic type *T* is *indexical* in nature, speakers do not use sentences embedding such an expression to express properties of objects of type *T*. (As noted above, speakers do not seem to use sentences embedding "that" to express a view about the appropriate referent of "that" in their context.) Khoo (2017) introduces a model that provides an appealing explanation of this fact (as well as an appealing characterization of the phenomenon of indexicality in general). On Khoo's analysis, the values of indexical expressions (here understood to include demonstratives) are very tightly anchored to "objective" features of the context—features that are not freely modulated by speech acts. (As mentioned in section 3.2, it doesn't appear that a speaker can use "that's a beautiful car" to express a prescription governing the use of "that.") Nonindexical context-sensitive expressions (which Khoo dubs "quasi-indexical") work differently: speakers can use these expressions in a way that is not deferential to objective features of the context, as a way of modulating the features of the context toward which a referential use would ordinarily be sensitive.

25 It is unlikely that speakers have what Harris (2017b) calls "central access" ("central" in the sense of Fodor 1983) to this sort of entity, and so it is unlikely that speakers will have central access to -abstracts in which the -term binds variables over such entities. Nothing in my account assumes central accessibility, in the relevant sense (e.g., I do not assume that speakers have de dicto intentions to express the semantic contents that figure in my account).

very well-suited to accounting for Specific + Descriptive readings of gradable adjectives in the positive form (and, with a little tinkering, Nonspecific + Descriptive readings too)—it would hardly be extravagant to add that agents might also target candidate thresholds with *normative judgments*: agents might (indeed, obviously do) regard certain candidate thresholds as *appropriate* (or *inappropriate*) ways of drawing the boundary between tall and not-tall.

In general, we will say that an agent regards a threshold $\theta(\delta)$ **as appropriate** iff $\theta(\delta)$ is consistent with her sortal attitudes toward any object x such that $\delta(x)$ is defined—so, for example, an agent regards $\theta(deg_{tall})$ as an appropriate threshold for tallness iff, for any x with some degree of tallness, the agent regards x as tall only if $deg_{tall}(x) > \theta(deg_{tall})$. It seems unlikely that the state of regarding $\theta(\delta)$ as appropriate can be understood as representing some way the world or context could actually be: $\theta(\delta)$ is not plausibly treated as a candidate for actuality (here again see MacFarlane 2016). On the face of things, the functional role of this sort of state is exhausted by its characterization (or, perhaps, determination or regulation) of an agent's sortal attitudes.[26]

Let $\Theta_{x,c}$ designate the set of thresholds that are consistent with an agent x's normative judgments about what is an appropriate way of drawing the boundary between tall and not-tall in c.[27] I have claimed that, in a context in which "Steph is tall" receives a nondescriptive interpretation, a speaker semantically expresses the following property of higher-order thresholds: $\lambda\theta.deg_{tall}(Steph) > \theta(deg_{tall})$ (henceforth abbreviated λ_θ). It is straightforward to say how an addressee representable with $\Theta_{addr_c,c}$ will respond to a speaker who semantically expresses this

---

26 Similarly, in the modal/conditional domain, there is evidence that the functional role of attitudes toward modal/conditional sentences cannot be understood as a representation of some way things could be. Yalcin (2011, 2012) adduces cognitive evidence that the class of possible descriptive contents fails to cover a core range of uses of epistemic and probabilistic talk. Charlow (2016b) and Russell and Hawthorne (2016) muster evidence from formal epistemology (Triviality Results) to suggest that a core range of uses of conditional and modal sentences *cannot* be assigned a propositional semantic value in context. And, of course, there is the famous result of Gibbard (1981), which shows that, if the indicative conditional is a two-place operator that respects modus ponens and import-export, any proposition expressed by the indicative conditional would have to be equivalent (pace the facts) to the proposition expressed by the material conditional.

27 This is not intended as a thesis about the content of the state of regarding a certain threshold as appropriate or eligible (although I am sympathetic to such an account for prioritizing modals; see Charlow 2018, Sec. 5). The account given in this paper is meant to remain neutral on questions about how to represent the content of this sort of state of mind, as well as questions about how to formally represent the update an addressee performs when she updates on a normative judgment bearing on what is an appropriate way of drawing the boundary between tall and not-tall in c. I do not assume that update goes via intersection of set-theoretically represented contents (although I utilize intersective operations to represent certain features of such an update). More generally, I do not assume that it is the job of linguistic theorizing to characterize a way of updating on a piece of semantic content. As I have argued elsewhere, requiring a theory to characterize such an update function will ultimately mean writing a (epistemological) theory of rational attitude revision into our representation of semantic competence—something to be avoided (Charlow 2014, 2016a).

property (if she comes to accept the speaker's utterance): she will no longer regard candidate thresholds that fail to satisfy λ_θ as appropriate, and so the set of thresholds that are consistent with her normative judgments about what is an appropriate way of drawing the boundary between tall and not-tall in c' (where c' is a context posterior to c, in which $addr_c$ has come to accept the speaker's utterance c) will be given by:

$$\Theta_{addr_c,c'} = \Theta_{addr_c,c} \bigcap \lambda_\theta$$

As with practical prescriptions (as expressed by imperatives), cognitive prescriptions (as associated with a nondescriptive interpretation of "Steph is tall") are associated with the illocutionary act of proffering a property for cognitive adoption—a property that is not assumed to always correspond to a state of representational belief in the truth of some proposition. States that do not correspond to states of representational belief divide into practical (planning, action-guiding) states and nonpractical states (e.g., regarding some threshold as an appropriate way of sorting agents into tall and not-tall) that play a *determinative role in fixing an agent's sortal attitudes*. States of the latter kind do not play a determinative role with respect to an agent's practical states, e.g., her intentions (although of course the sortal attitude associated with thinking Steph tall can *cause* an agent to have intentions, e.g., the intention to pick Steph in our pickup game).

This section has argued that speakers can, and do, semantically express properties of contextual parameters, and that this kind of *locutionary* act can be married to a satisfying account of the *illocutionary* function of such locutionary acts—that is, proffering this property for adoption by the addressee in the context. The account was illustrated with gradable adjectives but could be extended with a bit of effort to the "informational" parameter against which modals are semantically evaluated.

Because semantic determinacy is orthogonal to prescriptiveness, the model I have described here does not bear directly on the issues of indeterminacy canvassed above. This is deliberate.[28] Speakers can, for this reason, semantically express a *range of candidate properties* for adoption by their addressees. Recall the following case:

> [Context: A/B could be referring to either the car type or token.]
>   A: That might be a Ferrari.
>   B: No, but it might be a Maserati.

The account here will analyze this case as one in which B semantically expresses more than one candidate property for adoption by A—the property of not ruling out possibilities in

---

28 This is not to say that these issues are unconnected: handling the problems arising from the phenomenon of semantic underspecification *and* those arising from nondescriptive uses will require admitting exceptions to Referentialist metasemantics.

which the car token is a Maserati, as well as the property of not ruling out possibilities in which the car's type is the type *Maserati*. Assuming that the car token is a Maserati iff its type is *Maserati*, it seems B's communicative aims are realized regardless of which property A adopts. I see no reason to say that the metasemantics of natural language demands that B express one of these properties but not the other, when the realization of B's communicative aims does not.

6. Coordination and Expressivism

I have argued that speakers semantically can use semantically "underspecified" language to express properties of semantic parameters (like states of information and degree thresholds), and thereby express an attitude — loosely, the attitude of (being representable as) satisfying this property. How and why does *language* provide for this sort of thing?

So far as the *function* of language is concerned, I would not try to improve on the story told in Gibbard (1990) (which is similar to, but in certain respects more general than, the story told in Lewis 1969).[29] These introductory remarks give a good flavor of the account:

> The need for complex coordination stands behind much of the way language works in our thoughts, in our feelings, and in social life. It figures centrally in our emotional dispositions, especially for such morally significant emotions as outrage, guilt, shame, respect, moral admiration, and moral inspiration. Matters of coordination, in the picture I shall sketch, stand squarely behind the psychology of norms, and hence behind what is involved in thinking something rational or irrational. Primitive human life is intensely social. In the conditions under which we evolved, anyone's prospects for survival and reproduction depended crucially on the beneficial human bonds he could cultivate. Human cooperation, and coordination more broadly, has always rested on a refined network of kinds of human rapport, supported by emotion and thought.
>
> (Gibbard 1990, 26)

The practice of proffering properties (of semantic parameters) for cognitive acceptance or adoption—regardless of whether the property is one whose psychological instantiation is equivalent to representing the world a certain way—is a practice that plausibly facilitates coordination in causally and behaviorally significant features of informational, sortal, and motivational psychological states. It is no surprise that, if the account of this paper is on the right track, language users would avail themselves of a dedicated type of content for performing this sort of speech act.

---

29 For recent appeals to coordination in a pragmatic account of nondescriptive language use, see Yalcin (2011, 2012) and Charlow (2015).

This is an Expressivist theory (for reasons I hope will be at least somewhat apparent to the reader). Notice in particular the following core presupposition of this account: that there is a difference in meaning between the speech act of *expressing* (for example) a particular sortal attitude (the content of which I represent as a constraint on, or property of, degree thresholds) and the speech act of *saying that* one has that same attitude (the content of which I represent with an ordinary proposition, to the effect that one's own threshold for considering an object, say, to be tall has a certain characteristic, say, exceeding 183cm).

The account, however, departs from traditional Expressivist theories like Gibbard's, in one or two significant ways, on which I will end the paper by reflecting. First, and most obviously, Gibbard, perhaps owing to a background commitment to the Humean Theory of Motivation, recognizes two "kinds" of "content": propositional content (content which bears on what the world is like, modeled with sets of possible worlds) and planning content (content that bears on how to plan, modeled with sets – equivalently, properties—of Normative Systems or Hyperplans). For Gibbard (and, we have seen, for MacFarlane as well) non-propositional content is generally theorized as *planning* content. This paper has argued that non-propositional content comes in more varieties than planning content. Although we do express planning attitudes (as well as descriptive beliefs) with language, these are not the only types of states of mind (which are not equivalent to descriptive belief in the truth of a propositional content) that we use sentences of natural language to express. Far from it. Indeed, although I have not argued for it here, I am drawn to the thesis that, *whenever* a sentence's semantic value in context is semantically "parametrized" — whether to a state of information (as with epistemics), an experiencer (as with experiential language, e.g., predicates of personal taste like "tasty"), a degree threshold (as with gradable adjectives), a plan (as with practical language, like imperatives and deontic modals), etc. — an utterance of that sentence can be used by a speaker to express a property of the relevant parameter, and thereby to express an attitude that can be modeled using a set of such parameters. Expressivism's insights can be fruitfully extended and generalized to many different types of language and language use, provided we are willing to entertain the sort of "polymorphism" (in essence, type-heterogeneity) about semantic content at which I am gesturing here. (On content polymorphism for epistemics, see my 2020.)

Second, Gibbard has offered a broadly Gricean account of the function of the speech act *expressing an attitude* (see also Gibbard 2003, 78ff.):

> Suppose Caesar tells Cleopatra, "I was captured by pirates in my youth." Why might he do this? Assume he is simply informing her about his youth; the story, then, will be something like this. He wants her to know about his capture by pirates. He thinks she lacks true belief on the subject, but he believes that she thinks him sincere and that she thinks him an authority on events of his youth. Here to be sincere is to express only beliefs one actually has, and to be an authority on something is to be quite unlikely to be mistaken about it. Caesar thus intends to get Cleopatra to believe that he was captured by pirates in his youth, and to do so in the

following manner. He utters words that conventionally purport to express, on the part of any speaker, a belief that he was captured by pirates in his youth. He intends her to come to accept that he has that belief, and to do so in virtue of her recognition of this intention. Since she takes him to be sincere, she has reason to accept, upon hearing his words, that he does believe that he was captured by pirates in his youth. Since she thinks him an authority on his youth, she concludes from his believing it that he indeed was captured by pirates in his youth.

(Gibbard 1990, 85).

On Gibbard's broadly Gricean account, a speaker *S* expresses state of mind *M* in order to get her addressee to form the belief that *S* is in *M* (a belief that will, under the right conditions, get the addressee to adopt *M* herself). Coordination in attitude is a natural effect (in a context in which the addressee recognizes the speaker as an authority on the relevant subject-matter) of the addressee forming a *belief about the speaker's state of mind*, on the basis of the speaker making a linguistic performance that indicates her possession of that state of mind.

Here, however, is a difficulty with this explanation.[30] We *have* dedicated linguistic devices (i.e., attitude ascriptions) for telling our addressees what states of mind we are in: instead of expressing his belief that he was captured by pirates in his youth (by asserting that he was), Caesar can self-ascribe the belief (by reporting to Cleopatra that he believes he was captured by pirates in his youth). If the speaker's aim is coordination in attitude, and coordination in attitude is explained by an addressee's belief about the speaker's attitude, speakers could realize the same communicative aim by reporting themselves to be in the relevant state. This, I will argue, presents a threat to the claimed explanatory role of non-propositional (e.g. planning-type) semantic content in Gibbard's broader theory.

To better see the threat, let us consider a "Subjectivist" alternative to an Expressivist account of normative claims. According to the Subjectivist alternative, when someone asserts that *x* is rational, they are semantically expressing a proposition about their own planning attitudes: roughly, the proposition that, according to their plans, *x* is permitted. (That is to say, according to the Subjectivist, they are semantically ascribing a plan permitting *x* to themselves.) Piggybacking on Gibbard's pragmatic theory, the Subjectivist might say the speaker typically does this in order to get their addressee to come to accept that their planning attitudes are this way; insofar as the addressee regards the speaker as an authority on whether to have plans that permit *x*, the addressee will have reason to share/adopt the speaker's planning attitude toward *x*. This Subjectivist theory works much the same as Gibbard's. And so one may begin to wonder what explanatory role the assignment of *non-propositional semantic content* is supposed to fill in Gibbard's semantic and pragmatic theory for normative language.

---

30 For a related critique, see Schroeder (2008, Section 4).

To forestall an obvious reply, it is true that the Subjectivist theory does not directly account for the (evident) difference in meaning between an attitude ascription in the mold of (1) and a corresponding attitude ascription in the mold of (2).

(1) Beth believes (says/agrees/disagrees) that x is rational.
(2) Beth believes (says/agrees/disagrees) that her plans permit x.

It does not follow that this difference in meaning is *incompatible* with the Subjectivist theory. Subjectivism, as stated above, is purely a thesis about the semantic content of an utterance of "x is rational" (relative to a context of utterance). It incurs no direct commitments regarding the semantic content of such a clause, as *embedded* under an attitude or illocutionary verb. To see this more clearly, consider a version of Subjectivism according to which:

- At a context c providing a variable assignment g, "x is rational$_n$" expresses the proposition that $g(n)$ permits x.
- By default, $g(n)$ is the planning attitude of c's speaker.

It is a live (indeed quite plausible) possibility in semantic theory that attitude and illocutionary verbs *quantify over* (and thereby shift) variable assignments (in addition to quantifying over more familiar objects like possible worlds) (see e.g. Santorio 2012). Adapting the idea, it is a live possibility, for the Subjectivist, that the truth condition of "Beth believes that x is rational" is roughly that, for any planning attitude n (such that n is compatible with Beth's plans), n permits x. This truth condition — which we can gloss as "Beth's plans permit x" — is evidently distinct from the truth condition of "Beth believes that her plans permit x". Note that this form of Subjectivism agrees with Gibbard about the content of Beth's belief that x is rational: that its content is best represented, roughly, with a *plan* (rather than a proposition).

Expressivism, then, or Subjectivism? I will try to explain why I still incline toward Expressivism (although Gibbard's theory of semantic interpretation means he will struggle to distinguish his Expressivist theory in similar fashion). The Subjectivist disagrees with the Expressivist about the *semantic content* of Beth's assertion that x is rational. Their reasons (as I imagine them) are something like this: on our best theory of assertion (Stalnaker's), the essential effect of an assertion is to update the Common Ground with a propositional content. Now, there is no doubt that a speaker who asserts that x is rational *does* generally make it Common Ground that *their* plans permit x. Since the proposition that the speaker's plans permit x is generally part of the post-assertion Common Ground anyway in such cases, it seems to make good theoretical sense (for the Subjectivist) to say that such a speaker semantically expresses (locutes) the proposition that their plans permit x in the course of asserting (illocuting) that x is rational. In contexts where we have reason to interpret the speaker as endorsing this feature of their plan as a basis for rational coordination, prescription-type interpretations of claims like "x is rational" will arise in broadly the same way as on Gibbard's theory.

This form of Subjectivism shares with Gibbard's theory an apparent aversion to the idea of what I'll call *intrinsically practical content* — content the apprehension and acceptance of which can "directly" constrain a planning state (perhaps, again, owing to background Humean assumptions). For both Gibbard and the Subjectivist, interpreters who accept speaker's claim that *x* is rational must be represented as reasoning their way from beliefs about the speaker's communicative intention to a (self-directed) normative judgment that subsequently constrains the interpreter's planning state, deploying something like the following syllogism.

(1) The speaker intends for me believe that their plans permit *x*.
(2) Given (1), the speaker intends for me to have plans that permit *x*.
(3) Given (2), my plans ought to permit *x*.
(4) So, my plans ought to permit *x*. [Conclusion: revise plans to permit *x*.]

Compare this to the following, rather simpler, account of the reasoning involved in accepting a speaker's claim that *x* is rational.

(1) The speaker expresses a way of planning that permits *x*.
(2) To accept the way of planning the speaker expressed, my plans must permit *x*.
(3) So: revise plans to permit *x*.

Expressing (we might also say "proffering") a way of planning that permits *x*, as I understand the notion, does not imply that one intends their addressee to believe that their plans permit *x*. This is a more flexible, less representationally committed, paradigm for understanding what internal representations an interpreter deploys when they accept a speaker's claim that *x* is rational. Notice, for instance, that the paradigm is easily fitted to cases of "selfless" direction, in which it is Common Ground that a speaker is expressing the view that *x* is permitted (a way of planning that permits *x*), despite having plans that prohibit *x*. The Gibbard-Subjectivist paradigm is not.

Another reason to want this kind of flexibility is the *natural ubiquity* of prescription-type content. Creatures with limited meta-representational capabilities comprehend prescriptions, by interpreting apprehended (non-linguistic) signals as attempts to constrain their behavior. A warning call is interpreted as carrying a prescriptive message — *(you should) tread carefully!* Accepting the message, so interpreted, means that the addressee's internal state is organized in such a way that they are behaviorally disposed to tread carefully. The intentional content of a pain experience can be productively, if partially, theorized as prescriptive — *(you need to) stop this!* The subject of a pain experience is theorized as the recipient of this message; accepting the message means having a plan that requires stopping the pain. In cases like these, must we represent the subject as arriving at a decision to accept the message's instruction as reasoning with representations about the "intentions" of

the message's producer? I think not: the subject simply *apprehends an instruction*, or way to plan, and decides — possibly, but not necessarily, after engaging in higher-level reasoning about the source of the instruction — whether or not to adjust their plans accordingly. Of course, I do not deny that interpreters often (perhaps always) utilize representations about the source of the instruction in trying to identify *which* instruction the source "means" to be transmitting. (Is the pain being sent by the stomach or the heart? To whom is the warning call directed?) Once an interpreter determines the content of the relevant instruction, the interpreter decides to accept it or not. While it is *possible* (depending, of course, on their representational capabilities) for the interpreter at this stage to engage in further practical reasoning, to try to determine whether or not it is a *good idea* to accept the interpreter's instruction, it does not seem to be a prerequisite.

So both Gibbard and the Subjectivist are, I believe, mistaken about the manner in which prescriptive (including planning) content is generally apprehended and accepted. In Gibbard's account (and the Subjectivist's), generating a prescriptive interpretation of a message is a *side-effect* of forming a specific belief about the internal state of the message's source. In our account, speakers express properties of cognitive parameters, and they do so in order to proffer those properties (directly) for cognitive adoption (acceptance) by their addressees. In such cases, a cognitive constraint — a *way of representing* on some question or issue — is inherent in the content that an addressee apprehends, when they apprehend the speaker's message, and the addressee must either accept the message (by representing the question that way) or not (otherwise). Language simply and directly provides speakers the tools for conveying cognitive prescriptions of sundry types and flavors, facilitating wide-ranging coordination in causally significant features of our internal states.

References

Asher, Nicholas, and Alex Lascarides (2001). "Indirect Speech Acts." *Synthese* 128: 183–228. https://doi.org/10.1023/A:1010340508140.

Barker, Chris (2002). "The Dynamics of Vagueness." *Linguistics and Philosophy* 25: 1–36. https://doi.org/10.1023/A:1014346114955.

Braun, David, and Theodore Sider (2007). "Vague, So Untrue." *Noûs* 41: 133–56.

Charlow, Nate (2011). "Practical Language: Its Meaning and Use." www.natecharlow.com/work/dissertation.pdf. Accessed August 2021.

——— (2014). "Logic and Semantics for Imperatives." *Journal of Philosophical Logic* 43: 617–64. https://doi.org/10.1007/s10992-013-9284-4.

——— (2015). "Prospects for an Expressivist Theory of Meaning." *Philosophers' Imprint* 15: 1–43.

——— (2016a). "Decision Theory: Yes! Truth Conditions: No!" In *Deontic Modality*, edited by N. Charlow and M. Chrisman, 47–81. Oxford: Oxford University Press.

——— (2016b). "Triviality for Restrictor Conditionals." *Noûs* 50: 533–64. https://doi.org/10.1111/nous.12111.

——— (2018). "Clause-Type, Force, and Normative Judgment in the Semantics of Imperatives." In *New Work on Speech Acts*, edited by D. Fogal, D. Harris, and M. Moss, 67–98. Oxford: Oxford University Press.

——— (2020). "Grading Modal Judgement." *Mind* 129: 769–807. https://doi.org/10.1093/mind/fzz028.

Dowty, David (1985). "On Recent Analyses of the Semantics of Control." *Linguistics and Philosophy* 8: 291–331. https://doi.org/10.1007/BF00630916.

Fodor, Jerry (1983). *The Modularity of Mind*. Cambridge: MIT Press.

Gibbard, Allan (1981). "Two Recent Theories of Conditionals." In *Ifs*, edited by W. Harper, R. Stalnaker, and G. Pearce, 211–47. Dordrecht: D. Reidel.

——— (1990). *Wise Choices, Apt Feelings*. Cambridge: Harvard University Press.

——— (2003). *Thinking How to Live*. Cambridge: Harvard University Press.

Harris, Daniel W. (2017a). "Imperative Inference and Practical Rationality" (Forthcoming in *Philosophical Studies*) https://doi.org/10.1007/s11098-021-01687-0.

——— (2017b). "Semantics without Semantic Content." (Forthcoming in *Mind and Language*) http://doi.org/10.1111/mila.12290.

Heim, Irene, and Angelika Kratzer (1998). *Semantics in Generative Grammar*. Oxford: Blackwell.

Jacobson, Pauline (1999). "Towards a Variable-Free Semantics." *Linguistics and Philosophy* 22: 117–85. https://doi.org/10.1023/A:1005464228727.

Kaufmann, Magdalena (2012). *Interpreting Imperatives*. Dordrecht: Springer.

Kennedy, Christopher (2007). "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives." *Linguistics and Philosophy* 30: 1–45. https://doi.org/10.1007/s10988-006-9008-0.

Khoo, Justin (2017). "Quasi-Indexicals." *Philosophy and Phenomenological Research* 100, 1: 26–53. http://doi.org/10.1111/phpr.12519.

King, Jeffrey C. (2014). "The Metasemantics of Contextual Sensitivity." In *Metasemantics: New Essays on the Foundations of Meaning*, edited by A. Burgess and B. Sherman, 97–118. Oxford: Oxford University Press.

——— (2017). "Felicitous Underspecification and Updates."

——— (2018). "Strong Contextual Felicity and Felicitous Underspecification." *Philosophy and Phenomenological Research* 97: 631–57. https://doi.org/10.1111/phpr.12393.

Kratzer, Angelika (1977). "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy* 1: 337–55. https://doi.org/10.1007/BF00353453.

——— (1981). "The Notional Category of Modality." In *Words, Worlds, and Contexts*, edited by H. Eikmeyer and H. Rieser, 38–74. Berlin: De Gruyter.

——— (1991). "Modality." In *Semantics: An International Handbook of Contemporary Research*, edited by A. von Stechow and D. Wunderlich, 639–51. Berlin: De Gruyter.

Lewis, David (1969). *Convention*. Cambridge: Harvard University Press.

——— (1979). "A Problem about Permission." In *Essays in Honour of Jaakko Hintikka*, edited by E. Saarinen, R. Hilpinen, I. Niiniluoto, and M. B. Provence Hintikka, 163–79. Dordrecht: D. Reidel.

——— (1980). "Index, Context, and Content." In *Philosophy and Grammar*, edited by S. Kanger and S. Öhman, 79–100. Dordrecht: D. Reidel.

MacFarlane, John (2011). "Epistemic Modals Are Assessment-Sensitive." In *Epistemic Modality*, edited by B. Weatherson and A. Egan. Oxford: Oxford University Press.

——— (2014). *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.

——— (2016). "Vagueness as Indecision." *Aristotelian Society Supplementary Volume* XC: 255–83. https://doi.org/10.1093/arisup/akw013.

Moss, Sarah (2013). "Epistemology Formalized." *The Philosophical Review* 122: 1–43. https://doi.org/10.1215/00318108-1728705.

——— (2015). "On the Semantics and Pragmatics of Epistemic Modals." *Semantics and Pragmatics* 8: 1–81.

Ninan, Dilip (2005). "Two Puzzles about Deontic Necessity." In *New Work on Modality, Vol. 51 of MIT Working Papers in Linguistics*, edited by J. Gajewski, V. Hacquard, B. Nickel, and S. Yalcin. Cambridge: MITWPL. http://semanticsarchive.net/Archive/WZlNTg0Y/.

Plunkett, David, and Timothy Sundell (2013). "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosophers' Imprint* 13: 1–37. http://hdl.handle.net/2027/spo.3521354.0013.023.

Portner, Paul (2004). "The Semantics of Imperatives within a Theory of Clause Types." In *Proceedings of SALT 14*, edited by Robert B. Young, 235–52. Ithaca: CLC Publications. http://semanticsarchive.net/Archive/mJlZGQ4N/.

——— (2007). "Imperatives and Modals." *Natural Language Semantics* 15: 351–83. https://doi.org/10.1007/s11050-007-9022-y.

——— (2018). "Commitment to Priorities." In *New Work on Speech Acts*, edited by D. Fogal, D. Harris, and M. Moss, 296–316. Oxford: Oxford University Press.

Roberts, Craige (2018). "Speech Acts in Discourse Context." In *New Work on Speech Acts*, edited by D. Fogal, D. Harris, and M. Moss, 317–59. Oxford: Oxford University Press.

——— (2015). "Conditional Plans and Imperatives: A Semantics and Pragmatics for Imperative Mood." In *Proceedings of the 20th Amsterdam Colloquium*, edited by T. Brochhagen, F. Roelofsen, and N. Theiler, 353–62. Amsterdam: University of Amsterdam: Institute for Logic, Language, and Computation.

Russell, Jeffrey Sanford, and John Hawthorne (2016). "General Dynamic Triviality Theorems." *The Philosophical Review* 125: 307–39.

Schroeder, Mark (2008). "Expression for Expressivists." *Philosophy and Phenomenological Research* 76: 86–116. https://doi.org/10.1111/j.1933-1592.2007.00116.x.

Searle, John R. (1975). "Indirect Speech Acts." In *Syntax and Semantics Volume 3: Speech Acts*, edited by P. Cole and J. Morgan, 59–82. New York: Academic Press.

Stalnaker, Robert (1978). "Assertion." In *Syntax and Semantics 9: Pragmatics*, edited by P. Cole, 315–32. New York: Academic Press.

——— (1984). *Inquiry*. Cambridge: MIT Press.

Starr, William (2020). "A Preference Semantics for Imperatives." *Semantics and Pragmatics*, forthcoming. http://dx.doi.org/10.3765/sp.13.6.

Stevenson, C. L. (1937). "The Emotive Meaning of Ethical Terms." *Mind* 46: 14–31.

Swanson, Eric (2006). "Interactions with Context." http://www-personal.umich.edu/~ericsw/Swanson,%20Interactions%20with%20Context.pdf. Accessed August 2021.

——— (2016). "The Application of Constraint Semantics to the Language of Subjective Uncertainty." *Journal of Philosophical Logic* 45: 121–46. https://doi.org/10.1007/s10992-015-9367-5.

Veltman, Frank (1996). "Defaults in Update Semantics." *Journal of Philosophical Logic* 25: 221–61. https://doi.org/10.1007/BF00248150.

Yalcin, Seth (2011). "Nonfactualism about Epistemic Modality." In *Epistemic Modality*, edited by A. Egan and B. Weatherson, 295–332. Oxford: Oxford University Press.

——— (2012). "Bayesian Expressivism." *Proceedings of the Aristotelian Society* CXII, Part 2: 123–60. https://doi.org/10.1111/j.1467-9264.2012.00329.x.