Ergo AN OPEN ACCESS JOURNAL OF PHILOSOPHY

Underspecification and Communication

RAY BUCHANAN
University of Texas, Austin

It has recently been argued that our use of vague language poses an intractable problem for any account of content and communication on which (i) the things we assert are propositions and (ii) understanding an assertion requires recognizing *which proposition* the speaker asserted. John MacFarlane has argued that this problem concerning vague language is itself a species of an even more general problem for such traditional accounts—the problem posed by "felicitous" underspecification. Repurposing certain ideas from Allan Gibbard, MacFarlane offers a novel theory of vagueness, *plan-expressivism*, as an account that can handle both vagueness and such underspecification in communication.

In this paper, I argue that despite its many virtues, plan-expressivism fails as a general account of meaning and communication. In particular, I show that when it comes to the problem of felicitous underspecification, plan-expressivism fares no better than the more traditional accounts of content and communication it is intended to replace. Along the way, I argue that the problem of felicitous underspecification puts considerable pressure on an assumption that is at the core of both plan-expressivism and the accounts of its more traditionalist rivals – namely, that in cases of successful linguistic communication, we must always be able to find something – whether a propositional content or a conversational update – that the speaker is intending to share with her audience.

Keywords: felicitous underspecification; vagueness; content; communication; context-sensitivity

1. Introduction

It has recently been argued that our use of vague language poses an intractable problem for standard accounts of content and communication according to

Contact: Ray Buchanan <lraybuchanan@gmail.com>

https://doi.org/10.3998/ergo.6922

which (i) the things we assert are propositions and (ii) understanding an assertion requires recognizing *which proposition* the speaker asserted. According to John MacFarlane, this problem concerning vague language is itself a species of an even more general problem for such traditional accounts—the problem posed by "felicitous underspecification".[1] Repurposing certain ideas from Allan Gibbard, MacFarlane offers a novel theory of vagueness, *plan-expressivism*, as an account that can handle both vagueness and such underspecification in communication.

Despite its numerous virtues, the plan-expressivist account of vague content and communication fails. It is no better suited to account for the general problem posed by felicitous underspecification than are the traditional accounts of content and communication that it is intended to replace. Towards this end, I'll show that *even if* we were to assume that the contents of our vagueness-related propositional attitudes are themselves Gibbardian contents of the kind favored by MacFarlane, the problem posed by such underspecification remains. One upshot is that, *contra* MacFarlane, the phenomena of vagueness and of communicative underspecification require different accounts; even if he is right about vagueness, more needs to be said regarding underspecification. Another upshot is that expressivists in any domain (regarding, say, moral discourse or vagueness) need a more fully worked out story regarding what it means to "express" an attitude and how exactly this relates to the theory of meaning and communication. Along the way, I hope to get clearer on the source of felicitous underspecification and why theorizing about this phenomenon at the abstract level of "updates" on common ground threatens to obscure the phenomenon rather than illuminate it.

I begin with a general discussion of the problem posed by felicitous underspecification (§2) and why MacFarlane thinks vague language is a subspecies thereof (§3). After sketching the positive plan-expressivist picture of vagueness, along with some of the reasons it seems so tempting, (§3), I then show why the problem of underspecification remains (§4). I conclude with some forward-looking remarks on why the issues concerning communication that worry MacFarlane will ultimately require "clouds of (vague) contents" whether we identify those contents with structured entities containing vague constituents or with the Gibbardian sets of world/hyperplan pairs he favors (§5).

2. The Source of Felicitous Underspecification

In a cooperative communicative exchange, you speak with the goal of being understood and your audiences engages with the goal of understanding. But what is it to understand an utterance? Pinning down necessary and sufficient

1. See MacFarlane (2016; 2020a; 2020b; 2020c).

conditions has proven difficult, but we might (at least provisionally) agree on at least this much: understanding an utterance requires successful recognition of both the kind and content of the speech act(s) the speaker has performed by uttering what she did. Recognizing that it was an assertion, rather than a question or a command, is required. But, also, the audience must successfully recover the *content* of the speech act; *what* was asserted, commanded, or so on. Insofar as asserting that *p* requires the speaker to mean, or intend to communicate, *p*, understanding an assertion requires the hearer to recover a content that the speaker intends for her to entertain.

How is such understanding ever achieved? Fortunately, our words provide excellent evidence of the speech acts we use them to perform. In the simplest of cases, the content of your speech act might be nothing over and above the standing, conventional meaning of your words. Focus on assertion.[2] Suppose that you utter (1) as an intentional expression of (and only of) your belief that two is prime:

(1) Two is prime.

If your hearer can (somehow) safely assume that you are speaking literally, she can simply "read off" *what* you asserted directly from your specific choice of words on that occasion. But, of course, normally more work will be required on the part of your audience than just that since what we mean and assert so often goes beyond what we can reasonably expect our words by themselves to deliver. For example, I can literally utter (2) as an assertion of (*) that Resnik and Passlof got married *to each other* even though that is not itself given by the standing meaning the words I've used (or the semantic significance of their particular syntactic configuration):

(2) Resnik and Passlof got married.

Moreover, in the right contextual setting, I can reasonably expect for you to recognize that I have asserted (*) by (2); suppose, for e.g., that we mutually know that Resnik and Passlof are at the gallery opening, that you just asked me who in attendance has recently gotten married, and that I am feeling helpful and would like to help answer your question, etc. In short, we can often rely on our audience's competence with the words we've used *and* facts about the context of utterance to correctly recognize what we have asserted.

So far, we have spoken simply of the *content* of an assertion, i.e., *what* the speaker asserts. But what kind of thing is that? It is standardly assumed that the

2. I assume that *asserting* is the genus of which *saying, telling*, and *literally meaning* are species.

things we assert are *propositions*—entities that determine possible worlds truth-conditions. According to some, a proposition can simply be identified with the set of possible worlds at which it is true; according to others, they are more fine-grained, structured entities that somehow encode such truth conditions. For present purposes, we don't need to settle these issues since the assumption we will rely is something that all would agree to—namely, that propositions have classically-construed (non-relativized) truth-conditions. Call this minimal, standard assumption regarding what we assert "Content$_a$" (short for *the Content of Assertion Thesis*):

(Content$_a$) The things we assert are propositions.[3]

Putting Content$_a$ together with the earlier claim that understanding an assertoric speech act requires the speaker's audience recognizing (i) that it was an assertion (the kind of speech act) and (ii) the content thereof, at least in part on the basis of the speaker uttering what she did, we arrive at a version of what I call *the Standard View of Meaning and Communication* (Buchanan 2010). As MacFarlane points out, an immediate consequence of such a view is that communication succeeds only if the hearer grasps "the truth-conditions intended by the speaker" (MacFarlane 2020a, 597).[4]

Notice that for the proponent of the Standard View, a case such as (2) already suffices to show that what we assert can be *underspecified* relative to the particular words we've used; the truth-conditions of my assertion are underspecified by the words I uttered. The proponent of the Standard View will claim, however, that this variety of underspecification is unproblematic for her account of successful communication since the gap between our words and what we assert can so easily be bridged by the knowledge the speaker and hearer can reasonably expect one another to have about each other and the conversation in which they are engaged. What is left underspecified by our words, becomes specified once we take into consideration the relevant facts about the context of utterance.

In recent years, however, several theorists have argued that there *are* cases of underspecification that genuinely threaten the Standard View.[5] For example, in Buchanan (2010), I argued that that there are cases of fully successful communication in which both the speaker's words, *and* the surrounding facts about

3. See the discussion of the "Content Thesis" in my (2010). MacFarlane calls this thesis "Classical Contents" (MacFarlane 2020b, 619).

4. MacFarlane characterizes the Standard View as requiring that (in the assertoric case) communication is successful "just in case the hearer grasps the truth-conditions intended by the speaker" (597). Here, however, we will only need the weaker (and less potentially contentious) necessary condition stated above.

5. See, e.g., Abreu Zavalata (2021), Buchanan (2010), Buchanan & Ostertag (2005), Gillies & Von Fintel (2011), and King (2018; 2021) for more.

the context of utterance, fail to determine any particular proposition to be what she meant or asserted (or as something the hearer must take her to have asserted if she is to understand). In such cases, it seems that there is a multiplicity of non-truth conditionally equivalent "candidates" for what the speaker meant/asserted with nothing to privilege one over the others as being what the speaker asserted. Even taking the speaker's intentions, her words, and the context into consideration, underspecification remains, but is acceptable and does not hinder smooth communication. The remaining underspecification is, to borrow King's helpful terminology, "felicitous" (2018; 2022). Let's say that underspecification is *catastrophic* when it leads to communicative failure—when there is no possibility of the speaker's utterance being understood by the hearer—and that it is *felicitous* otherwise. In what follows, I will be concerned exclusively with the latter, felicitous variety of underspecification. Hereafter, "underspecification" should be read with this qualifier in mind.

To illustrate the phenomenon, consider one of my examples from (Buchanan 2010). While getting ready for a beginning of term (blowout) party, Chet utters (3) to his fellow student, Tim:

(3) Every beer is in the bucket.

What Chet asserted depends on the relevant domain restrictions on the quantifiers 'every beer' and 'the bucket'.[6] But, in this case, there are numerous, truth-conditionally inequivalent "candidates" for what the speaker asserted. For example:

P_1 Every beer *we bought at the bodega* is in the bucket *in the backyard*.
P_2 Every beer *we will serve at the party* is in the bucket *decorated in pirate motif*.
P_3 Every beer *at the apartment* is in the bucket *next to the hot tub*.
P_4 Every beer *for our guests* is in the bucket *filled with ice*.
P_5 Every beer *we bought at the bodega* is in the bucket *next to the hot tub*.
P_6 Every beer *at the apartment* is in the bucket *in the backyard*.[7]

I argued that, even if (as is doubtful) the speaker had any particular one of these propositions "in mind" in uttering what she did, she cannot reasonably intend to communicate or assert it as opposed to any other of the candidates. In such a case, how could she reasonably expect her audience to get just that content on the basis of her utterance? But, likewise, on the side of the audience, there doesn't seem to be any particular candidate that the hearer *must* entertain if she is to

6. I am assuming a Russellian treatment of definite descriptions.

7. Here the context-invariant meaning of the sentence is understood as something akin to an "gappy" propositional template with admissible candidates to be "fillings in." See Bach (2001) and Harris (2022) for more.

understand the utterance (she needn't just get P_1, because entertaining P_4, or the conjunction of P_5 and P_3, and so on would have just as well sufficed, and so on). After arguing against the obvious possible responses (the speaker asserted the conjunction, or disjunction of the candidates, or a vague content with vague constituents corresponding to the domain restrictions, or simply 'All *those* beers…,' and so on), I concluded that the speakers' communicative intentions exhibit a certain *generality* and *indifference* that cannot be captured by supposing that they meant, or asserted, that, or any other particular proposition.[8] On this basis, I argued that Content_a should be rejected in favor of a view on which such an utterance is associated with a range or "cloud" of contents that is consonant with the speaker's communicative intentions (more anon).[9]

More recently, King (2018, 2022) has offered numerous examples of this phenomenon and a book-length, novel treatment concerning how we might accommodate it in a more general picture of context-sensitivity, update, and common ground. Consider one of his many examples. Suppose you are at the beach watching a group of surfers and you turn to your friend and utter (4):

(4) Those guys are good.

There need not be any particular plurality of surfers that the speaker is referring to by 'those guys'; As MacFarlane puts it, in discussing King's example:

> The speaker may not have a definite intention that settles whether
> (a) *those guys* = Abe, Bob, Cindy, Maria, Zeke, or
> (b) *those guys* = Abe, Bob, Cindy, Maria, Sid.
> Despite this, the assertion of (5) can be felicitous, and understanding it does not seem to require deciding between (a) and (b). (2020b, 618)

Though the complex demonstrative lacks a unique semantic value in context, this does not seem in any way to hinder the smooth flow of communication. In such a case, "failure to coordinate on a determinate 'supplement' for a context-

8. In Buchanan (2010), I didn't consider the possibility of Stalnakerian diagonalization. The arguments that MacFarlane offers against this proposal in (2020a, 612–616) help to close this lacuna.

9. Gillies and von Fintel (2010) argue for a similar conclusion, focusing on epistemic modals such as in (i):

> (i) Passlof might have been a student of de Kooning at Black Mountain College.

Epistemic modals must be evaluated for truth with respect to information states; but unlike rival contextualist accounts, they claim that context does not, and need not, uniquely specify *the* information state relative to which the modal should be assessed. In using an epistemic modal, a speaker can 'put into play' a cloud of propositions even though she would not be entitled to assert each of the contents in that cloud (2010).

sensitive term does not seem to be required for communicative success" (2020b, 611). Here, and in other cases of felicitous underspecification, "singling out a unique semantic value is not crucial to the communicative aims of the speakers" (King 2018, 651).

Myself (Buchanan 2010) and King (2018; 2022) both argue that underspecification can arise in virtually *any* case of context-sensitive language use. And despite some potentially significant disagreements about the details, we both reject the simple model of communication sketched above, in favor of an account that allows for a perfectly literal, assertoric utterance to be associated with a plurality, or a cloud, of contents, rather than any unique content that the speaker and hearer must "coordinate on." According to us both, there is often no (single) propositional content to be identified as *the* content of the speaker's utterance.

As we are about to see, MacFarlane claims that our use of vague language gives rise to a problem for the standard view of meaning and communication that is a "special case of the more general problem of explaining what King has called felicitous underspecification" (2020a, 616). Moreover, MacFarlane agrees that in such cases there need not be any unique content of *the standard truth-conditional variety* to be found. But rather than adopt a non-standard, cloudy account of vagueness and communication in response, MacFarlane innovates in the theory of content. Roughly put, MacFarlane argues that, in the relevant range of cases, we *can* identify a unique content as what is asserted, albeit not a content of the kind you were probably expecting.

3. Vagueness and Plan-Expressivism

If vague expressions are context-sensitive, it should be unsurprising if they can exhibit underspecification. Consider the following case MacFarlane offers as an illustration. Suppose that a chemist and her assistant are looking at a row of test tubes arranged the in order of height from tallest (25 mm) to shortest (10 mm):

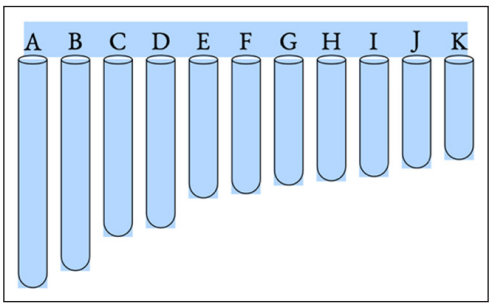

Figure 1.1: The test tubes.
(MacFarlane 2020a, 598)

If the chemist utters "the tall test tube contains hydrofluoric acid," it is plausible enough that the speaker intended, and can be recognized as having intended, to communicate a particular proposition: roughly, that Tube A [in Figure 1.1] contains hydrofluoric acid. This can be so, even if there is no mutually understood, agreed upon "cut off" for what counts as 'tall.' But, now suppose the chemist says (5):

(5) Some of the tall test tubes contain hydrofluoric acid.

MacFarlane:

> … [I]f successful communication requires the hearer to recognize the truth-condition the speaker intends to communicate, it is puzzling how the utterance of [(5)] *could* be a successful communication. For recognizing the truth condition that the chemist means to get across in uttering [(5)] would seem to require knowing where she puts the cutoff point for 'tall.' (MacFarlane 2020a, 599)[10]

Even if the speaker did have a particular cut-off for 'tall' in mind, how could she reasonably intend her utterance of (5) to put her audience in a position to secure uptake? Plausibly, the speaker did not intend to communicate any particular cut-off involving proposition, nor does her hearer need to recover any such thing from the utterance in order to understand it.

Vague expressions are often context-sensitive in more than one sense. For example, a gradable adjective such as 'tall' is contextually-sensitive to the specification of a relevant class—*for an NBA star*, *for a building*, etc. But even if we fix the comparison class, there are contextually varying norms regarding what would count as a *typical* case of being *tall for an F* (Graff Fara 2000). And as MacFarlane points out, there is still room for further contextual variation for "multi-dimensional" vague predicates such as 'smart' (2020a, 610). In short, there is a significant gap between what is explicitly given by the context-invariant meaning of 'tall' and what your audience would have to get from your utterance in order to recover anything that could plausibly be identified with "the proposition" that you have asserted. MacFarlane is skeptical that this gap can be filled in

10. MacFarlane argues that even if there is no mutually recognized cutoff, communication need not fail, as it clearly would if, for example, the chemist were to have said "those test tubes contain hydrofluoric acid" without "making any attempt to clarify which test tubes she is demonstrating" (2020a, 599).

Of course, not every conversation can happily proceed with the cutoff left underspecified in this way. Indeed, it is even easy to imagine various ways the conversation involving (5) might further evolve such that mutually settling a cutoff would seemingly be required for successful communication. For MacFarlane's purposes, however, it is enough to show that, at least sometimes, the cutoff can remain felicitously underspecified.

all cases of successful communication involving vagueness. Nor need it be: the remaining underspecificity can (at least sometimes) be felicitous.

But MacFarlane is no fan of "clouds" of classically construed, truth-conditional contents. Indeed, he argues at length that the (then) extant cloudy accounts of communication and content should be rejected since they cannot provide plausible accounts of conversational update (MacFarlane 2020b). These critical arguments against rival accounts are important, but for now, let's focus on his positive, non-cloudy view on which we can (he claims) identify a single, unique asserted content even in cases of underspecification.

According to MacFarlane, vagueness is literally a result of indecision; in particular, indecision with regard to how, exactly, to apply our words (and concepts). MacFarlane claims our usage of vague language reflects such indecision. For example, we might be resolute in applying the term 'stubble' to uniform, distributions of 2 mm hair on the lower half of someone's face, but not to 10 mm ('beard' would clearly apply instead). But what about 5 mm? Here we are simply undecided. Our ambivalence about borderline cases is not (*contra* the Epistemicist) a matter of doxastic uncertainty about the location of the (putative) cut-off, it is a matter of *practical* indecision (MacFarlane 2020c, 647).[11] Since "vague statements are partly practical, partly doxastic" (647), they cannot be understood as merely providing constraints on the possible-worlds truth-conditions; they also provide constraints on plans of action.

As a first step towards MacFarlane's implementation of this insight concerning the connection between planning and vagueness, consider his example of (6) below. Suppose that it is already mutual knowledge between us that Richard's height is somewhere between 5'9" and 6'1", but we aren't sure exactly where in that range his actual height falls. As MacFarlane points out, if I literally utter (6), I need not be intending to rule out any particular worlds *or* any particular thresholds for 'tall':

(6) Richard is tall.

Rather, I might simply be aiming to link "the question of which thresholds govern 'tall' to the question of Richard's actual height"; I am proposing that, moving forward, we should count anyone as tall as, or taller than, Richard as 'tall'. In such a case, "the intended effect on the common ground cannot be modeled as a change either to a factual common ground (a set of possible worlds) or to a non-actual scoreboard—including, for example, a range of *delineations* that set thresholds for the gradable adjectives." (2020c, 644).

11. As MacFarlane argues, we don't use the language of subjective uncertainty to talk about borderline cases, as would be expected if we were dealing with cases of doxastic uncertainty. Rather. we talk about what can, or should, "count as" as falling within the extension.

What is needed, MacFarlane claims, is to understand the common ground as not simply a set of possible worlds (or, for that matter, a set of possible worlds along with a (disparate) set of delineations), but instead as a set of world/delineation *pairs*. With this shift, it is then possible to understand an (assertoric) update as a proposal of how to modify common ground; roughly as a proposal about which world/delineation pairs to keep or to eliminate. For example, in the case of (6), perhaps the proposed update is to "rule out all world/delineation pairs <*w*,*d*> at which Richard's height in *w* falls short of the threshold established for 'tall' by *d*"(2020c, 644).

But how exactly should we understand this new conception of common ground given that there is virtually never a mutually known, contextually-determined, precise threshold for the applicability of a vague predicate?[12] The key, MacFarlane claims, is to give up thinking of a delineation (that sets a threshold) in a straightforwardly factualist manner, i.e., as some sort of fact about the relevant discourse. Rather, we should look to Gibbard's non-factualist, expressivist account of normative judgment for inspiration.

According to Gibbard (2003), a claim such as (7) is best understood as an 'expression' of a planning state:

(7) I ought to go to the Frankenthaler exhibit this week.

To judge (7), Gibbard claims, is to adopt a certain plan of action. When I judge (7), I rule out any plan that is not consistent with my getting to the exhibit this week. An ordinary plan such as this will invariably be underspecified. Though I may already have decided that I'll go there tomorrow by taxi, that leaves open whether I'm going at 7 or 8, whether I'll wear a suit and tie, and so on. To model the content of such a planning state, Gibbard introduces the notion of a *hyperplan*:

> A *hyperplan*, we can stipulate, covers any occasion for choice one might conceivably be in, and for each alternative open on such an occasion, to adopt the plan involves either rejecting that alternative or reject rejecting it. In other words, the plan either forbids an alternative or permits it. (2003, 56).

Hyperplans are "fully determinate contingency plans concerning every possible circumstance and resolving **all indecision**" (MacFarlane 2020c, 647, bold mine). An ordinary, less than fully resolved, plan such as the one Gibbard thinks I am expressing by (7) can then be modeled by the set of world/hyperplan pairs compatible

12. See (2020c, 657–661) for more details. MacFarlane argues there could be a determinate threshold set by the common ground even though we do not, and cannot, mutually know what it is.

with it. The basic logical operations – conjunction, disjunction, negation, and so on—can be given for these Gibbardian contents in the standard Boolean way.

Since MacFarlane thinks the vagueness of 'tall' ultimately stems from our indecision regarding where to draw a cutoff, his appeal to Gibbard's apparatus is natural. We should, he claims, "treat the delineation in our world/delineation pairs as a fully specific plan for setting a threshold for gradable adjectives" (MacFarlane 2020c, 648). Since a hyperplan will resolve all possible indecision, it will resolve any indecision regarding where to a draw a cutoff for 'tall' as well (i.e., it will settle the needed delineation). If so (more anon), we can ignore the needed comparison class, we can treat (6) as having a fixed, content-invariant content:

(6b) {<*w*,*h*>: the height of Richard in *w* ≥ the threshold for *tall* determined by *h*}

The "significance of asserting this content" will then depend on the given state of the common ground between speaker and hearer. If, for example, we were already agreed that 5'9" counts as tall then, then my utterance of (6) will be a proposal to add the factual information that Richard is at least 5'9. However, if we already mutually know Richard's height—it's 5'11"—I might, in effect, be making a proposal that, moving forward, anyone as tall as, or taller than, Richard, will count as 'tall' (2020c, 648–649).

MacFarlane convincingly argues that plan-expressivism has many of the virtues of rival accounts without their familiar drawbacks. Notably, the plan-expressivist requires no inscrutable, unknowable boundaries (*contra* the Epistemicist), but can nevertheless recover (a version of) bivalence (2020c, 651).[13] Likewise, the plan-expressivist can give a pleasing response to the Sorites Paradox, as well as an account of our intuitions of 'tolerance' in terms of practical constraints on our plans, that doesn't require the all-too-fortuitous, un-noticed contextual shifts in the applicability of vague expressions (*contra* the Contextualist), or the rejection of classical laws of inference (as is required on Global Supervaluationism).[14] Moreover, MacFarlane makes a plausible case that plan-expressivism can be extended to an account of the contents of vague thoughts

13. In particular, the plan-expressivist can accept a restricted version of bivalence: namely, if *u* is an utterance that says that *p*, then *p* or not *p* (see MacFarlane 2020c, 651).

14. MacFarlane argues that the plan-expressivist should reject the induction step in Sorites reasoning (for example, if someone with n hairs on their head is bald, then so too is anyone with n-1 hairs). Accepting the induction step would be tantamount to accepting a "fully determinate plan" for using the relevant vague concept, "one that leave no room for further refinement or needs" (MacFarlane 2020c, 656). There are, as MacFarlane argues, compelling practical considerations against adopting such fully resolved plans.

For a discussion of Global Supervaluationism and classical rules of inference, see Williamson (1996: ch. 5); for more on Contextualism, see Graff Fara (2000) and Raffman (1996).

in addition to vague language. The application of out concepts in practical and theoretical deliberation are subject to planning, he claims, and, hence, amendable to the plan-expressivist treatment (MacFarlane 2020c, 660–661).

For present purposes, however, I want to focus on the question whether plan-expressivism helps with the problems of communication and content with which we began. According to MacFarlane, his account can be generalized to *all* cases of underspecification (2020c, 649):

> In [2020a], I argued that the problem about vagueness and communication was a special case of the more general problem of felicitous underspecifiction. For gradable adjectives, we will need delineations in addition to worlds, for epistemic modals, we will need information states, for plural definites, we will need group boundaries, and so on. But we can think of all of these things as determined by a hyperplan. After all, a hyperplan is a maximally specific plan, one that encodes plans for everything—including, *a fortiori*, these aspects of language use.

Is this correct? No. In what follows, I'll argue that *even if* we were to assume that the contents of our vagueness-related propositional attitudes are themselves Gibbardian contents of the kind favored by MacFarlane, the problem posed by underspecification remains.

4. Why the Problem Remains

According to some, cases of underspecification motivate "cloudy" accounts of assertion and communication. Such theorists (in their own respective ways) would deny one or more of the theses that MacFarlane calls *Classical Pragmatics*:

Classical Pragmatics

I. The content of an assertion is a (single) proposition.
II. Uptake consists in recognizing the proposition asserted.
III. If the assertion is accepted, its content is added to the conversational common ground.

MacFarlane argues that the extant cloudy pictures on offer have failed to provide a plausible account of "conversational update" in place of the classical picture. We will return to (some of) these critical arguments later, but for now let's focus on MacFarlane's argument that underspecification and vagueness require no departure from Classical Pragmatics. Rather, he claims that "we can keep

the idea that the content of an assertion is a single proposition, and that the proposed update is to add this proposition to the common ground" (2020b, 619), so long as we give up the thesis that we earlier dubbed (Content_a):

(Content_a): The things we assert are propositions.

Once we make this proposed shift to Gibbardian contents, he claims, we can hold onto the idea that there is "a single content that is asserted by the speaker and recognized by the hearer" (2020b, 640).

For the sake of argument, let's follow the plan-expressivist in (a) taking the source of vagueness to be practical indecision in our plans for using our words and deploying our concepts. Moreover, let's provisionally assume (b) that the contents of our vagueness-related thoughts, as well as all of our planning-oriented mental states, have Gibbardian sets of world/hyperplan pairs as their contents, and (c) that in the relevant range of cases, we can identify a unique set of such Gibbardian contents as the common ground between the speaker and hearer. I am dubious of each of these assumptions (especially (b) and (c)), but I don't want to pursue these worries here. For now, I want to show that even if we grant MacFarlane these claims, underspecification remains.

Expressivists in other domains often claim that the content of an utterance is simply "the mental state it *expresses*." If we were happy enough with this slogan without further elaboration, it might seem as if we have conceded the case to the plan-expressivist. After all, we are assuming that for *any* vagueness-related thought in the mind of the speaker there will be a corresponding Gibbardian content and that the common ground of any relevant conversation will be appropriately structured for such a content-in-mind to interact with and "update." If so, aren't we done already?

We are not done, however, and I am certain that MacFarlane would agree. Crucially, expressivists must provide an account of what it means for an utterance to "express" a mental state. Expressing can't, for example, merely be a matter of providing evidence, by your utterance, that you have the relevant mental state. My utterance of the words 'Marfa is closer to Montana than Michigan' is evidence that I believe that there are at least three words in English beginning with the letter 'M', but—elaborate stage setting aside—this fact has nothing to do with the content of any speech act I performed in uttering it, or anything that my audience must recover to understand me.

Fortunately, MacFarlane is much clearer on this issue than his fellow Expressivists in other domains. Both in his argument against rival accounts and in the presentation of his positive view, he assumes that the speaker must (in the relevant assertoric cases*) intend her audience to recognize what she is asserting*. And that much seems clearly correct. "Expressing" a judgment, in the sense

relevant to assertion, at least minimally requires that the speaker produce her utterance *as* evidence of that judgment; evidence that will, by her lights, put her audience in a position to recognize the content thereof (and with that what she is asserting). Hence, even if we grant MacFarlane assumptions (a) and (b) above, we must do more. Having a Gibbardian content in mind is one thing; asserting it another.

There is no difficulty in finding a Gibbardian content to correspond to the context-invariant meaning of (8):[15]

(8) I should smoke.
(8b) {<w,h>: h calls for my smoking in situation w}

It is not impossible, but it is implausible, however, that anyone has just *that* Gibbardian content "in mind" when they make a judgment they'd express by (8). Though (8b) forbids any plan on which I don't smoke *something*, *sometime, somehow*, *in some way*, it leaves open much too much to plausibly be the content of any plan of mine. I simply don't have *any* planning state that would be satisfied should I end up, say, smoking loose leaf notebook paper from a pipe while outdoors in subzero temperatures.[16] Though I am undecided about numerous of the exact features of my plans to smoke, I have already ruled out smoking catnip outside the local coop in the rain, or borrowing my uncle's beloved corncob pipe for the occasion, or … I might not be fully decided on what to do in all contingencies, but presumably *these* plans to smoke are already ruled out. Rather, I simply plan to smoke a cigarette sometime in the near future. I might be undecided on exactly how, or where, or what clothes I'll be wearing at the time, but I am resolute that it should not be a menthol, or poisoned, or covered

15. Let's set aside the intransitive sense of the verb 'smoke', and, with that, any plans that would involve me, or some aspect of my person (e.g., clothing or hair), emitting smoke or vapor.

16. This case suggests that an agent can plan to ϕ, but **not** *have a plan* to ψ, even though it is not possible to ϕ without ψ-ing; our plans are not closed under necessary consequence. Notice that any world in which I successfully execute my plan *to smoke a cigarette sometime soon* will be a world in which I *smoke something, somehow, someway, in some manner*, but (as I suggested above) I do not have *any* plan that would be satisfied should I end up puffing on a corn cob pipe full of pencil shavings. Given that our cognitive attitudes are not closed under necessary logical consequence, it should be unsurprising that our planning states aren't either.

An anonymous referee has suggested that the plan-expressivist might argue that I *do* have a completely unrestricted plan to smoke (period); it's just that this plan fails to capture my overall planning state in its full specificity. In my view, it is unclear what reason there would be for maintaining this, except for the fact that the unrestricted plan is a logical consequence of my own, more specified planning state. The matter is, however, too complex to fully adjudicate here. We might look to a related debate between Braun (2015), Graff Fara (2013), and Lycan (2012) on desire (and ascriptions thereof) for inspiration. But, either way, there are (as argued below) serious difficulties for thinking that what we assert in the relevant range of cases, can be identified with such weak, unrestricted contents (even if we think there are hyper-undecided planning states with such contents).

in butter, or…. Though no one is a fully decided hyper-planner, virtually no one is ever completely undecided regarding all of the contingencies of their plans either.

Let's assume that whatever, exactly, the contours of my undecided plan might be, there will be a unique Gibbardian content C that fits; a set of world/hyperplan pairs that perfectly capture my less than fully undecided state, such as it is. For now, notice two things. First, C can't plausibly be identified with (8b) above, and, though it might be share*able*, it is not likely to be the exact content of anyone's else's judgment. And secondly, there will be indefinitely many other Gibbardian contents that differ only slightly from C, but which rule out (or permit) other plans that are not already excluded (or permitted) by C.

A problem begins to emerge when we turn from judgment to assertion. What exactly am I asserting by uttering (8)? Insofar as the unconstrained Gibbardian content we get from the context-invariant meaning of (8) is an implausible candidate for the content of my (or anyone else's) planning state, it is a poor candidate to be the content of the judgment I am seeking to express by (8). In any normal context, if my audience were merely to have taken me to have asserted my commitment to such a completely unconstrained plan for smoking, it'd be a sign of pathology, rather than understanding my utterance.[17] (For example, imagine my audience responds 'Yeah, you should. Take my lighter and this dried Catnip!')

If, however, I asserted my commitment to some Gibbardian content other than (8b), what exactly is it? It is not plausibly the exact *content* of my planning state, C above, since I cannot reasonably intend my audience to get just that. There are indefinitely many "nearby", but distinct, Gibbardian contents that differ only slightly from C (for example, consider such contents that forbid plans on which I smoke in the next four minutes, rather than five, or where I am wearing a baseball cap rather than my ridiculous velvet smoking hat, or…). The worry is that even taking into account everything we mutually know about each other (e.g. you know I smoke tobacco, that I look eager to smoke soonish, that I am currently wearing my dedicated [and sometimes misunderstood] smoking hat, etc.), how can I reasonably intend my utterance to put you in a position to get that I am expressing my commitment to any particular Gibbardian content (to the exclusion of the nearby competitors) on the basis of my utterance of (8)? If the things we assert are unique Gibbardian contents that our audience must recover in order to understand us, successful communication would seem a momentous achievement in such a case. But it isn't, and that is the worry.

17. To be clear, I'm not claiming that it is impossible to intend (or be recognized as having intended) a completely un-constrained Gibbardian content; rather, I'm claiming that this is just not so in the normal case.

"Wait a minute!," a Plan Expressivist might ask (command?), "If you've already agreed that the common ground itself can be modeled in terms of mutually accepted Gibbardian contents, why can't we just identify what the speaker asserted with the result of 'updating' that common ground with the Gibbardian content corresponding to the context-invariant meanings of the relevant words (and their syntactic arrangement)?" No, not if we are assuming (along with MacFarlane) that asserting that *p* requires intending for your audience to entertain *p*, at least in part, on the basis of your linguistic efforts. In a nutshell, the reason for this is that no facts about your words, the common ground, or what your audience might (even reasonably) take you to be intending can make it the case that you were so intending. Quite generally, whether or not you have the relevant (communicative) intention is a matter of your psychological constitution, then and there; it's simply not the sort of thing that depends on facts about your audience or how they interpret you.[18]

What the speaker asserts depends, in part, on what she intends to convey. But crucially, the facts about what the speaker intends are not themselves metaphysically determined by common ground and the context-invariant meanings of her words. (Of course, the latter are an important part of the *epistemic* considerations that we rely on coming to know what the speaker asserted, but that is different.)

Unsurprisingly, there is no difficulty in finding cases in which the content of the speaker's assertion comes apart from what you would get from intersecting the common ground with the content—Gibbardian, or otherwise—of the context-invariant meaning of her words. For example, suppose that you have happened upon the diary of a co-worker, Tom, and against your better judgment you glanced at some of the entries. Turns out that Tom is seriously considering leaving TechCore as he can no longer stand your colleague, Tiff, and her constant discussion of her numerous dietary restrictions. Moreover, it is also clear that he thinks that everyone else in the office is just as annoyed with Tiff as he is, though he is actually quite mistaken about this. (Your colleagues have, in fact, taken little notice of Tiff.) At the company retreat, you see Tom trying to make his way through a bowl of regrettable chili prepared by the boss's husband, while Tiff is talking at him about her recent decision to no longer eat potatoes. A colleague, noticing the dark, distant look on Tom's face, asks him 'What's going on? You ok?'. Tom raises his face from the (abominable) bowl of chili and utters:

(9) I've had enough.

18. Of course, we craft our utterances to be understood and this requires us to form our communicative intentions in ways that take into account what we think our audiences will be in a position to take away from the evidence we provide them. It is, however, the speaker's beliefs about how she'll be interpreted that are relevant to the formation of her communicative intentions, not how, as a matter of fact, she will be interpreted. See Harris (forthcoming) on communicative planning.

Sure, the chili is terrible and everybody knows it. But, that isn't what Tom is talking about. Though your coworkers—relying on Tom's words and the information in the common ground—don't understand this, you do. Tom is asserting (something to the effect of) that he's had enough of Tiff's comments, or her presence, or maybe even his job at TechCore.[19]

The crucial point for now, however, is that you can't identify what Tom's asserting in this case with what you'd get from simply intersecting the Gibbardian content corresponding to the context-invariant meaning of (9) with the common ground between Tom and his audience; that "update" does not deliver anything that he is intending to convey, much less assert, by uttering (9). And again, the fact that your colleagues might reasonably take him to be asserting, say, *that he has had enough of the awful chili* cannot make it the case that he did so assert. He did not intend, and hence did not assert, any such thing; Tom doesn't care about the quality of the chili, and he would even be fine with having more (so long as it is not in Tiff's presence).[20]

We cannot just "read off" the content of the speaker's assertion from common ground and lexical meaning. Rather, we must look for a content that the speaker plausibly intends to convey and can reasonably expect her audience to recognize in the context. The worry above regarding (8) is that there is no plausible such candidate Gibbardian content. Moreover, this worry threatens to generalize to all cases of felicitous underspecification. For example, return to the case of domain restriction, (3):

(3) Every beer is in the bucket.

Corresponding to the displayed sentence-type, there will be a unique Gibbardian content (3*):[21]

(3*) {<w,h>: The x: Bucket_h(x)(Every y: Beer_h(y)): (y is in x) in w}

19. As in other cases of felicitous underspecification it is hard, if not impossible, to say precisely what was conveyed (more anon).

20. See Goldstein and Kirk-Giannini (2022) for a critical discussion of "Contextology"—the science of conversational update. I agree with these authors that "the laws governing the evolution of context in response to assertion must make essential reference to the private information states of interlocuters" (2022: 3187) and that this thwarts any view that would identify post-update context with the intersection of pre-update context and the content of the speaker's assertion. Cases such as (9) show something different: namely, that you can't read off the content of the speaker's assertion solely from the facts about common ground and the context-invariant meaning of her words.

21. Where, roughly put, we can read the restriction on the quantifiers as 'according to the domain determined by hyperplan h.'

As before, we can reasonably assume that (3*) doesn't plausibly capture the exact contours of anyone's mental state, much less a content the speaker would be seeking to express by uttering (3): again, no speaker uttering (3) will be hyperdecided, nor are they likely to be completely undecided either. Notice that in our earlier case involving quantifier domain restriction, for each of his classically-construed "candidates" (and combinations thereof), there will be a distinct Gibbardian content. For example, corresponding to P_1 (from earlier) there is a Gibbardian content G_1 that forbids (among other things) any plan for using 'every beer' so as not to include *beer bought at the bodega*; corresponding to the disjunction of P_5 v P_6, G_2 forbids any plan for using 'every beer' so as not to include beer to be served at the party or beer in the apartment; and so on. But, whatever the exact contours of the speaker's undecided state, suppose that there is a particular Gibbardian content that fits exactly, call it $G_{special}$.

As best I can see, the same argument that MacFarlane would offer against thinking that the speaker meant any unique classically individuated content can now be repeated for the plan-expressivists favored contents as well. In particular, how can the speaker reasonably expect his utterance of (3) to put his audience in a position to recognize a specific one of these candidate Gibbardian contents—$G_{special}$ included—rather than any of the of the other candidates? Insofar as asserting a content—be it a set of worlds *or* a set of world/hyperplan pairs—requires that the speaker intend her audience to entertain that content on the basis of the evidence her utterance provides in context, we are still left with a multiplicity of candidates. To the extent that the "generality and indifference" of the speaker's meaning-related intentions precluded us from identifying any one classical content as *the* one the speaker asserted, we are equally precluded from finding a unique, asserted Gibbardian content.

We can bolster the case by considering the matter from the hearer's perspective. MacFarlane assumes that successful communicative "uptake" requires that the speaker's audience come to recognize what's asserted. But ask yourself: which Gibbardian content *exactly* must the hearer get in order to understand? Is there any one, unique such content that she must entertain in order to understand the utterance? This seems doubtful. Suppose that the Gibbardian content the speaker has in mind is one consistent with counting as relevant (i) all the beer *bought at the bodega*, (ii) all *the beer for their guests*, and (iii) all the beer *in the apartment*, but rules out any other beer. It seems far too demanding to require that the hearer must arrive at just this Gibbardian content with precisely those contours. Imagine, for example, that the hearer mistakenly thinks that the speaker has a secret stash of beer in the attic that is being saved for a true emergency situation but that the party later that night is no such occasion. As such, imagine she arrives at a different Gibbardian content on the basis of (3): for e.g., she takes the speaker to be proposing a plan that would count as relevant (i) and (ii) above, but

not (iii). As best I can tell, this is no bar on successful communication. The foregoing argument could simply be repeated for any particular contextually, relevant Gibbardian candidate you pick. If no such particular content must be recognized by your audience, then no Gibbardian content is the content of your assertion.

The worry is not limited to cases of quantifier domain restriction. It can arise for virtually any context-sensitive construction, including cases involving vague, comparative adjectives, that MacFarlane focuses on. Even if we spot the plan-expressivist his favored non-factualist construal of the relevant delineation, we still need the context to provide a comparison class for it to operate on. Earlier, we saw MacFarlane claim that—if we ignore the needed comparison class—we can treat your utterance of (6) as having a single, fixed Gibbardian content, (6b):

> (6) Richard is tall.
> (6b) {<w,h)> the height of Richard in w ≥ the threshold for *tall* determined by h}

In a footnote, he adds:

> Of course, different thresholds may be relevant to different uses of 'tall' in a discourse: someone who is tall for an academic may not be tall for an athlete. One way to handle this is to say that a delineation maps the semantic value of a gradable adjective to a threshold, but allow that 'tall' can have a different semantic value when used with a different threshold. (MacFarlane 2020c, 649, fn 11)[22]

The suggestion is plausible, but (again) I am not optimistic that the speaker's intentions and the totality of further facts about the context will always manage to winnow the candidate restrictions on the delineation down to uniqueness. First, note that even ignoring the needed comparison class, (6b) will not correspond to the content of any mental state that a speaker might plausibly have, or be interested in asserting; though not hyper-decided, the speaker is not fully *un*decided either (i.e., plausibly she will be in a state—undecided thought it is—that makes at least some further restrictions on the set of world/hyperplan pairs that are relevant to her judgment that Richard is tall). But second, once we try to accommodate the need for a contextually provided comparison class, we find a

22. It is an interesting question whether *all* literal uses of vague comparative adjectives require a contextually "given" comparison class (or vague cloud, thereof). After all, as MacFarlane emphasizes, sometimes our communicative goal in using, say, 'tall', is to help *establish* a comparison class, e.g., by tying its conversational applicability to the mutually recognized height of this, or that, person or object. Be that as it may, comparative adjectives often *do* seem to require contextually relevant comparison classes and in what follows, we will limit our attention to such cases.

multiplicity of candidates and nothing to choose between them. If so, there will be no unique mutually understood update to the common ground. Suppose we are watching Richard on the field, in the scrum, during the annual Execs versus Consultants rugby match at our corporate retreat. Commenting on Richard's comparatively small stature in comparison to the other hulking players on the field, you say "Richard looks tiny!" to which I reply, by uttering "I know, but Richard *is* tall." To get a unique Gibbardian content, we need to settle the choice of comparison class. But which such classes specifically are compatible with my undecided plans for using 'tall', such as they are? And which can I expect for you recognize? The relevant candidate restrictions would presumably exclude, say, tall *for a rugby player on the field*, but might include many other related, but inequivalent options such as being tall: *for someone from corporate, for an executive, for a consultant, for someone who attended yesterday's mandatory seminar, "Accelerating Innovation: Thinking Inside, and Outside, the Box," for someone at this galvanizing company retreat*, and with some imagination the list can be extended indefinitely.[23] Again, even if we assume that there is a Gibbardian content that perfectly fits my own indecision on this matter, how could I expect for you to get just that merely from uttering 'Richard *is* tall' in this context?[24]

MacFarlane argues that felicitous underspecification raises worries for the standard account of assertion, communication, and the classical conceptions of content. But, if the foregoing discussion is correct, this problem can't be solved merely by giving up standard, truth-conditional contents in favor of Gibbardian ones.

5. Conclusion

There is an urge that perhaps we should resist. Roughly put, this urge is to find something that we might identify as *the* content, or *the* update, in cases of underspecification. Insofar as the arguments of MacFarlane and others put pressure on

23. If challenged ("What do you mean he's *tall*?"), I might happily cough up any one of the specific candidate completions. "Did you mean tall *for an executive*?", I'd say "Sure"; likewise, were you to ask "… tall for *someone from the corporate office*?", and so on for the other candidates. The fact that I would happily "fall back" on any one, or more, of candidates, "suggests that no single such candidate, or set of candidates, perfectly capture [my] communicative intentions' in uttering what I did." (Buchanan 2010: 310).

24. Is it indeterminate which specific Gibbardian content the speaker asserted (though, supervaluating, she did determinately assert something)? No. We have been assuming (with MacFarlane) that assertion requires *intending to communicate*, and in these cases, the speaker *definitely* does *not* so intend. Hence, it is not indeterminate (i.e., not definitely the case that she asserted it and not definitely the case that she did not). See Buchanan & Ostertag (2005) and Schiffer (2020) for this style of argument. Moreover, if the plan-expressivist were at all inclined (at this late point in the dialectic) to go supervaluationist, there is the question of why he did not do so much earlier (and potentially obviating the need for his view to begin with).

the assumption that what we mean/assert in such cases cannot be identified with a classically construed proposition, the solution is not, I think, to try to find some other single entity to offer in its place.[25]

Stepping back, notice that the phenomenon of underspecification arises due to the gap between the context-invariant meaning of the speaker's words and what she meant by using them on a particular occasion. But, of course, this same gap is not limited to cases of literal speech such as those discussed in introducing the phenomenon (§2). For example, consider the following well-worn case from Grice (1975) in which a professor writes a one sentence long letter of recommendation for a student applying for graduate school:

(9) Jones is punctual and has excellent penmanship.

It is clear enough that the professor is seeking to raise a red flag about Jones's candidacy for the program, but what *exactly* is he conversationally implicating in this case? There are many non-equivalent candidates: *that Jones isn't bright*, or *that he is a poor student*, or *that he has no talent in the area for which he's applying*, or *that he shouldn't be admitted*, and so on.

As Heck (2006, 27) nicely puts it, even if the professor "does intend to communicate some such message, I see no reason to suppose that there has to be some *particular* such message he intends to communicate; even if there were, he could not reasonably suppose that his audience could work out which specific message that was." The same gap—between our words, the context, and what exactly we can expect our audience to entertain—gives rise to the multiplicity of non-equivalent candidates for what the speaker indirectly means (i.e., implicates). Here, however, I take it that no one has been seriously tempted to somehow try to reconstruct from these candidates any one unique entity to celebrate as *the thing* the speaker really meant or implicated.[26] What exactly *did* the speaker implicate, then? In such a case, perhaps the best we can do is to simply appreciate the cloudy nature of the speaker's speech act (implicating), noting that the candidates above are clearly "in the cloud," but it might be indeterminate what else (if anything) is included. Though these seem to be the facts as we find them, we somehow manage just fine.

25. I am guilty of this, too, in Buchanan (2010)—though see 368–369, fn 39 for a discussion of whether underspecification might ultimately require us to give up the "entity-sharing" model of communication. See Abreu Zavaleta (2021) for this criticism of my earlier view. Abreu Zavaleta himself attempts to find a unique, *assessment-sensitive*, content by packing the relevant underspecificity into the circumstance of evaluation, rather than into the asserted content itself. In my view, identifying "what" we assert with Abreu Zavaleta's stripped-down, domain-restriction free, contents is no more plausible than simply identifying them with the context-invariant meanings (or characters) of the relevant sentence-type, but showing this will have to wait for another occasion.

26. Nor should we be so tempted in cases of metaphorical speech; see Camp (2006).

Perhaps a similar attitude is needed for the cases of literal speech exhibiting underspecification. Don't try to find *the* content—whatever that might be. Rather, appreciate the cloudy contribution of the speaker such as it is. We are, of course, open to theorize about the shape of the cloud in various ways. One question: what kind of contents comprise the cloud? Do we need structured propositions? Or Gibbardian contents? Both? Or perhaps some other variety of content? A second question: allowing that it will oftentimes be indeterminate—what if anything can we say about, what is in, and what is out, of the cloud? As I argued in my (2010) piece, in cases of literal speech, we might require that anything in the cloud 'fit' the (potentially) gappy, context-invariant meaning of the sentence uttered. Likewise, we could further require that any content "in the cloud" is such that, by the speaker's lights, if her audience were to get just that, she'd be in a position to recognize her downstream perlocutionary aims (e.g., where to get a beer). Even if we give up the search for a single content in such cases, there is still plenty of theorizing to be done.

King (2018; 2021) raises, but explicitly does not try to provide, an answer to the question of "what we assert'' in such cases. This is, in my view, a wise choice, not a missed opportunity. I am inclined to think that in cases of felicitous underspecification there simply isn't any unique thing that we assert, and that King's question doesn't have an answer. Unfortunately, I do not know how to happily square these closing thoughts about speech act content with any standard (broadly Stalnakerian) picture of common ground and conversational update. Whereas MacFarlane (and others) might see this as a reason to reject cloudy accounts of communication and content, I take it, instead, as a clear illustration of one the shortcomings of these highly idealized models. The familiar talk of "conversational update" helps to suggest a picture on which the speaker and her audience must somehow manage to come to mutually recognize that a particular update to the common ground has been made, if communication is to be successful. In my view, cases of felicitous underspecification help to show that successful communication requires no such thing.

Acknowledgments

An earlier version of this paper was presented at a workshop for the Lang-Cog Group, Philosophy Center, at the University of Lisbon (September 2022), as well as a workshop on Language, Planning, and Cooperativity at the CUNY Graduate Center (November 2022). I would like to thank the audiences at both events for their helpful feedback, with special thanks to: Sam Berstler, Elizabeth Camp, Ainhoa Fernández, Laura Delgado, Manuel García-Carpintero, Alex Grzankowski, Peter Hanks, Devon Morse, Stephen Neale, Craige Roberts,

Stephen Schiffer, Henry Schiller, Arthur Sullivan, and Elmar Unnsteinsson. I would also like to thank my colleagues Sinan Dogramaci, Daniel Drucker, Jon Litland, and Mark Sainsbury, as well as Dan Harris and Gary Ostertag, for giving me invaluable feedback on an earlier draft of this piece. And finally, I'd like to thank the two anonymous referees for *Ergo* for their help in improving the paper.

References

Abreu Zavaleta, Martín (2021). Communication and indifference. *Mind and Language*, 36 (1), 81–107.

Bach, Kent (2001). You don't say? *Synthese*, *128*(1), 15–44.

Braun, David (2015). Desiring, desires, and desire ascriptions. *Philosophical Studies*, *172*(1), 141–162.

Buchanan, Ray (2010). A Puzzle about Meaning and Communication. *Noûs*, *44*(2), 340–371.

Buchanan, Ray and Gary Ostertag. (2005). Has the Problem of Incompleteness Rested on a Mistake? *Mind*, *114*(456), 889–913.

Camp, Elizabeth (2006). Metaphor and that certain 'je ne sais quoi'. *Philosophical Studies*, *129*(1), 1–25.

Gibbard, Allan (2003). *Thinking How to Live*. Harvard University Press.

Gillies, Thony and Kai Von Fintel (2011). 'Might' Made Right. In Andy Egan and Brian Weatherson (Eds.), *Epistemic Modality* (108–130). Oxford University Press.

Goldstein, Simon and Cameron Domenico Kirk-Giannani (2022). Contextology. *Philosophical Studies*, *179*(11), 3187–3209.

Graff Fara, Delia (2000). Shifting Sands: An Interest-Relative Theory of Vagueness. *Philosophical Topics*, *28*(1), 45–81.

Graff Fara, Delia (2013). Specifying Desires. *Nous*, *47*(2), 250–272.

Grice, Paul (1975). Logic and Conversation. *Syntax and Semantics: Speech Acts 3*, 41–58.

Harris, Daniel (2022). Semantics without semantic content. *Mind and Language*, *37*(3), 304–328.

Harris, Daniel (forthcoming). Gricean Communication and the Evolution of Language. In B. Geurts and R. Moore (Eds.), *Evolutionary Pragmatics*. Oxford University Press.

Heck, Richard Kimberly (2006). Reason and Language. In C. Macdonald and G. Macdonald (Eds.), *McDowell and His Critics*. Blackwell, 22–45. (Originally published under the name "Richard G. Heck, Jr").

King, Jeffrey C. (2018). Strong Contextual Felicity and Felicitous Underspecification. *Philosophy and Phenomenological Research*, *97*(3), 631–657.

King, Jeffrey C. (2021). *Felicitous Underspecification: Contextually Sensitive Expressions Lacking Unique Semantic Values in Context*. OUP Press.

Lycan, William G. (2012). Desire Considered as a Propositional Attitude. *Philosophical Perspectives*, *26*, *Philosophy of Mind*, 201–215.

MacFarlane, John (2016). Vagueness as Indecision. *Aristotelian Society Supplementary Volume*, *90*(1), 255–283.

MacFarlane, John (2020a). Indeterminacy as Indecision, Lecture I: Vagueness and Communication. *Journal of Philosophy*, *117*(11/12), 593–616.

MacFarlane, John (2020b). Indeterminacy as Indecision, Lecture II: Seeing through the Clouds. *Journal of Philosophy*, *117*(11/12), 617–642.
MacFarlane, John (2020c). Indeterminacy as Indecision, Lecture III: Indeterminacy as Indecision. *Journal of Philosophy*, *117*(11/12), 643–667.
Raffman, Diana (1996). Vagueness and context-relativity. *Philosophical Studies*, 81, 175–192.
Schiffer, Stephen (2020). Expression-Meaning and Vagueness. In A. Sullivan (Ed.), *Sensations, Thoughts, Language: Essays in Honor of Brian Loar*. Routledge.
Williamson, Timothy (1996). *Vagueness*. Routledge.