

Communication and Variance*

Martín Abreu Zavaleta[†]

1 Introduction

According to standard assumptions in semantics, (a) users of a language have implicit beliefs about the truth-conditional contents of assertoric utterances in that language, and (b) they often have the same such beliefs.¹ For example, it is assumed that if Anna and John are competent English speakers and the former utters ‘grass is green’ in conversation with John, they will both implicitly believe that Anna’s utterance has the truth-conditional content that grass is green; that if Anna utters ‘I like apples’ in conversation with John, both of them will believe that Anna’s utterance has the truth-conditional content that Anna likes apples; etc.

These assumptions play an important role in an intuitively compelling and fairly standard picture of communication. According to that picture, successful communication through an assertoric utterance requires knowing what the speaker intended to communicate through that utterance, and is typically achieved thanks to our shared beliefs about the truth-conditional content of the utterances we make.

*Thanks to Ben Holguín, Ian Grubb, Rose Flinn, Tienmu Ma, Chris Scambler, and two anonymous reviewers for very helpful comments. For helpful discussion, thanks to audiences at Institut d’Études Avancées de Paris, New York University, The New School for Social Research, and Università degli Studi di Torino. Special thanks to Cian Dorr, James Pryor, Stephen Schiffer, and Erica Shumener for their insightful feedback on several drafts of this paper.

[†]Please cite as Abreu Zavaleta, Martín

¹These theses follow from the following standard assumptions. First, that knowledge entails belief. Second, that competent language users know the meaning of utterances in the languages they are competent in. Third, that knowing the meaning of an utterance requires knowing its truth-conditions. For a version of this last assumption, see for instance Heim and Kratzer (1998), who start their famous textbook by stating “To know the meaning of a sentence is to know its truth-conditions” (p.1). Portner (2005) also takes knowledge of truth-conditions to be the starting point for semantics: “The knowledge of meaning involves (at least) knowledge of the conditions under which a sentence is true, and those under which it’s false” (p. 13). Larson and Segal (1995) motivate similar assumptions in their initial discussion of the relation between meaning and truth (pp. 5–7). Though Heim and Kratzer talk about knowledge of a sentence’s (as opposed to an utterance’s) truth-conditions, they do so only for the sake of simplicity, since they are ignoring the influence of context on a sentence’s semantic content. Some authors (e.g. Fodor (1983), Pettit (2002)) think that the mental states underlying linguistic competence aren’t beliefs or knowledge, but something else, such as subpersonal mental states or states of a modular input-output system. Versions of the arguments I will develop below apply to those views as well, but I will leave this generalization for future work. Thanks to an anonymous reviewer for helpful discussion.

For example, suppose John wants to know what Carla's hobbies are, and Anna wants to inform him that Carla likes to run. According to the standard picture of communication, Anna will make an utterance with the truth-conditional content that Carla likes to run. Normally, according to the standard picture, if communication is successful John will come to know that Anna wanted to inform him that Carla likes to run. He will come to know this on the basis of two beliefs: first, that Anna's utterance has the truth-conditional content that Carla likes to run, and second, that Anna believes that her utterance has the truth-conditional content that Carla likes to run.^{2, 3}

This paper argues that ordinary language users rarely agree on what the truth-conditional content of an utterance is. More precisely, it argues for *Variance*:

Variance: Nearly every utterance is such that there is no proposition that more than one language user believes to be that utterance's truth-conditional content.

Provided that successful communication through literal assertoric utterances is as common as it intuitively seems, if the participants in a linguistic interaction seldom have the same beliefs about the truth-conditional content of the utterances made in the course of the interaction, successful communication cannot be typically achieved due to shared beliefs about those utterances' truth-conditional content. Nor can successful communication typically require that we know exactly what the speaker intended to communicate through her utterance—if, for example, John doesn't believe that Anna's utterance of 'Carla likes to run' has the truth-conditional content P, he will have little reason to believe that Anna intended to communicate P (as opposed to some other proposition, Q) to him through her utterance of 'Carla likes to run'.⁴

² Without intending to do exegesis here, it would be fair to attribute a version of the standard picture of communication to philosophers in the tradition started by Grice (1989a,b)—a tradition that includes Strawson (1970, 1964), Schiffer (1972), Bach and Harnish (1979) and, to some extent, Stalnaker (1999c). According to philosophers in that tradition, in making an assertoric utterance a speaker means a proposition (or propositions), and the audience understands the utterance only if she recognizes the proposition(s) the speaker meant. If at least one of the propositions the speaker meant (in Gricean terms, the proposition the speaker *said*) determines the truth-conditions of the uttered sentence, we should expect that if the audience understands the speaker's utterance, then speaker and audience believe the uttered sentence to have the same truth-conditions. Heck (2002, pp. 6–8), Evans (1982, p.22), and Dummett (2010) have attributed versions of the standard picture of communication to Frege. See also Portner (2005, pp.21-2) for an endorsement of the picture sketched in this paragraph.

³ Defenders of the standard picture normally add further requirements for successful communication. For example, they may claim that successful communication between Anna and John normally requires not only that they both know that 'Carla likes to run' as Anna used it has the truth-conditional content that Carla likes to run, but also that they both know that the other knows this, that they know that they know it, and so on. For the purposes of the present discussion, we can do without such additions to the standard picture.

⁴ Authors like Stalnaker (1999a) sometimes remark that communication will be possible in some cases even if speaker and

I will elaborate on Variance’s consequences for communication and other linguistic phenomena later on. For the time being, here is the structure of the paper. I start by arguing for Variance and stating some clarifications about it (section 2). Then I examine some strategies for resisting the case for Variance appealing to metaphysical and psychological naturalness, to the notion of common ground, and to social externalism, respectively (section 3). I argue that those strategies fail. Towards the end of the paper I illustrate what is at stake in explaining the nature of communication by characterizing three independent ways in which a conversation can be communication-like (section 4). As I argue, seemingly natural strategies for giving a unified explanation of those ways for conversations to be communication-like—appealing, respectively, to diagonalization and metalinguistic negotiation—are unsuccessful (section 5). I conclude with a brief discussion of Variance’s consequences for other debates in philosophy of language (section 6).

2 The case for Variance

Variance is the thesis that nearly every assertoric utterance is such that there is no proposition that more than one language user believes to be its truth-conditional content. Before presenting the case for Variance, it is worth clarifying the notion of truth-conditional content. Inspired by Lewis (1975), we can define truth-conditional content as follows (where U is an utterance and P is a proposition):

TC-content: U has the truth-conditional content P just in case: the unique sentence S that U is an utterance of, the unique language L that S is in, and the unique context c that U occupies are such that, necessarily, S is true as used in c in L if and only if P is true.⁵

For the sake of simplicity, we can understand context in the purely formal sense due to Kaplan (1989); i.e. as a sequence of values for the various parameters a sentence’s content—the proposition it expresses—could depend on. An utterance’s truth-conditional content is the proposition resulting from providing the linguistic meaning (treated as a function from contexts formally understood to propositions) of the

audience don’t agree over the truth-conditional content of the utterances used to communicate. This suggests that perhaps the standard view should be taken as a view of *ideal* communication. Unfortunately, this observation alone will not take us very far, since, as we will see in section 5, it is not obvious how to extend the standard view to non-ideal cases. Furthermore, what the truth of Variance suggests is that non-ideal cases of communication are so widespread that offering a view of non-ideal communication is far more pressing than Stalnaker seems to acknowledge. Thanks for an anonymous reviewer for discussion.

⁵ Sentences here should be thought of as disambiguated sentences.

sentence that utterance is an utterance of with the context the utterance occupies.

Given the present definition of truth-conditional content, and on pain of contradiction, if U is an utterance and the propositions P and Q are not necessarily equivalent, at most one of P and Q is U's truth-conditional content. Equivalently, if we think of propositions as sets of possible worlds, at most one possible-worlds proposition is U's truth-conditional content.

With this in mind, here is the main point in favor of Variance. Nearly every assertoric utterance is such that, for any language user, there are enormously many extremely similar and equally eligible truth-conditional contents that language user could easily have believed the utterance to have. Given the large number of such extremely similar and equally eligible truth-conditional content candidates each of those utterances has, it would be extremely unlikely for there to be a proposition that more than one language user believes to be that utterance's truth-conditional content. Thus, absent further explanation, there is no such proposition.⁶

For instance, consider gradable adjectives. According to a popular line of research, the positive form of a gradable adjective (e.g. 'heavy', as opposed to 'heavier') expresses a relation between the degree to which an object possesses the gradable property measured by the predicate and a contextually determined standard of comparison (See Cresswell 1977, Heim 2000, Kennedy and McNally 2005, Kennedy 2007). For example, the predicate 'heavy' is taken to express the property of having a degree of heaviness greater than or equal to a contextually determined standard of tallness; the predicate 'relevant' is taken to express the property of having a degree of relevance greater than or equal to a contextually determined standard of relevance; etc. More generally, given a gradable adjective, we can think of the relevant standard as a cutoff point along the scale the predicate measures, such that objects which have the property to a degree at least as great as the cutoff count as instances of the predicate.

Now, for any gradable adjective F and any utterance U involving F, there are many slightly different and equally natural contexts which any language user could easily have believed U to occupy. Each of those slightly different contexts will determine a slightly different standard of F-ness and, accordingly,

⁶ See Abreu Zavaleta (2018). See also Dorr and Hawthorne (2014) for a related argument to the effect that, if the propositions expressed by most utterances in ordinary language depend on microphysical facts, they depend very sensitively on those facts. See also Schiffer (1981a) for a related argument to the effect that the content of an utterance can't include reference to specific modes of presentation, and Buchanan (2010) for a related argument against Gricean accounts of speaker meaning. I intend to discuss the differences between some of these arguments and the one I present here in future work.

a different truth-conditional content for U.⁷ Given the great number of equally eligible contexts any language user could have taken U to occupy, it would be extremely unlikely for there to be one which more than one language user takes U to occupy. And, since each of those contexts determines a different truth-conditional content for U, it would be extremely unlikely for there to be a proposition which more than one language user takes to be U's truth-conditional content.

For example, suppose Anna points to a certain box she has asked John to move and utters

(1) That box is heavy.

There are many extremely similar and equally eligible contexts of utterance (formally understood) which Anna and John could easily have taken Anna's utterance of (1) to occupy, all of which determine different standards of heaviness. For example, let's suppose that degrees of heaviness come in a scale from 0 to 1. Then one such context will be one in which the box Anna pointed to counts as "heavy" just in case it is heavy to degree .7 or higher, another will be one in which it counts as "heavy" just in case it is heavy to degree .71 or higher, etc. Given the large number of those contexts, it is very unlikely that Anna and John take Anna's utterance of (1) to occupy exactly the same one. Accordingly, it is just as unlikely that there is a proposition Anna and John both take to be the truth-conditional content of Anna's utterance of (1).

Considerations like this hold for nearly every context-dependent expression, including location adverbs like 'here' or 'there'. Nearly every utterance involving a context-dependent expression is such that there are many equally eligible contexts any language user could easily have taken the utterance to occupy, all of which determine different values for the parameters on which the utterance's truth-conditional content depends—hence, all of which determine different truth-conditional contents for that utterance.

Note, however, that the kind of considerations I just presented do not apply only to utterances

⁷ Kennedy (2007, section 3) denies that this is true of so-called "absolute" gradable adjectives, which he takes to come with fixed standards. For instance, according to Kennedy, in order for an object to fall in the extension of 'impure' it suffices that it has *some* minimal degree of impurity which remains constant throughout all contexts; in order for an object to fall under the extension of 'straight', it must be completely straight; etc. I am skeptical of Kennedy's claims: if my only purpose is to drink water that won't poison me, I will be willing to take an utterance of 'that water is pure' to be true even if the water in question has one milligram of sodium, but I will be less willing to take a similar utterance to be true in the context of a delicate chemical experiment. Of course, it might be that the range of acceptable standards for absolute gradable adjectives is more constrained than the range of acceptable standards for tallness or expensiveness, but there are reasons to think that the standards for absolute gradable adjectives can change with context nevertheless. I hope to discuss Kennedy's arguments elsewhere.

involving context-dependent expressions. For example, the verb ‘to run’ is not usually taken to be context-dependent, but similar observations are true of it. For instance, suppose that Anna utters

(2) Carla likes to run.

Though the verb ‘to run’ is not usually taken to be context-sensitive, there are many similar, equally eligible truth-conditional contents any language user could easily have taken Anna’s utterance to have. For instance, different people may disagree about the speed at which someone must move in order for her activity to be called ‘running’ (as opposed to ‘jogging’), and this will result in their having different beliefs about the truth-conditional content of Anna’s utterance.⁸ Furthermore, people may even disagree on whether an activity is to be called ‘running’ solely on the basis of pace; they may think, for example, that whether an activity is to be called ‘running’ depends on the putative runner’s top speed, age, fitness level, and so on.

Because of the very large number of equally natural ways of drawing the line between activities that are called ‘running’ and activities that are not, there are many equally natural truth-conditional contents any language user could easily have believed Anna’s utterance to have. Absent further explanation, it would be extremely unlikely for any two language users to believe that utterance to have exactly the same truth-conditional content.

More generally, observations of the kind I just presented hold for (but are not limited to) utterances involving predicates whose application depends on the properties an object has along one or more sufficiently fine-grained dimensions.⁹ In order to see this, let F be a predicate of that kind and j be a proper name. For any language user, there will be a very large number of extremely similar and equally eligible properties such that that language user could easily have believed that that utterance has the truth-conditional content that j (the object denoted by j) has that property—each of those properties will correspond to a slightly different cutoff point along the dimensions on which F ’s application depends. Given the huge number of such properties, it would be extremely unlikely for there to be one such that

⁸ See JBiz et al. (2010) for an example of this kind of disagreement about the definition of the verb ‘to run’.

⁹ Continuous dimensions are the prime example of a fine-grained dimension, but even dimensions with a finite number of points can satisfy this requirement. For instance, even if the application of F depends on an object’s properties along a dimension with only 10,000 points, the present argument would go through. All the argument requires is that there are enough points along those dimensions for there to be a large number of very similar and equally natural properties (all with a plausible claim to be the one expressed by F), each corresponding to slightly different cutoff points along those dimensions. Predicates whose application depends on an object’s properties along continuous dimensions (e.g. gradable adjectives) are prime examples of this kind of predicate, but they are not the only such examples.

more than one language user believes that U has the truth-conditional content that j has that property.

For example, suppose that Anna believes that F expresses the property corresponding to the cutoff point x along one of the dimensions relevant to F 's application; she could just as easily have believed that F expresses the property corresponding to the slightly different cutoff point x' , that it expresses the property corresponding to the slightly different cutoff point x'' , etc., and the same is true of every language user—given the extreme similarity between x , x' , x'' , and so on, even small differences in people's evidence, their attunement to the evidence, and so on, will make a difference as to which of those points they takes to be the cutoff relevant to F 's application. Now, take an arbitrary utterance U of $\lceil j \text{ is } F \rceil$. Given the huge number of properties any language user could easily have taken F to express, it would be extremely unlikely for there to be a property such that more than one language user believes that U has the truth-conditional content that j has that property. The same applies, *mutatis mutandis*, to terms that express relations rather than properties.

I will discuss three potential ways of resisting the case for Variance in a moment, but for the time being it's important to state a clarification. I have written as if the truth-conditional contents language users believe utterances to have divide the totality of logical space between the possible worlds in which the utterance is true and those in which it is not. In my discussion of the example, I assumed that beliefs about an utterance's truth-conditional content can be characterized through the use of precise vocabulary (e.g. 'heavy at least to degree 0.7'), and in the generalization of that kind of example I assumed that the properties that could figure in somebody's beliefs about an utterance's truth-conditional content made sharp cutoff points along some relevant dimension. Call this the *precision assumption*.

Some people may deny that the truth-conditional contents ordinary speakers believe utterances to have divide logical space in the way the precision assumption requires, and question the case for Variance on that basis. According to those people, the truth-conditional contents people ordinarily believe utterances to have are instead vague.

There is no generally accepted treatment of vague propositions, but the appeal to vague propositions does not make a difference to the case for Variance. In particular, it is compatible with the case for Variance that in the schema *A believes that U has the truth-conditional content P* (where 'A' refers to an agent and 'U' refers to an utterance), P be a vague proposition. All the case for Variance requires is that for most utterances there are very many equally natural *vague* truth-conditional contents any language

user could easily have believed each of those utterances to have, given which it would be extremely unlikely for any two language users—including the participants in the conversation in which the utterance is made—to believe any of those utterances to have the exact same (vague) truth-conditional content.^{10,11}

3 The case against Variance

In stating the case for Variance I have assumed that ordinary language users' beliefs about the truth-conditional content of a given utterance are not perfectly correlated—i.e. that those beliefs can come apart. This lack of perfect correlation, together with the vast number of truth-conditional contents any language user could easily have believed an utterance to have, makes it very unlikely for there to be a proposition which more than one language user believes to be a given utterance's truth-conditional content.

To get a better idea of this lack of perfect correlation in people's beliefs, consider the following example. Suppose there are two normal mercury thermometers submerged in the same liquid, and that that liquid has a temperature of around fifty degrees. There is a huge number of readings each of the two thermometers could easily have produced, corresponding to the points in the line between (say) the marks signaling forty-nine and fifty-one degrees. Given the huge number of readings the two

¹⁰ Taking for granted a treatment of vague propositions along roughly supervaluationist lines, here is how the case for Variance would apply if truth-conditional contents are vague. Let's say that a vague proposition is a set of classical propositions. Necessarily, a vague proposition P is true just in case every proposition in P is true, false just in case every proposition in P is false, and indeterminate otherwise. Then we can see that, in most cases, there will be many vague propositions one could easily have taken to be the truth-conditional content of a given utterance. For instance, suppose that Anna believes that her utterance of (1) has truth-conditional content P, where P is the vague proposition containing exactly the following classical propositions: that the box she pointed at is heavy at least to degree 0.6, that it is heavy at least to degree 0.61, and so on, up to the proposition that the box Anna pointed at is heavy at least to degree 0.7. There are many vague propositions extremely similar to this one which John could easily have taken to be the truth-conditional content of Anna's utterance. For example, the vague proposition containing exactly the following classical propositions: that the box Anna pointed at is heavy at least to degree 0.61, that it is heavy at least to degree 0.62, and so on up to the proposition that the box Anna pointed at is heavy at least to degree 0.72. Given the huge number of these vague propositions, and that none of them seems more natural than the rest, it would be extremely unlikely for any two language users to believe exactly the same one to be the truth-conditional content of Anna's utterance. Thanks to an anonymous reviewer for discussion.

¹¹ Note that appealing to vagueness in the identity conditions of vague truth-conditional contents will not take defenders of the standard picture of communication very far. Strictly speaking, all we need in order to reject the standard picture is that, for most utterances, there is a large enough number of equally natural (or otherwise eligible) yet *not definitely identical* truth-conditional contents any language user could easily have believed each of those utterances to have. Given the huge number of not-definitely-identical truth-conditional content candidates each of those utterances has, it would be extremely unlikely for any two language users to *definitely* have the same beliefs about those utterances' truth-conditional contents. If this is true, then the standard view of communication would predict that definite cases of successful communication are extremely rare.

thermometers could easily have produced, they are unlikely to produce exactly the same one. Given the huge number of extremely similar candidate readings, even small differences in calibration, the exact temperature of the liquid immediately surrounding each thermometer, the amount of mercury in each thermometer, and so on, are likely to produce a difference in the exact reading the thermometers produce (or even in the reading they produce up to, say, a hundredth of a degree). Thus, though the two thermometers will often give very similar readings, they are very unlikely to give the exact same one.

In stating the case for Variance, I have assumed that any two people's beliefs about an utterance's truth-conditional content are as imperfectly correlated as the readings of the two thermometers from the example. Given the large number of plausible and extremely similar truth-conditional contents any two language users could easily have believed an utterance to have, even small differences in those people's evidence, those people's attunement to the evidence, and the information each of those people deem relevant to determining the utterance's truth-conditional content, among other factors, are likely to produce differences in the exact truth-conditional content those two people believe an utterance to have. Thus, if people in fact have beliefs about an utterance's truth-conditional content, they will often have similar such beliefs, but they are unlikely to have the exact same one.

Opponents of Variance face the challenge of explaining how ordinary people's beliefs about an utterance's conditional content could be so closely correlated that the fact that one of them has a certain belief about an utterance's truth-conditional content makes it very likely that other language users have that exact same belief as well. According to what I take to be the most promising strategy for answering this challenge, such close correlation arises because the facts (partly) responsible for one language user's beliefs about an utterance's truth-conditional content are also (partly) responsible for other language users' beliefs about that utterance's truth-conditional content.¹² Call this the *dependence strategy*.

For instance, suppose Anna utters 'Carla is tall' in conversation with John, and she believes that her utterance has the truth-conditional content that Carla is at least 1.8m tall. According to the present

¹² Facts may be responsible for other facts in a constitutive or a causal sense (perhaps among others). The strategies I will consider here are implemented in terms of constitutive dependence, but it shouldn't be difficult to see that the same remarks apply to implementations using the causal notion of dependence. As we will see in the discussion below, the reason the present strategy fails is that, given the large numbers of different propositions language users could easily have believed to be an utterance's truth-conditional content, even small differences in the facts (causally or constitutively) responsible for someone's beliefs about an utterance's truth-conditional content will make a difference in those people's beliefs about that utterance's truth-conditional content.

strategy, some of the facts that contribute to determining *Anna's* beliefs about her utterance's truth-conditional content—say, that a certain standard for tallness is more salient than others given the assumptions Anna and John make for the purposes of the conversation—also contribute to determining *John's* beliefs about that utterance's truth-conditional content. Thus, it is not a mere coincidence that Anna and John both believe that the truth-conditional content of Anna's utterance is the proposition that Carla is at least 1.8m tall: according to the present strategy, there is a single set of facts that is partly responsible for the fact that each of Anna and John have that belief. According to this strategy, because in all linguistic interactions there are certain facts which contribute to determining *all* of the participants' beliefs about the truth-conditional content of the utterances made in the course of the interaction, it is not a mere coincidence that those participants have the same beliefs about those utterances' truth-conditional content.

The rest of this section examines three versions of the dependence strategy and argues that they fail. I will focus on versions of the strategy attempting to explain how, at the very least, the participants in a conversation in which a given utterance is made must have the same beliefs about that utterance's truth-conditional content. By appealing to naturalness, common ground, and social externalism, respectively, those versions attempt to explain how it is that the participants in typical conversations can come to have the same beliefs about the truth-conditional content of the utterances they make.

3.1 Naturalness

Some people may believe that certain propositions are more natural than others in a metaphysical or psychological sense. According to them, one potential reason why someone may believe that a given proposition P—as opposed to, say, Q—is the truth-conditional content of a given utterance U is that P is a more natural proposition than the alternatives, or that P is a more natural candidate for the truth-conditional content of U. According to this line of thought, it is because of the greater naturalness of P that any two people who participate in the same conversation are likely to believe that P is U's truth-conditional content. The idea is that, for nearly every utterance, there will be a proposition which is more natural candidate for being that utterance's truth-conditional content than the rest; on account of that naturalness, the thought goes, language users will tend to agree that that more natural proposition is that utterance's truth-conditional content.

There is some plausibility to the idea that, from a metaphysical or a psychological perspective, some propositions are more natural candidates to for being the truth-conditional content of certain utterances than others. However, I am skeptical that standard or otherwise readily available accounts of psychological or metaphysical naturalness can be extended so as to make the dependence strategy plausible. I am skeptical, for example, that any one of the propositions

(3) Carla likes to move on her feet at a pace of ten minutes per mile

(4) Carla likes to move on her feet at a pace of ten minutes and four seconds per mile

(5) Carla likes to move on her feet at a pace of nine minutes and fifty-eight seconds per mile

is a more natural candidate for being the truth-conditional content of Anna's utterance of 'Carla likes to run'.

Let's start with metaphysical naturalness. Some philosophers have claimed that some properties are more natural than others, and that such naturalness makes them easier to latch on to as the contents of our concepts and predicates.¹³ If this is true, it seems, certain truth-conditional contents must be more easily assigned to an utterance than others on account of their naturalness. For example, because the property of being water is more natural than the property of being something that is water before January 20, 2185 or transparent thereafter, it is more likely that a speaker will believe that an utterance of 'there is water in the glass' has the truth-conditional content

(6) there is water in the glass

than that it has the truth-conditional content

(7) there is something in the glass that is water before January 20, 2185 or transparent after January 20, 2185.

Appeal to metaphysical naturalness may help explain why there is rarely disagreement over the truth-conditional content of utterances that involve only terms expressing very natural properties, such

¹³ See e.g. Lewis (1983), Sider (2011), among others.

as ‘electron’, ‘quark’, and other terms used by the most fundamental sciences. But metaphysical naturalness won’t help with utterances involving terms that do not express very natural properties.¹⁴ It is very unlikely, for example, that any one of (3), (4) and (5) is metaphysically more natural than the others. So an appeal to metaphysical naturalness does not explain why there should be a proposition ordinary language users will take to be the truth-conditional content of Anna’s utterance of ‘Carla likes to run’. More generally, an appeal to metaphysical naturalness does not explain why, for a high number of utterances, there should be a proposition various language users take to be that utterance’s truth-conditional content.

Though psychologists and philosophers don’t talk about psychological naturalness, it is not too difficult to make sense of the idea that some propositions may be more natural than others in a psychological (rather than metaphysical) sense. In particular, we can illustrate the notion of psychological naturalness by considering some of the cognitive biases that developmental psychologists postulate to explain the acquisition of common nouns. Here are some examples of such biases:

- **Whole-object bias.** According to Markman (1990), language learners assume that new nouns denote whole objects, rather than any of their parts.
- **Shape bias.** According to Landau et al. (1988, 1998), in determining the meaning of a noun, language learners assign a larger weight to similarity in shape than to similarity along other perceptual dimensions, such as color or texture.
- **Taxonomical bias.** According to Markman (1990), Markman and Hutchinson (1984), language learners assume that nouns denote objects of a given kind (e.g. cats, dogs) rather than thematic relations between objects (e.g. cause, recipient, beneficiary, agent, experiencer, etc.).

For example, suppose a child has never heard the word ‘dog’ and is presented with a picture in which a brown dog is chewing a bone. Pointing to the dog, roughly in the direction of its head, an experimenter tells the child ‘Look! There is a dog’. Given the whole-object bias, the child will be more likely to believe that that utterance’s truth-conditional content is

¹⁴Terms that do not express especially natural properties include some so-called “natural-kind terms”. For example, it is unlikely that there is a most natural way of drawing the line between members and non-members of a given species, so natural-kind terms like “dog” or “cat” are not natural in the metaphysical sense required by the present strategy.

(8) there is a dog

than that it is

(9) there is a dog's head.

Given the shape bias, the child will also be more likely to believe that the utterance's truth-conditional content is

(10) there is a dog-shaped thing

than that it is

(11) there is a brown thing.

Finally, given the taxonomical bias, the child will be more likely to believe that the utterance's truth-conditional content is (8) than that it is

(12) there is a pair of something and a bone such that the former
chews the latter.

Insofar as, due to the biases I described above, language learners are more likely to believe that the utterance of 'there is a dog' has the truth-conditional content (8) or (10) than that it has the truth-conditional content (9), (11), or (12), (8) and (10) are psychologically more natural than (9)–(12).

I am skeptical that the resulting notion of psychological naturalness can be used to single out one of the candidate truth-conditions in the cases that concern us as psychologically more natural than the rest. The whole-object, shape, and taxonomical biases don't distinguish between the different truth-conditional contents two people could believe an utterance of 'Carla likes to run' to have. For example, the propositions (3)–(5) don't differ from one another in whether they categorize objects by shape, in being about parts of objects rather than whole objects, or in whether they attribute an object the property of belonging in a certain class.

This is not to say that new cognitive biases can't be discovered. It may be that new psychological research leads to postulating new cognitive biases which makes one of (3)–(5) (among others) more natural than others. But given that the candidate truth-conditional contents for many utterances can differ from one another along various dimensions, and that there is an enormous number of candidates

differing from one another along each of those dimensions, it is unlikely that those new psychological biases will make one of them psychologically more natural than the rest—as the present strategy would require. In order for the strategy to succeed, there should be cognitive biases capable of discriminating between the candidate truth-conditional contents along every dimension of variability, and they should be strong enough to deem one of those candidates more natural than the others along every dimension. I am skeptical that cognitive biases of such discriminating power need to be postulated in order to explain any general psychological phenomenon.

To summarize, if the present version of the dependence strategy is to succeed, we need to explain why any one of the truth-conditional content candidates for most utterances would be more natural than the rest. As I have argued, it seems unlikely that readily available accounts of metaphysical or psychological naturalness can help provide those explanations. Absent such explanations, the present strategy offers no reason to reject the case for *Variance*.

3.2 Common ground

According to a popular picture of assertion, an assertion is a proposal to update the common ground—the set of propositions that all the participants in a conversation presuppose for the purposes of the conversation, that they all presuppose that they presuppose, etc.¹⁵ That common ground is itself supposed to determine the content of an assertion,¹⁶ and it is supposed to be such that if a proposition is in the common ground, the participants in the conversation know that it is so or, at the very least, presuppose it.¹⁷

Some people may think that this picture of assertion is independently plausible and use it to formulate a version of the dependence strategy. According to that version of the strategy, the common ground

¹⁵ The idea that in every conversation there is a set of propositions all the participants presuppose, that those participants presuppose that they presuppose, and so on, can be traced back at least to Schiffer (1972), is clearly present throughout Lewis (1979), and drives much of Stalnaker's discussion in his (1999c) and subsequent work.

¹⁶ As Stalnaker (2009) puts it,

[A]n assertion is, in effect, a proposal to shrink the context set [the set of possible worlds compatible with every proposition in the common ground] with the content of the assertion. But the context set represents the information that is presumed to be available for the interpretation of the speech act, and if the asserted content is not determined by this information, then the addressee will not be in a position to tell what is being proposed. (p.407)

¹⁷ See Hawthorne and Magidor (2009) for objections to this picture of assertion in connection to this kind of transparency assumption, Stalnaker (2009) for a defense of the transparency assumption, and Hawthorne and Magidor (2010) for a reply.

determines the truth-conditional content of the assertoric utterances made in the course of a conversation. Since all the participants in the conversation know what is in the common ground, they will all agree on the truth-conditional contents of the utterances made in the course of the conversation.

I will focus on the case of context-dependent sentences. According to the present strategy, in conjunction with the linguistic meaning of a context-dependent sentence, the common ground will determine the truth-conditional content of an utterance of that sentence. For example, suppose that Anna utters ‘I like apples’ in a conversation with John, and believes that her utterance’s truth-conditional content is the proposition that she (Anna) likes apples. According to the present strategy, Anna’s belief is partly determined by two factors: first, her belief that the English word ‘I’ denotes whoever utters it; second, her knowledge that she and John both presuppose that in the context of their conversation it was she who uttered ‘I’, that she and John both presuppose that the other presupposes it, that they both presuppose that they both presuppose that they presuppose it, etc. Since both Anna and John know that ‘I’ refers to whoever utters it, and that in the context of their conversation it is common ground that Anna uttered it, both Anna and John will believe that Anna’s utterance of ‘I like apples’ has the truth-conditional content that Anna likes apples.¹⁸

More generally, the thought goes, a person’s beliefs pertaining to the truth-conditional content of an utterance of a context-dependent sentence are determined by her knowledge of the common ground and her knowledge of those sentences’ linguistic meaning. According to this strategy, since in most conversations the participants know the linguistic meaning of the uttered sentences and what is in the common ground, in most cases they will come to have the same beliefs about the truth-conditional content of the utterances of context-dependent sentences made in the course of the conversation.¹⁹

I don’t think this strategy can be used to resist the case for Variance. One potential worry concerns whether the participants in a conversation ordinarily have the same beliefs about the linguistic meaning of the context-dependent sentences they utter, but I will set that worry aside throughout my discus-

¹⁸ See Stalnaker (1999b) for elaboration of this picture in connection with the standard Kaplanian treatment of context-dependence. According to Stalnaker, “Since the relevant contextual parameters must be available, and presupposed to be available, they will be incorporated into the speaker’s presuppositions, and so will be represented by the set of possible situations that constitute the context set.”(1999b, p.10)

¹⁹ It is not straightforward to develop a version of the present strategy that addresses the argument for Variance in cases of context *independence*. So even if the present strategy succeeds (which, as I will soon argue, it does not), it would not count against the observations pertaining to context-*independent* expressions presented in section 2.2.

sion.²⁰ What I want to point out is that, for most context-dependent sentences, *if* the participants in a conversation make presuppositions precise enough to determine truth-conditional contents for utterances of those sentences, it is very unlikely that they make the same such presuppositions. In the same way in which there are many equally natural candidates for the truth-conditional content of most utterances, there are many equally natural (or otherwise viable) sets of propositions any person could have presupposed for the purposes of a given conversation, each of which determines different values for the parameters on which the truth-conditional content of an utterance of a context-dependent utterance depends. Given the enormous number of such sets, it is extremely unlikely that any two participants in a conversation will presuppose the propositions in exactly the same such sets.

Take for example the case of gradable adjectives. As I said above, according to some of the most promising semantic analyses of gradable adjectives, the truth-conditions of utterances involving such adjectives depend on contextually determined standards of comparison. Ordinary conversations rarely include explicit remarks about standards of comparison, so if the participants in the conversation are to reach the same conclusion about which standard of comparison should determine the truth-conditional content of a given utterance, they must do so on the basis of perceptually available evidence. However, since there are small variations in how different people perceive their surroundings, it will be rare for any two people to reach the exact same conclusion about which standard should determine the utterance's truth-conditional content.

For example, suppose Anna and John are waiting for Carla and see her walking towards them from afar. Looking at Carla, Anna starts a conversation with John by uttering 'Carla is tall. I hadn't noticed it before'. Since the conversation started with Anna's utterance, if the common ground has information about the relevant standard for tallness, it can't be in virtue of any explicit remarks about that standard. Suppose further that neither Anna nor John know what the other knows (or what the other knows about what the other knows, and so on) about the heights of objects outside their perceptual range, so they can't reasonably presuppose that they will make the same presuppositions about the heights of those objects. Because neither Anna nor John know what the other knows about the heights of objects

²⁰ The worry arises from the thought that, for many (perhaps most) context-dependent sentences, there is a huge number of different linguistic meanings—understood as functions from Kaplanian contexts to possible-worlds propositions—any language user could have associated with each of those sentences. Given the huge number of such different linguistic meanings, it is unlikely that any two language users associate the exact same linguistic meaning with any of a wide variety of sentences.

outside their perceptual range, they also can't reasonably make any presuppositions about the standard of tallness to be assumed for the conversation on the basis of the heights of those objects. Thus, if information about the reigning standard of tallness is to be part of the common ground, it must make it into the common ground on the basis of what is perceptually available to both Anna and John.

Yet Anna and John also can't reach a common conclusion about the reigning standard of tallness on the basis of their perceptual evidence. Since there are small variations between most people's perceptual apparatuses, what looks to Anna to be a certain height may look to John to be a slightly different height—e.g. what looks to Anna to be 1.8m tall may look to John to be 1.82m tall. Given this difference in the way things look to Anna and John, it is very unlikely that they will make exactly the same presuppositions about the heights of the objects within their perceptual range. Thus, on the assumption that presuppositions about the standard of tallness are arrived at on the basis of presuppositions about the heights of certain objects, it is also unlikely that Anna and John will make the same presuppositions about the standard of tallness relevant to their conversation.

I assumed above that neither Anna nor John knew what the other knew about the heights of objects in their perceptual range. Now that I have pointed out that it's unlikely Anna and John could reach the same conclusion about an object's height through ordinary perceptual evidence (i.e. without the aid of a measuring instrument), we can weaken that assumption. Even if Anna and John know that they are both opinionated about the height of a certain object outside their perceptual range, it is unlikely that they will be able to make the same presuppositions about that object's height. Ordinarily, they will come to make presuppositions about the object's height just by looking at it, without the aid of any measuring device; given the perceptual differences between Anna and John, it is unlikely that they will make the exact same presuppositions about that object's height.

Adjectives like 'tall' are the best case for the defender of the common-ground strategy. At least in principle, the participants in a conversation could reach the same conclusion about the relevant standard of tallness on the basis of perceptual evidence available to all of them. Yet I have argued that, even for a word like 'tall', the participants in a conversation are unlikely to make the exact same presupposition about the relevant standard in the context of the conversation.

It is even more difficult to see how the participants in a conversation could all make the same presuppositions about the relevant standard for adjectives like 'rich', 'nice', 'cheap', 'expensive', 'relevant',

and perhaps the vast majority of gradable adjectives. In those cases, it is even less plausible to think that all the participants in a conversation will make the same assumptions that could determine a relevant standard for the adjective. In the case of ‘rich’ or ‘wealthy’, for instance, it is unlikely that all the participants in a conversation will make the same assumptions about the average wealth of people in a certain population, let alone about the wealth of particular people.

Here is my conclusion from the discussion of common ground. It is unlikely that the participants in a linguistic interaction make the exact same assumptions pertaining to the values of the contextual parameters which determine the truth-conditional content of utterances involving gradable adjectives. Absent such agreement about the values for the relevant contextual parameters, the common-ground strategy is not in a position to explain why all the participants in a conversation would come to have the same beliefs about the truth-conditions of utterances involving gradable adjectives. Similar remarks apply to the other cases of context-sensitivity.

The observations I have presented so far are not meant to show that the notion of common ground is theoretically useless. The assumption that every conversation has a common ground is a fruitful idealization in formal pragmatics—the study of the reasoning processes that language users engage in in the context of a conversation—and nothing of what I said attempts against making that idealization when engaging in formal pragmatics. My point is instead that the strategy that uses common ground is not a convincing way of resisting the case for Variance.²¹

3.3 Social Externalism

According to social externalism, an individual’s social environment may produce differences in the contents of that individual’s thoughts, beliefs, and so on.²² For example, if social externalism is true, people with the same brain states may have different beliefs because they belong to different communities. For instance, following Burge (1979), suppose Anna thinks she has arthritis in her thigh. According to Burge, if Anna had been in the same brain state but belonged to a community which used the word ‘arthritis’ to denote *tharthritis* rather than arthritis, Anna’s thought would have been about *tharthritis*

²¹ See Lederman (forthcoming) for further challenges to the notion of common knowledge.

²² A prominent defender of social externalism is Burge (1979, 1986), but the general idea may be traced back to Putnam (1975). See Williamson (2007) for appeals to social externalism in objections to analyticity. Ludlow (1995, 1997), Pollock (2015), Wikforss (2001), among others, criticize social externalism.

and not about arthritis.

Some philosophers may want to appeal to social externalism to resist the case for Variance. They may think that, given social externalism, it is very likely that members of the same linguistic community have the same beliefs about the truth-conditional content of most utterances. They may reason as follows. If social externalism is true, then part of what determines someone's beliefs about the truth-conditional content of a given utterance is her social environment, such as the linguistic community that person belongs to. Thus, they may claim, any two members of the same linguistic community are very likely to have the same beliefs about the truth-conditional content of ordinary utterances on account of belonging to the same linguistic community. Since most people who engage in conversations with each other belong to the same linguistic community—say, the community of speakers of the language in which the conversation takes place—they will have the same beliefs about the truth-conditional content of the utterances involved in their conversation.

Even if social externalism is true, this line of reasoning overestimates the extent to which social environment can determine what we believe. Social externalism states that the linguistic community a person belongs to plays a role in determining her beliefs by determining the content of her internal states—for instance, her brain states. But different members of the same linguistic community may have different beliefs on account of having different internal states—such as brain states—regardless of the truth of social externalism. For example, even if Anna and John belong to the same community, they may have different beliefs about the truth-conditional content of Anna's utterance of 'Carla likes to run' (i.e. one may think that its truth-conditional content is (3), and the other that it is (4)) as a result of the fact that one of them is in brain state b_1 and the other is in brain state b_2 . Thus, social externalism notwithstanding, the fact that two people belong to the same linguistic community does not by itself make it likely that they have the same beliefs about an utterance's truth-conditional content, since the fact that two people belong to the same linguistic community does not make it likely that they have the same internal states (e.g. brain states) or internal states that are equivalent in the relevant ways.

We can get a better grasp of this point and of the way in which social environment determines the content of our beliefs about an utterance's truth-conditional content by using the language of thought hypothesis. Roughly, according to that hypothesis, to believe a proposition P is to stand in a certain

relation to a sentence in one's language of thought whose content is P.²³ To use a metaphor due to Schiffer (1981b), we can think of that relation to a sentence as that of having a token of that sentence in one's belief box.

Call the language of thought *mentalese*. If social externalism is true, then social environment determines the contents of our beliefs by determining the contents of the mentalese sentences tokened in our respective belief-boxes. So, for example, Anna's social environment determines that if Anna has a token of the mentalese sentence *M* in her belief-box, she believes that the truth-conditional content of her utterance of 'Carla likes to run' is (3). In this way, Anna's social environment partly determines her beliefs about her utterance's truth-conditional content by determining the content of the mentalese sentence about that utterance's truth-conditional content which are tokened in her belief-box.

What social environment does *not* determine is which of a wide range of mentalese sentences (whose contents are determined by social environment) pertaining to an utterance's truth-conditional content is tokened in Anna's belief-box. Yet, for nearly every utterance in natural language, there are as many mentalese sentences expressing different propositions about that utterance's truth-conditional content as possible truth-conditional contents for the utterance; given the huge number of such mentalese sentences, it is unlikely that any other person will token exactly the same mentalese sentence as Anna in his belief box. Since the same is true of any two language users, social externalism by itself is not enough to resist the case for Variance.

Defenders of the present strategy may attempt to address the issue as follows. To begin with, individuate expression-types of natural language so as to count the disambiguation of the word 'bank' that denotes a river bank as one expression type, and the disambiguation of the word 'bank' that denotes a financial institution as another. Defenders of the present strategy will assume, first, that users of a language usually associate expression-types in natural language with expression-types in mentalese. Second, they will assume that those mentalese expression-types combine very much in the way their natural-language counterparts combine to form mentalese counterparts of natural-language sentence types. Third, they will assume that if two people belong to the same linguistic community and they token mentalese sentence-types that are counterparts of the same natural-language sentence-type, the tokens of those mentalese sentence-types have the same content (i.e. express the same proposition). Fi-

²³ See Field (1978) and Fodor (1987).

nally, they will assume that someone's beliefs about the truth-conditional content of a natural-language utterance are given by content statements (stated in mentalese) linking that utterance with the mentalese counterpart of the uttered sentence—i.e. mentalese sentences of the form *U has the truth-conditional content that S*, where 'S' is to be replaced with a sentence of mentalese.²⁴

For example, suppose that Anna and John are members of the same linguistic community, and that Anna utters 'Carla likes to run' in the course of a conversation with John. Suppose further that Anna associates the English words 'Carla', 'likes', 'to', and 'run' with the mentalese words 'C', 'l', 't', and 'r',²⁵ respectively, and that John associates the same English words with the mentalese words 'C'', 'l'', 't'', and 'r'', respectively. According to this view, Anna will token the mentalese sentence

(13) Anna's utterance of 'Carla likes to run' has the truth-conditional content that Cltr

in her belief-box, whereas John will token the mentalese sentence

(14) Anna's utterance of 'Carla likes to run' has the truth-conditional content that C'l't'r'.

Since Anna and John belong to the same linguistic community and the mentalese sentence types 'Cltr' and 'C'l't'r'' are both counterparts of the English sentence 'Carla likes to run', tokens of (13) and (14) in Anna and John's respective belief boxes will express the same proposition. This means that Anna and John will have the same beliefs about the truth-conditional content of Anna's utterance of 'Carla likes to run'.²⁶

As appealing as this strategy may seem when it comes to context-independent sentences, it lacks the resources to resist the case for Variance when it comes to context-dependence. The problem can be stated as a dilemma. On the one hand, if expression-types of natural language are individuated so as to count all utterances of a sentence like 'that box is heavy' (regardless of the circumstances in which the

²⁴ Cf. Abreu Zavaleta (2018).

²⁵ Expressions written in this font are expressions of mentalese.

²⁶ Note that, if the present strategy is to have any plausibility, it must leave room for the possibility of misunderstandings between members of the same linguistic community. For example, the view must allow that, in certain circumstances, one of the participants in a conversation may believe that an utterance of 'Anna is next to the bank' has the truth-conditional content that Anna is next to a certain financial institution, while another believes that the same utterance has the truth-conditional content that Anna is next to a certain river bank. In the present implementation of the strategy, this flexibility is achieved by taking mentalese sentence-types to be counterparts of natural-language sentence-types, and by individuating expression-types of natural language so that the disambiguation of 'bank' that denotes financial institutions and the disambiguation of 'bank' that denotes river banks count as different expression-types. If expression-types in natural language are individuated in this way, it is in principle possible for two people to associate mentalese types that are counterparts of different disambiguations of 'bank', which in turn allows the view to predict that misunderstandings are possible (though, according to this view, very uncommon).

utterances are made) as tokens of the same type, the present view is at odds with context-dependence. It predicts that people believe every utterance of a given context-dependence sentence to have exactly the same truth-conditional content, regardless of the circumstances in which the utterance takes place.

On the other hand, if we individuate natural-language expression-types so as to count utterances of ‘that box is heavy’ with different truth-conditional contents as tokens of different types, the present strategy does little to resist considerations from multiple candidates. For instance, suppose we individuate sentence-types so as to count tokens of ‘that box is heavy’ evaluated with respect to different Kaplanian contexts as tokens of different types. Given the enormous number of different sentence-types ‘that box is heavy’ could be disambiguated between, it is unlikely that Anna and John token sentences in their respective belief-boxes linking an utterance of ‘that box is heavy’ with mentalese counterparts of the exact same disambiguation. The same holds, *mutatis mutandis*, for other context-dependence sentences susceptible to multiple-candidate considerations—such as sentences involving gradable adjectives.

To see this, suppose there are standards of heaviness $h_0, h_1, \dots, h_{10000}$, and that each of them determines a different truth-conditional content for an utterance U of ‘that box is heavy’ when given as input to that sentence’s linguistic meaning. Given the number of such standards, there will be at least as many different sentence-types ‘that box is heavy’ could be disambiguated between, each with different truth-conditions. And, given the number of sentence-types ‘that box is heavy’ could be disambiguated between, it is unlikely that there are mentalese counterparts of the exact same such sentence-type such that Anna’s belief-box has a token of the mentalese sentence resulting from substituting one of them for ‘S’ in

(15) Utterance U of ‘that box is heavy’ has the truth-conditional content that S,

and John’s belief-box has a token of the mentalese sentence resulting from substituting the other for ‘S’ in the same schema. Thus, the present strategy does not explain why it would be likely that the participants in a conversation have the same beliefs about the truth-conditional content of an utterance of ‘that box is heavy’. The same goes for other context-dependent sentences.²⁷

²⁷ One possible reaction is to assume that sentences of mentalese are themselves context-sensitive (see Gross (2005) for discussion). But this only pushes the problem back. For suppose that Anna tokens Utterance U of ‘that box is heavy’ has the truth-conditional content that S in her belief box, and John tokens Utterance U of ‘the box is heavy’ has the truth-conditional content that S’, where S and S’ are to be replaced by context-sensitive mentalese counterparts of ‘that box is heavy’. Now the issue is whether both sentences should be interpreted with respect to the same contextual parameters in virtue

If what I have said so far is correct, the present strategy is not enough to resist multiple-candidates considerations regarding context-dependent sentences. Those considerations are enough to support a version of Variance restricted to context-dependent sentences, which in turn would provide sufficient grounds for rejecting the standard view of communication. Since the present strategy does not succeed in upholding that standard view of communication, it is not clear what could motivate it as an account of people's beliefs about the truth-conditions of context-*independent* sentences in the first place.

4 Three kinds of communication-likeness

If *Variance* is true, successful communication can't be typically achieved due to the fact that the participants in a conversation have the same beliefs about the truth-conditions of the utterances they make. Nor can successful communication require, even in normal circumstances, that the participants in a conversation infer from each other's utterances exactly what the other intended to communicate. For the remainder of this paper, I want to explore the question of what successful communication *does* require. In particular, I will discuss three different ways in which a linguistic interaction may be communication-like—i.e. ways in which a linguistic interaction may resemble ideal cases of successful communication (cases in which speaker and audience have the same beliefs about the truth-conditions of the utterances they make)—despite the truth of *Variance*. What I want to propose is that there is no unique natural phenomenon that we talk about when we talk about successful communication; rather, there are several independent features a conversation may have which make it communication-like. As we will see later on, this leads to the adoption of a pluralistic picture of communicative success.

A linguistic interaction can be communication-like in that (a) it allows its participants to exchange certain pieces of relevant information with each other, in that (b) the participants in the interaction can have certain kinds of agreements or disputes about the truth of the utterances made in the course of the interaction that are not merely verbal, or in that, (c) given the right circumstances, the interaction may put (some of) its participants in a position to achieve certain goals in a distinctly non-coincidental way.

of being counterparts of the same English sentence. If the answer is yes, then misunderstandings will be impossible, since it will be impossible for any two people to disagree over the truth-conditional content of any context-sensitive sentence. If the answer is no, we still face the problem of explaining why, given the huge number of equally eligible contextual parameters, the sentences that replace S and S' would be evaluated with respect to the exact same one. Thanks to an anonymous reviewer for helpful discussion.

These ways for a linguistic interaction to be communication-like are not exhaustive and, as we will see soon, they can come apart from one another.

Here is an example. Suppose Anna believes that the word ‘green’ denotes things of colors 2–5, whereas John believes that it denotes things of colors 3–6 (see figure 1). Anna is talking to John about the house she recently moved to, which is color 4. Prompted by John’s question about the house’s color, Anna utters:

(16) It is green.

Assume for the sake of simplicity that both Anna and John believe that ‘it’ as it occurs in (16) refers to *H*, Anna’s house.²⁸ Because of their respective beliefs about the word ‘green’, Anna believes that (16) as she used it is true if and only if *H* is one of colors 2–5, whereas John believes that (16) is true if and only if *H* is one of colors 3–6.²⁹ Anna uttered (16) intending to communicate to John that her house is one of colors 2–5, but, as a result of hearing Anna’s utterance, John thinks that Anna wanted to communicate that her house is one of colors 3–6.³⁰



Figure 1: Some colors

Despite the difference in Anna and John’s beliefs about (16)’s truth-conditions and John’s mistaken belief about what Anna intended to communicate, Anna and John’s conversation resembles ideal cases of successful communication in the following ways. First, through her utterance of (16), Anna transmitted to John information about *H*’s color: that *H* is neither color 1 nor color 7, that *H* is not red, and so on. All this information is exactly about the subject matter Anna was talking about—i.e. *H*’s color—and in that sense Anna transmitted relevant information to John through her utterance of (16).

Second, given the circumstances in which Anna and John’s interaction takes place, Anna and John can have certain kinds of agreements and disputes about (16)’s truth that are not merely verbal. For

²⁸ This assumption is controversial in the literature on anaphoric pronouns, but the controversy does not matter for present purposes. See King and Lewis (2017) for an overview of related issues.

²⁹ I am assuming for simplicity that the participants in a conversation have determinate beliefs about the truth-conditions of ordinary utterances. See Abreu Zavaleta (2018) for discussion of cases in which this assumption doesn’t hold.

³⁰ Notice that, for all John takes for granted for the purposes of the conversation, Anna’s house may be any one of colors 1–7. Thus, the proposition that Anna’s house is one of colors 2–5 and the proposition that Anna’s house is one of colors 3–6 are not equivalent relative to the common ground in Anna and John’s conversation.

example, suppose that shortly after the conversation takes place John disputes the truth of (16) by uttering ‘that’s not true!’, and that, pointing to H in a photograph, Anna replies ‘yes it is, my house is *this* color (color 4)!’. In such a case, assuming John has normal vision and given his beliefs about (16)’s truth-conditions, John could not reasonably reply with something like ‘Wait a moment! The fact that the house is that color does not make what you said true! That color is not green!’. That such a response would not be available to John suggests that Anna and John’s dispute is not purely verbal.

Finally, in the right circumstances, Anna’s interaction with John would put the latter in a position to achieve certain goals. For example, suppose Anna invites John to her new house and John does not know the exact address, though he does know the house’s block; as it happens, there are only two houses on Anna’s block: Anna’s, which is color 4, and another house, which is color 1. If (16) is true by *John’s* lights, the thing to do is to go to the house that is one of colors 3–6. Since Anna’s house is in fact one of colors 3–6, believing that (16) is true will put John in a position to reach Anna’s house (given his beliefs about (16)’s truth-conditions). Moreover, given the circumstances in which Anna and John’s conversation took place, the differences in the truth-conditions Anna and John believe (16) to have do not matter for the purposes of reaching Anna’s house in the present scenario: since Anna’s house is in fact color 4 and the other house is color 1, John could have reached Anna’s house regardless of whether he had believed that (16) is true if and only if H is one of colors 3–6 or that it is true if and only if H is one of colors 2–5. This is not a mere coincidence: since Anna’s house is color 4 and the other house is color 1, going to the house that is one of colors 3–6 is the same action as going to the house that is one of colors 2–5.

It might be tempting to think that these three ways for an interaction to be communication-like can all be explained in terms of similarity between propositions. According to this view, it is because the propositions

(17) H is one of colors 2–5 and

(18) H is one of colors 3–6

are similar enough that Anna transmitted information to John about H ’s color, that Anna and John can have agreements or disputes about (16)’s truth which are not merely verbal, and that, in the circumstances discussed above, John could reach Anna’s house thanks to his interaction with Anna.

That temptation should be resisted. (17) and (18) are as similar (given any ordinary notion of similarity) in any possible circumstance as they are in any other, yet the three ways for an interaction to be communication-like I just discussed can come apart depending on the circumstances in which the interaction takes place. For example, in a scenario in which Anna's house is color 2 rather than 4, and the other house is color 6 rather than 1, (17) and (18) would still have entailed that Anna's house is not red, and in that sense Anna would have still transmitted information to John about the color of her house through her utterance of (16). In that scenario, however, Anna and John could not have had the kind of non-merely-verbal agreements and disputes about (16)'s truth they could have had in the original scenario. If in the new scenario John disputed (16)'s truth and Anna attempted to settle the issue by pointing to the house's color (i.e. color 2) as what makes (16) true, it would have been perfectly reasonable for John to reply (given his beliefs about the meaning of 'green'): 'Wait a moment! The fact that the house is that color does not make what you said true! That color is not green!'—which suggests the dispute would have been purely verbal. Hence, since (17) and (18) are as similar to each other in this new scenario as they were in the original one, that similarity does not suffice for Anna and John to have non-merely-verbal agreements or disputes about (16)'s truth.

Nor does the similarity between (17) and (18) suffice for Anna and John's interaction to put John in a position to reach Anna's house in this new scenario, or for the differences between those two propositions not to matter for the purpose of reaching Anna's house. If (16) is true by *John's* lights, the thing to do is to go to the house that is one of colors 3–6. Since Anna's house is in fact color 2, believing that (16) is true will not put John in a position to reach Anna's house (given his beliefs about (16)'s truth-conditions): if John believes that (16) is true, he will go to the house color 6, which is not Anna's house. Furthermore, unlike in the original example, the differences between (17) and (18) *do* matter for the purposes of reaching Anna's house. Since Anna's house is in fact color 2, if John had believed that (16) is true if and only if *H* is one of colors 2–5, believing that (16) is true would have put him in a position to reach Anna's house. Yet (17) and (18) are as similar in this new scenario as they were in the original one.

Similarity between propositions does not explain the communication-like features of a conversation we are interested in, but I believe those features should nevertheless be explained in terms of various relationships between propositions. I outline such explanations below, but a complete analysis falls out

of the scope of this paper.³¹

Let's start with the transmission of relevant information. We can capture this feature in terms of a relationship between propositions: two people can exchange relevant information with one another through a declarative utterance just in case the propositions those people associate with the uttered sentence—i.e. the propositions those people respectively take to be true if and only if the sentence is true—have a relevant entailment in common. Here I will define relevant entailment using a version of Fine's truthmaker semantics (see his 2016, 2012, MS), but it could also be defined using Yablo's (2014) version. In truthmaker semantics, a proposition's possible truthmakers and falsitymakers are taken to be possible states of affairs, which can stand in parthood relations with other states of affairs. A proposition's possible truthmakers (falsitymakers) are the possible states of affairs which are fully relevant to the determination of the proposition's truth-value, and which would guarantee the proposition's truth (falsity) were they to obtain. For example, a possible truthmaker for (17) is a state of affairs in which *H* is color 2, and a possible falsitymaker for the same proposition is a state of affairs in which *H* is color 1. A proposition's subject matter is represented as the set of that proposition's possible truthmakers and falsitymakers.³²

Using these notions, we can define the following notion of *relevant entailment*: P relevantly entails Q if and only if (i) every possible truthmaker for P has a possible truthmaker for Q as a part, and (ii) every element in Q's subject matter is part of an element in P's subject matter. By the lights of this notion of entailment, (17) and (18) both entail that *H* is not red, that it is not color 1 or color 7, etc. Thus, the present view accurately predicts that Anna transmitted relevant information to John through her utterance of (16).³³

Let us turn now to agreements and disputes. I said above that, in the scenario in which Anna's house is color 4, if John disputed the truth of (16) and Anna attempted to settle the issue by pointing to her house's color (color 4), John could not have reasonably replied something like 'the fact that the

³¹ See Abreu Zavaleta (2018) for further development of these views.

³² Roughly, this corresponds to the idea that a proposition's subject matter is determined by the states of affairs that play a direct role in making a proposition true or false.

³³ Here is a rough explanation of why (17) and (18) both relevantly entail that *H* is not red. All states of affairs in which *H* is one of colors 2–5 (the possible truthmakers for (17)) and all states of affairs in which *H* is one of colors 3–6 (the possible truthmakers for (18)) are themselves state of affairs in which *H* is not red (and so, parts of states of affairs in which *H* is not red). Second, every state of affairs that is either a truthmaker or a falsitymaker for the proposition that *H* is red concerns exclusively *H*'s color, and as such, is either a truthmaker or a falsitymaker for (17), and either a truthmaker or a falsitymaker for (18). Thus, both (17) and (18) relevantly entail that *H* is not red.

house is that color does not make what you said true!'. On the other hand, a similar response would be perfectly reasonable in a scenario in which Anna's house is color 2.

I propose to use truthmaker semantics to account for this difference in Anna and John's disputes in the two different scenarios: what distinguishes those two disputes is the fact that, in the original scenario but not in its variant, the propositions Anna and John (respectively) associate with (16) have the exact same truthmakers. Because in the original scenario (17) and (18) have the same truthmakers, Anna and John's dispute could have been settled merely by pointing to those facts. In contrast, in the second scenario, the fact that makes (17) *true* (the fact that Anna's house is color 2) is the fact that makes (18) *false*, and so Anna and John's dispute would not have been resolved regardless of whether they knew the facts responsible for the truth-value of (16) by their respective lights.

More generally, what distinguishes the kind of disputes and agreements I called *non-purely verbal*—i.e. the kind illustrated by Anna and John's dispute in the original scenario—is the fact that, in those disputes or agreements, the proposition the disputants or agreeers associate with the sentence whose truth is under discussion have the same truthmakers or the same falsitymakers relative to the scenario (i.e. possible world) in which the agreement or dispute takes place. According to this view, whether a dispute about the truth of a sentence is purely verbal depends both on which propositions the disputants associate with the sentence whose truth is under dispute *and* on the circumstances in which the dispute takes place. In that sense, the present view is a kind of externalism about non-purely verbal disputes.

Finally, I said that Anna and John's original conversation resembled ideal cases of successful communication in that, given the circumstances in which said conversation took place, it put John in a position to reach Anna's house in a non-coincidental way. In particular, I said that, given the circumstances in which Anna and John's conversation took place, the differences between the propositions they associate with (16) do not matter for the purpose of reaching *H*. Put informally, the differences between two propositions don't matter for the purposes of achieving a goal in given circumstances whenever, in those circumstances, any way of doing what one ought to if one of those propositions is true is a way of doing what one ought to if the other proposition is true, and vice versa. For example, given that Anna's house is color 4 in the original circumstances in which Anna and John's interaction takes place, any way of going to the house that is one of colors 2–5 is a way of going to the house that is one of colors 3–6.

We can capture this intuition more formally by representing a goal as a set of (possibly overlapping) sets of possible worlds, each of which corresponds to a way of achieving the goal depending on how the world is. For example, the goal of reaching H will contain a set of worlds in which H is color 2, one in which it is color 3, one in which it is 5m tall, etc. Those sets of worlds correspond to ways to reach Anna's house: in worlds in which H is color 2, one can reach H by going to the house that is color 2; in worlds in which H is 5m tall, one can reach it by going to the house that is 5m tall, etc.

If we represent goals as sets of sets of possible worlds, we can say that a proposition is *useful* for the achievement of a goal in a given possible world w just in case (a) it relevantly entails a proposition which is true in exactly the possible worlds in some union of sets in the goal, and (b) w is a member of at least one of those sets. Using the present framework, we can capture what it is for the differences between two propositions not to matter for the achievement of a given goal in terms of the notion of *practical equivalence*. Say that a set in a given goal is *maximally specific* just in case it has no proper subsets which are also members of the goal. Two propositions are *practically equivalent* relative to a possible world w and a goal just in case every maximally specific set in the goal of which w is a member and in which one of those propositions is true is a set in which the other proposition is true, and vice versa. Intuitively, this corresponds to the idea that every way of doing what one ought to do in w in order to achieve the goal if one of the two propositions is true is a way of doing what one ought to do in order to achieve the goal if the other proposition is true. Where P is the proposition the audience associates with a given sentence, and Q is the proposition the speaker associates with the same sentence, hearing an utterance of that sentence puts the audience in a position to achieve a certain goal g in a non-coincidental way relative to a given possible world w just in case: (a) P is useful to the achievement of the goal in w , and (b) every proposition relevant to the goal which is relevantly entailed by P is *practically equivalent* relative to g and w to a proposition relevant to the goal which is relevantly entailed by Q , and vice versa.

The fact that the three communication-like features I just examined can come apart suggests we should adopt a pluralistic picture of communicative success. In particular, I want to propose that there is no unique natural phenomenon that we talk about when we talk about successful communication; instead, there are several independent communication-like features a conversation may have, none of which is more fundamental than the rest. From this pluralistic perspective, the assumption that the

participants in a conversation have the same beliefs about the truth-conditions of the utterances they make obscures the differences between the various communication-like features a conversation may have and the facts that account for their presence. This is so because, if the participants in a conversation have exactly the same beliefs about the truth-conditions of the utterances they make, conversations between them will tend to have all or most of the conversation-like features I introduced, among others.

According to the pluralistic conception of communicative success, *Variance* does not threaten the idea that we successfully communicate with one another through most of our conversations. Instead, it illuminates the way to a more complex conception of communicative success, according to which communication as the standard picture conceives it can be thought of as a limit or ideal of communicative success. Ordinary conversations rarely, if ever, reach that ideal, but they still succeed in having some or all of the communication-like features I have described in this section, among others.

Should our theorizing require a notion of communicative success, we can define various notions of communicative-success on the basis of those communication-like features. For instance, in one sense of success, communication will be successful just in case the participants in a conversation exchange relevant information through the conversation. In another sense of success, communication is successful just in case participants in the conversation are able to engage in the kind of non-verbal agreements and disputes about the truth-value of certain utterances (i.e. the utterances made in the course of their conversation) I discussed above. In the third sense of success, communication is successful as long as the differences between the propositions the participants in a conversation associate with a given sentence do not matter for certain practical purposes. These three senses of communicative success need not be exhaustive and, as the examples show, can come apart from each other depending on the circumstances in which a linguistic interaction takes place. Thanks to the analysis I just outlined, we can make concrete predictions about the cases in which each kind of communicative success can be achieved.

5 Alternatives to pluralism?

Not everyone will be attracted to the pluralistic approach. Perhaps, they may think, there are versions of the standard picture that allow us to explain all the ways there are for conversations to be

communication-like while retaining the idea that there is a single natural phenomenon that communicative success amounts to. Before concluding, I will examine two views of this kind.

What characterizes these views is the idea that there is a special kind of content speakers can “share” with their audiences despite the truth of Variance, even if that content is not the utterance’s truth-conditional content. According to the first of these strategies, inspired by Stalnaker (1999a), even if normal language users don’t agree on the truth-conditional content of the utterances they use to communicate, there is still a proposition the audience can take the speaker to have expressed: the diagonal proposition associated with the sentence the speaker utters. According to the second strategy, even if speaker and audience fail to agree about what an utterance’s truth-conditional content is, they can agree on certain metalinguistic commitments that the utterance imposes on the conversational score. The rest of this section discusses these proposals.

5.1 Diagonalization

According to Stalnaker’s model of assertion, we can represent the body of information the participants in a conversation presuppose as a set of possible worlds: the set of worlds in which all those presuppositions are true. Stalnaker calls this set of worlds the *context set*, and adds that there are three “essential conditions” the context set must meet in cases of rational communication:

1. A proposition asserted is always true in some but not all of the possible worlds in the context set.
2. Any assertive utterance should express a proposition, relative to each possible world in the context set, and that proposition should have a truth-value in each possible world in the context set.
3. The same proposition is expressed relative to each possible world in the context set.

(Stalnaker 1999a, p.88)

For our purposes, it will suffice to focus on the first and last principles. The first principle rules out cases in which an assertion is uninformative, either because what is asserted is already presupposed, or because it is incompatible with what is presupposed. The third principle amounts to the requirement that the participants agree on which proposition is asserted.

Stalnaker acknowledges that there may be some cases of seemingly rational communication in which these requirements are not satisfied if what is asserted is the literal meaning of a sentence. To use a modified version of Stalnaker’s example, suppose Anna and John know that there’s a woman next door, but don’t know whether it is Carla or Xóchitl. After hearing the woman talk, Anna thinks she recognizes Xóchitl’s voice, and utters:

(19) That is Xóchitl.

Assuming that the semantic content of ‘that’ is just its referent, this sentence will express different propositions relative to different worlds in the context set, thus violating the third principle. In worlds in which Carla is the woman in the room next door, (19) will express

(20) Carla is Xóchitl.

In worlds in which Xóchitl is the woman in the room, (19) will express:

(21) Xóchitl is Xóchitl

If c is a world in which Carla is the woman next door, and x is a world in which Xóchitl is the woman next door, we can represent the situation using table 1. In that table, the vertical axis represents worlds in their role of determinants of what an utterance expresses, and the horizontal axis represents them as the worlds in which the proposition expressed by an utterance is to be evaluated. Furthermore, note

	c	x
c	F	F
x	T	T

Table 1: The proposition expressed by (19) depending on how the world is, with the diagonal highlighted in green.

that the first principle is violated as well regardless of which of the two propositions is expressed. On the one hand, (20) is necessarily false, hence false in all worlds in the context set. On the other, (21) is necessarily true, hence true in all worlds in the context set. Either way, Anna’s utterance of (19) does not express an informative proposition.

How can we reconcile the intuition that this is a case of rational communication with the violation of Stalnaker’s first and third requirements? According to Stalnaker, Anna didn’t assert (20) or (21).

Instead, she asserted the “diagonal” proposition which is true in a given world w just in case the proposition expressed by (19) in w is true in w . Keeping fixed the referent of ‘Xóchitl’, this is the proposition that Xóchitl is the referent of ‘that’ as it is used in Anna’s utterance. This proposition is false in worlds in which ‘that’ refers to Carla, and true in worlds in which ‘that’ refers to Xóchitl (see tables 1 and 2). This proposition is (19)’s diagonal regardless of the world in which (19) is uttered, is true in some but not all the worlds in the context set, and is such that Anna and John are both in a position to know that Anna meant it through her utterance of (19). Thus, we can take Anna and John’s interaction to be a case of rational communication by taking Anna to have expressed the diagonal.

	c	x
c	F	T
x	F	T

Table 2: The diagonal proposition associated with (19) in Anna and John’s context.

How is this related to the troubles arising from Variance? According to the standard picture of communication, communication is successful only if the audience knows what proposition the speaker wanted to communicate, and is typically achieved thanks to shared knowledge of the truth-conditional content of the utterances used to communicate. Given the truth of Variance, communicative success cannot be typically achieved thanks to mutual knowledge of the truth-conditional content of the utterances used to communicate, nor can it require that the audience knows exactly what the speaker wanted to communicate—since, given that language users typically will lack shared knowledge of the truth-conditional content of the utterances they use to communicate, it is difficult to see how they could know exactly what the speaker wanted to communicate. Some theorists may be inclined to appeal to the diagonal to avoid these difficulties: they may claim that if what the speaker wants to communicate is always the diagonal proposition associated with the sentence she uttered, then as long as speaker and audience agree on what that diagonal is, it doesn’t matter that they fail to agree on the utterance’s truth-conditional content. Furthermore, those theorists may think that, since the participants in a conversation are always in a position to know what diagonal proposition is associated with the uttered sentence, they will always be in a position to know that the speaker meant to communicate the diagonal.

In this way, the present strategy endorses the part of the standard picture of communication according to which what it is for communication to be successful is for there to be a proposition the

speaker wants to communicate and which the audience knows the speaker wants to communicate. The difference is that, according to the present strategy, in most cases of successful communication that proposition is the diagonal associated with the sentence the speaker uttered, rather than the utterance's truth-conditional content.

For example, consider once more Anna and John's conversation in which Anna utters (16). Because Anna and John disagree about the meaning of the word 'green', they also disagree about the truth-conditional content of Anna's utterance: Anna believes that it is (17), and John believes that it is (18). According to the present strategy, and despite this disagreement, John is still in a position to know that Anna meant to express the diagonal proposition associated with 'it is green'. Assuming for simplicity that 'it' refers to H in every possible world, that proposition will be true in every possible world in which H has whatever property the word 'green' expresses in that world. For example, the diagonal will be true in a world in which 'green' expresses the property of being one of colors 2–5 and H is color 2; false in a world in which 'green' expresses the property of being one of colors 3–6 and H is color 2; etc. Since, for any given sentence, there is exactly one diagonal which it could plausibly be associated with, Anna and John should be in a position to agree on what the diagonal is regardless of their beliefs about the utterance's truth-conditional content. Thus, John will be in a position to know that Anna expressed this diagonal proposition through her utterance of (16).

I am skeptical that appealing to the diagonal will take us very far. While it is true that typical language users will be in a position to know that the speaker intended to communicate the diagonal proposition associate with the sentence she uttered, it typically won't suffice for successful communication that the audience knows that the speaker wanted to communicate the diagonal. For instance, suppose that (due to an unorthodox upbringing) Carla believes that the word 'green' applies to all and only things that are at least 5m tall. Wishing to inform Carla of her new house's color, Anna utters (16) ('it is green') in conversation with Carla.

Anna and Carla will then disagree about the truth-conditional content of Anna's utterance. Anna will believe that it is the proposition (17), and Carla will believe that it is

(22) H is at least 5m tall.

Despite this disagreement, Carla may still be in a position to know that Anna expressed the diagonal:

(23) H has whatever property ‘green’ expresses.

The problem is that the diagonal is too weak to serve Anna’s communicative purposes. Anna intended to give Carla information about H ’s color through her utterance of (16), but the diagonal entails no such information, since it can be true regardless of H ’s color. For example, the diagonal is true in all worlds in which ‘green’ expresses the property of being at least 5m tall and H is 6m tall, regardless of H ’s color. Thus, knowing that Anna expressed the diagonal doesn’t put Carla in a position to obtain any new information about H ’s color, contrary to Anna’s communicative intentions.

Nor does knowing that Anna expressed the diagonal put Anna and Carla in a position to have substantive disputes over the truth of Anna’s utterance. For example, suppose that immediately after Anna utters (16), Carla disputes the truth of Anna’s utterance by uttering ‘that’s not true’. If Anna and Carla agree that Anna’s utterance expresses (23), then their dispute will ultimately be over whether (23) is true. Yet Anna and Carla may disagree about (23)’s truth-value merely due to a disagreement over what ‘green’ means, and not due to a disagreement about facts about Anna’s house. This will be the case if, for example, H is in fact color 4 and 4m tall, and Anna and Carla both know this. In that case, Anna and Carla’s disagreement can only amount to a disagreement over what property ‘green’ refers to, rather than about the properties of Anna’s house. Thus, the fact that Anna and Carla agree that Anna’s utterance expresses (23) is not sufficient for their dispute over the truth of Anna’s utterance to be more than merely verbal.

I will leave further developments of the diagonalization strategy to its supporters, but the observations in the present section should make us skeptical that the diagonalization strategy can be worked out in a way that explains the different ways in which a conversation can be communication-like outlined in the previous section.³⁴

5.2 Metalinguistic negotiation

One recent development in the debate between contextualists and anti-contextualists concerns the idea of *metalinguistic negotiation*. Authors like Barker (2002), Khoo (forthcoming), Khoo and Knobe (2016), Sundell (2011) concede that there may be situations in which speaker and audience disagree

³⁴ Thanks to an anonymous reviewer for helpful discussion of the issues in this subsection

over the contextual standard relevant to the truth-conditional content of a given utterance. This will be the case if, for example, Anna utters ‘Carla is tall’, thinking that the relevant standard of tallness is one on which someone counts as tall just in case she is tall at least to degree .65, whereas John thinks that the relevant standard is one on which someone counts as tall just in case she is tall at least to degree 0.7.

According to these authors, in those cases the aim of a conversation may not be the transmission of information about Anna’s height in the standard way. Rather, part of what the speaker may attempt to do is to coordinate with her audience on what the relevant standard for tallness is. For example, if Anna and John both know that Carla is tall to degree 0.65, Anna may utter ‘Carla is tall’ to convince John that the relevant standard is one on which someone who is tall to degree 0.65 counts as tall. In those cases, it would seem that communication is successful despite a potential disagreement about the truth-conditional content of the utterance used to communicate. Call this the *metalinguistic negotiation strategy*.³⁵

There are many implementations of the present strategy, but here I will consider Khoo’s (forthcoming). Following Lewis (1979), Khoo starts with the notion of a *conversational score*: a set of indices, each of which we may think of as a sequence of a world and whatever other parameters the evaluation of a sentence may depend on—standards of tallness, of wealth, etc. The idea is that a conversational score keeps track of the information presupposed by the participants in a conversation (the common ground), but also of further elements relevant to the truth of regular assertoric utterances—standards of wealth, height, etc.

An utterance of the form ‘ a is F ’, where a is a proper name and F is a predicate whose application depends on standard g , is true at an index i just in case the object denoted by a satisfies the standard g_i determined by i in the possible world w_i determined by i . ‘ a is F ’s content at a context c will be the set of sets of possible worlds such that, for some index i in the conversational score determined by c , ‘ a is F ’ is true relative to that possible world and the standard g_i . Finally, the effect of an assertoric utterance of ‘ a is F ’ will be to modify the conversational score so that it contains only indices at which ‘ a is F ’ is true; to assert ‘ a is F ’ is to propose changing the conversational score so that ‘ a is F ’ is true at all of its indices. This is all summarized as follows (cf. Khoo forthcoming, pp. 14–15), where ϕ is a

³⁵ In relating authors like Barker, Khoo, and Sundell to this strategy, I do not mean to suggest that they would propose it as a solution to the problems arising from Variance. Rather, my aim is to examine a view that one could construct using the resources they employ for different purposes, and assess the viability of such a view.

sentence whose truth depends on a contextually supplied standard g , and \mathbb{S} is a conversational score:

Truth at an index: $\llbracket \phi \rrbracket^{c,i} = T$ iff $\llbracket \phi \rrbracket^{c,w_i,g_i} = T$

Content: $\llbracket \phi \rrbracket^c = \{f : \exists i \in \mathbb{S}_c (f = \lambda w. \llbracket \phi \rrbracket^{c,w_i,g_i} = T)\}$

Update value: $|\phi|^c = \{i : \llbracket \phi \rrbracket^{c,i} = T\}$

Assertion: To make an assertoric utterance of ϕ at c is to propose changing \mathbb{S}_c to $\mathbb{S}_c \cap |\phi|^c$.

How does this help give necessary and sufficient conditions for communicative success? An initial idea is that communication through an utterance of ϕ will be successful just in case the speaker means the content of ϕ (i.e. a set of propositions) and the speaker recognizes what content that is. Unfortunately, and for roughly the reasons that motivate Variance, this is a non-starter. ϕ 's content, as given by the definition above, depends on what particular indices are “alive” in a conversational score. In turn, this depends on what values for the parameters on which the truth of ϕ depends are considered alive. The problem is that, typically, it will be difficult for speaker and audience to agree on a set of live parameters for the purposes of the conversation.

For example, suppose that Anna utters ‘Carla is tall’ in conversation with John, and she thinks that at that point in the conversation the relevant parameters for ‘tall’ are those according to which someone is ‘tall’ just in case she is tall at least to degree 0.65, at least to degree 0.7, and everything in between. The problem is that there are many sets of parameters any language user could easily have taken to be alive in the conversational score just as easily; given the huge number of such sets of parameters, it would be extremely unlikely for any two people to agree on which of those sets contains exactly the parameters that are alive in the conversational score. Absent agreement on which parameters are alive in the conversational score, it will be extremely unlikely for any two people to agree on which content the utterance has. For instance, anyone could just as easily have believed that the relevant parameters are instead those according to which someone is ‘tall’ just in case she is tall at least to degree 0.66, at least to degree 0.72, and everything in between. Should John believe this, he will believe that Anna’s utterance has a content other than the one Anna herself believes her utterance to have.

In this way, even if we think of utterance content as Khoo does (as opposed to truth-conditional content), it is extremely unlikely for ordinary language users to agree on which content an utterance

has. However, there is a notion among the ones defined above on which ordinary language users can plausibly agree: update value. After all, as long as Anna and John agree that the application of ‘tall’ depends on a standard of tallness, they can agree that the update value of ‘Anna is tall’ is the set of indices such that Anna is tall relative to that index’s possible world and standard of tallness.

Motivated by this idea, someone may claim that communication through an assertoric utterance of ϕ is successful just in case the speaker means to update the conversational score (whatever the score is) with ϕ ’s update value, and the audience knows that the speaker means to do this. Since speaker and audience can easily agree on update value, ordinary language users will typically be able to satisfy this condition. For example, even if Anna and John disagree on what the conversational score is, they will likely agree that the update value of ‘Anna is tall’ is the set of all indices in which that sentence is true—i.e. the set of all indices i such that, in w_i , Anna is tall at least to degree g_i .

The problem is that, even in cases involving sentences like ‘Carla is tall’, it doesn’t seem that agreement on the update value of the uttered sentence suffices for communicative success. If Anna and John endorse radically different standards of tallness for the purposes of their conversation, we would not always think that their interaction was successful. For example, suppose Anna endorses a standard according to which Carla is tall just in case she is tall at least to degree 0.3—Carla may be a child, and Anna may think that Carla is tall for a child—while John endorses a standard on which Carla is tall just in case she is tall at least to degree 0.8—he may be recruiting for his amateur basketball team. This, I take it, would be a typical case of a misunderstanding.

Furthermore, the fact that Anna and John agree over the update value of Anna’s utterance of ‘Carla is tall’ is not sufficient for them to have disputes over the truth of Anna’s utterance that are more than merely verbal. For example, Anna and John may agree over Carla’s actual degree of tallness, yet disagree over the truth of Anna’s utterance. In that case, their disagreement will be merely over what standard of tallness to adopt, and not about Anna’s actual height or degree of tallness. In this sense, we may think of their dispute as a merely verbal one.³⁶

I take the present discussion to show that agreement on a sentence’s update value (as understood by

³⁶ As authors like Khoo and Sundell point out, there is a sense in which metalinguistic disagreements are genuine disagreements. But what we’re after is not just the notion of a “genuine” disagreement. Rather, we want to capture conflicts between people’s beliefs about the facts they take an utterance to be about; in this case, conflicts between Anna and John’s beliefs over Carla’s degree of tallness. See Abreu Zavaleta (2018, MS) for further discussion.

proponents of metalinguistic negotiation) is not sufficient for successful communication. Nor can it give us a unifying explanation of the various ways in which a conversation may be communication-like. This is enough to reject the metalinguistic approach as a unified account of communicative success. However, before concluding, I want to highlight some significant problems in generalizing the metalinguistic approach.

The first of these problems concerns indexicals like ‘here’ or ‘now’. Though the problems from Variance arise for utterances involving those terms as well, it is unclear how to generalize the present strategy so as to cover them. For example, suppose Anna utters

(24) Carla was here yesterday

intending to inform John of Carla’s location the day before the conversation took place.

The natural way of generalizing the metalinguistic negotiation strategy to cases like this is by taking (24)’s update value to be something like:

(25) $\{i : \text{in } w_i, \text{ Anna is at } l_i \text{ the day before } t_i\}$,

where l_i and t_i are a location and a time determined by i , respectively. According to the present strategy, as long as Anna and John agree that (25) is (24)’s update value, their conversation will be an instance of successful communication.

Intuition tells against this result. For example, suppose that Anna and John’s conversation occurs over the phone, that Anna is in her house, and that John is in his hotel room in Tokyo. As it happens, Anna and John disagree about the conversational score: Anna believes that it contains only indices in which the location is her house, whereas John believes that it contains only indices in which the location is his hotel room in Tokyo. As a result, Anna intends to make the conversational score such that it contains only indices whose possible world is one in which Carla was in Anna’s house the day before the conversation took place. Yet John will take Anna to intend to exclude indices in whose world Anna was not in John’s hotel room in Tokyo the day before the conversation took place. This would seem like a clear case of miscommunication.

Nor can the present account explain the transmission of relevant information or the possibility of having non-merely-verbal disputes in this case. First, (25) entails no concrete information about Carla’s location the day before Anna and John’s conversation. In particular, (25) doesn’t exclude possible

worlds in which Anna is in, say, Paris, Mars, or outside the Milky Way; instead, it only excludes *indices* such that, in the world determined by the index Carla was not at the location determined by the index the day before the time determined by the index. Second, since two people may agree about Carla's actual location without agreeing on what location is the one relevant to the conversational score, agreeing on the update value of Anna's utterance is not sufficient for having a non-merely-verbal dispute over that utterance's truth.³⁷

The second problem concerns terms that are not context-dependent. It is not obvious how to extend the present strategy to those cases. Where ϕ is a context-sensitive sentence whose truth depends on contextual parameter g , we took its update value to be the set of indices such that ϕ is true at the index's world relative to the value the index assigns to parameter g . But if ϕ is not context sensitive, its truth-value depends only on the index's world, regardless of the values the index assigns to other contextual parameters. Which is to say that if ϕ is not context sensitive, both its content and its update value are basically exhausted by its truth-conditional content: the set of possible worlds in which ϕ is true. Thus, given the truth of Variance, if ϕ is not context sensitive, ordinary language users will fail to agree on what ϕ 's update value is.

Defenders of the present strategy may try to address this issue by extending their treatment to context-independent expressions as follows. They may claim that indices not only determine contextual parameters, but also interpretation functions that determine the meaning of the context-independent predicates in a given sentence. On this approach, if a is a name and F is a context-independent predicate, $\lceil a \text{ is } F \rceil$ is true at index i just in case, in the possible world determined by i , the object denoted by a has the property assigned to F by i 's interpretation function. Finally $\lceil a \text{ is } F \rceil$'s update value will just be the set of indices at which it is true.

For example, let's assume that the verb 'to run' is not context-sensitive. Now consider the sentence

(26) Carla is running,

and the following three indices. Index 1 includes a world in which Carla is moving on her feet at about

³⁷ I want to insist that I am not attributing the present view of indexicals to authors who have defended the existence of metalinguistic negotiation as a natural phenomenon. Authors like Khoo (forthcoming), Khoo and Knobe (2016), Barker (2002), and other proponents of metalinguistic negotiation, are careful to claim that the strategy does not apply *tout court*, and often restrict their position to discussions of gradable adjectives or moral terms. Yet the problems arising from Variance are much more general than that, and our aim here is to see if the metalinguistic negotiation strategy can solve those problems.

10km/h and is not sneezing, and an interpretation function according to which ‘is running’ applies exactly to things that move on their feet at a speed of at least 7km/h. Index 2 includes the same world as index 1, in which Carla is moving on her feet at about 10km/h and is not sneezing, but an interpretation function according to which ‘is running’ applies exactly to things that are sneezing. Index 3 includes a possible world in which Carla is standing still and sneezing, and the same interpretation function as index 2, i.e. one on which ‘is running’ applies exactly to things that are sneezing. Then (26) is true at indices 1 and 3, but false at index 2. Accordingly, 1 and 3, but not 2, will be included in (26)’s update value.

Now, suppose that Anna utters (26) in conversation with John intending to give him information about what Carla is doing, and that they agree on (26)’s update value: they agree that its update value includes indices 1 and 3, but not 2. The problem with the present view is that (26)’s update value is too weak to serve Anna’s communicative purposes. Though Anna intended to give information to John about Carla’s activities, (26)’s update value entails no such information: it can be true at an index regardless of whether Carla is running, sneezing, or standing still, as long as the interpretation of ‘is running’ determined by the index suitably matches what Carla is doing in the index’s world. Thus, Anna and John’s agreement on (26)’s update value is not sufficient for Anna to give information to John about Carla’s activities.

Nor is agreement on (26)’s update value sufficient to put Anna and John in a position to have non-merely-verbal disputes over the truth of Anna’s utterance. For example, Anna and John may both believe that Anna moves at a speed of 10km/h and is not sneezing, but have a dispute over the truth of Anna’s utterance due to a disagreement over how to interpret ‘is running’—Anna may interpret it so that it expresses the property of moving on one’s feet at a speed of at least 7km/h, and John may interpret it so that it expresses the property of sneezing. In that case, their dispute would be merely verbal, since it would arise merely because Anna and John interpret ‘is running’ differently.

Finally, even pretheoretical intuitions seem to tell against the claim that agreement over (26)’s update value is sufficient for successful communication. For suppose Anna endorses an interpretation on which ‘is running’ expresses the property of moving on one’s feet at a speed of at least 7km/h, and John endorses an interpretation on which ‘is running’ expresses the property of sneezing. This, I take it, would be a typical case of miscommunication, yet the present approach predicts that communication

would be successful.

The discussion shows that appealing to metalinguistic negotiation does not explain the various ways in which a conversation may be communication-like, nor does it give us an adequate account of communicative success. That said, we should be open to the idea that the notion of metalinguistic negotiation may be useful in explaining certain ways for a conversation to be communication-like, with metalinguistic disagreements being the prime candidate. The point is just that metalinguistic negotiation does not explain *all* the ways in which a conversation may be communication-like.

6 Conclusion

In this paper I argued for *Variance*, the thesis that nearly every utterance is such that there is no proposition which more than one language user believes to be that utterance's truth-conditional content. The case for *Variance* relies on the observation that, for nearly every utterance, there are enormously many equally eligible propositions that language user could easily have taken to be that utterance's truth-conditional content. Given this observation, it is very unlikely that there is a proposition which more than one language user takes to be that utterance's truth-conditional content. I supported this line of reasoning with examples from the literature on context-dependence and examples of context-independent sentences subject to the same considerations. I also considered and rejected ways of resisting the case for *Variance* which appeal to the notion of common ground and to social externalism, respectively.

In the introduction I mentioned that a consequence of *Variance* concerns the nature of successful communication: if *Variance* is true, communication can't normally be successful thanks to shared beliefs about the truth-conditions of the utterances made in the course of an interaction, nor can communication normally require that the audience be in a position to tell exactly what the speaker wanted to communicate on the basis of the latter's utterances. Towards the end of the paper, I outlined an alternative picture of communicative success: the pluralistic picture. According to the pluralistic picture, there are many independent ways for a conversation to be communication-like, none of which require the participants in a linguistic interaction to have the same beliefs about the truth-conditions of the utterances they make. From the perspective of the pluralistic conception, part of the aim of a theory of communication is to understand the different ways in which linguistic interactions can be communication-like

and the facts in virtue of which those interactions can be communication-like in those ways. I have discussed some of those ways for a conversation to be communication-like in section 4, but no doubt others can be discovered. As I argued in section 5, seemingly natural unifying explanations of these features fail.

I would like to conclude by calling attention to two further consequences of *Variance* and the discussion of communicative success. One such consequence concerns the debate between contextualists and invariantists. Contextualism is the view that the semantic value of many words other than the standard indexicals (e.g. ‘I’, ‘here’, ‘now’) depends on the context in which they are uttered. According to Cappelen and Lepore (2005), if contextualism is true, then people who find themselves in different contexts would not understand the same utterances in the same way. According to them, that would make communication between people in different contexts difficult to explain.³⁸ But if, as I claim, *Variance* is true, there is no special problem of cross-contextual communication: any account of communication should explain how people communicate despite failing to have the same beliefs about the truth-conditions of the utterances they make.³⁹ In line with pluralism about communicative success, contextualists can attempt to explain the ways in which cross-contextual linguistic interactions can be communication-like without appealing to shared beliefs about a sentence’s truth-conditions.

Another consequence concerns a recent puzzle about speech reports due to Dorr and Hawthorne (2014). I take Dorr and Hawthorne to argue that facts about the proposition(s) someone semantically expresses through a given utterance are highly modally plastic: very tiny differences in the underlying microphysical facts would lead to differences in the proposition(s) someone semantically expresses through a given utterance. If this is so, Dorr and Hawthorne claim, most of our counterfactual speech reports are likely to be false. We can take a hint from pluralism about communicative success to explain why, even if the claim about plasticity holds, this need not affect the modal robustness of ordinary speech reports. In particular, one can adopt a view of speech reports according to which there are

³⁸ “If communicated contents are restricted to (or essentially tied to) specific contexts of utterance, then it is hard to envision how speakers who find themselves in different contexts can communicate, i.e. under such circumstances communication between contexts is thrown into doubt” (Cappelen and Lepore 2005, p.153)

³⁹ The same is true of an objection to holism about meaning often attributed to Fodor and Lepore. Jackman (2017) presents the objection as follows: “Strictly speaking, informative communication would be impossible [if meaning holism is true]. No one would mean the same thing by any of their terms unless they shared all the same beliefs, in which case, communication would be possible, but uninformative, and truly understanding the utterances of others would be impossible unless you already knew everything that they believed.” Yet if *Variance* is true, the problem of explaining what informative communication requires arises for nearly everyone.

different ways for them to be true: such reports can be true if they attribute to the speaker the saying of a proposition relevantly entailed by what the speaker semantically expressed, if they attribute the saying of a proposition that, in the possible world in which the reported speech act takes place and every possible world sufficiently similar to it, has the same truthmakers or the same falsitymakers as the proposition the speaker semantically expressed, etc.

A better understanding of *Variance* and the study of the various ways in which a conversation-like can thus play an important role in coming to understand other phenomena. Cross-contextual communication and counterfactual speech reports are two of them.

References

- Abreu Zavaleta, Martín. 2018. *Semantic Variance*. Ph.D. thesis, New York University.
- . MS. “Disagreement Lost.”
- Bach, K. and Harnish, R. 1979. *Linguistic Communication and Speech Acts*. MIT Press.
- Barker, Chris. 2002. “The Dynamics of Vagueness.” *Linguistics and Philosophy* 25:1–36.
- Buchanan, Ray. 2010. “A Puzzle About Meaning and Communication.” *Noûs* 44:340–371.
- Burge, Tyler. 1979. “Individualism and the Mental.” *Midwest Studies in Philosophy* 4:73–122.
- . 1986. “Intellectual Norms and Foundations of Mind.” *Journal of Philosophy* 83:697–720.
- Cappelen, Herman and Lepore, Ernest. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Wiley-Blackwell.
- Cresswell, Max J. 1977. “The semantics of degree.” In Barbara Partee (ed.), *Montague Grammar*, 261–292. New York: Academic Press.
- Dorr, Cian and Hawthorne, John. 2014. “Semantic Plasticity and Speech Reports.” *Philosophical Review* 123:281–338.
- Dummett, Michael. 2010. “Language and Communication.” In Darragh Byrne and Max Kölbel (eds.), *Arguing About Language*. Routledge.
- Evans, Gareth. 1982. *Varieties of Reference*. Oxford University Press.
- Field, Hartry. 1978. “Mental Representation.” *Erkenntnis* 13:9–61.
- Fine, Kit. 2012. “Counterfactuals Without Possible Worlds.” *Journal of Philosophy* 109:221–246.
- . 2016. “Angelic Content.” *Journal of Philosophical Logic* 45:199–226.

- . MS. “Truthmaker Semantics.” URL: https://www.academia.edu/10908756/survey_of_truthmaker_semantics, DOA: August 25, 2016.
- Fodor, Jerry A. 1983. *The Modularity of Mind*. Cambridge, MA: MIT Press.
- . 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. MIT Press.
- Grice, H. P. 1989a. “Meaning.” In *Studies in the Way of Words*. Harvard University Press.
- . 1989b. “Utterer’s Meaning and Intention.” In *Studies in the Way of Words*. Harvard University Press.
- Gross, Steven. 2005. “Context-Sensitive Truth-Theoretic Accounts of Semantic Competence.” *Mind and Language* 20:68–102.
- Hawthorne, John and Magidor, Ofra. 2009. “Assertion, Context, and Epistemic Accessibility.” *Mind* 118:377–397.
- . 2010. “Assertion and Epistemic Opacity.” *Mind* 119:1087–1105.
- Heck, Richard. 2002. “Do Demonstratives Have Senses?” *Philosophers’ Imprint* 2:1–33.
- Heim, Irene. 2000. “Degree Operators and Scope.” In Brendan Jackson and Tanya Matthews (eds.), *Semantics and Linguistic Theory 10*, 40–64. Ithaca, NY: Cornell University.
- Heim, Irene and Kratzer, Angelika. 1998. *Semantics in Generative Grammar*. Blackwell.
- Jackman, Henry. 2017. “Meaning Holism.” In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2017 edition.
- JBiz et al. 2010. “How fast do you have to be going to count it as running?” DOA: July 27, 2017.
- Kaplan, David. 1989. “Demonstratives.” In Joseph Almog, John Perry, and Howard Wettstein (eds.), *Themes From Kaplan*, 481–563. Oxford University Press.
- Kennedy, Christopher. 2007. “Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives.” *Linguistics and Philosophy* 30:1–45.
- Kennedy, Christopher and McNally, Louise. 2005. “Scale structure, degree modification, and the semantics of gradable predicates.” *Language* 345–381.
- Khoo, Justin. forthcoming. “Quasi Indexicals.” *Philosophy and Phenomenological Research* .
- Khoo, Justin and Knobe, Joshua. 2016. “Moral Disagreement and Moral Semantics.” *Noûs* 109–143.
- King, Jeffrey C. and Lewis, Karen S. 2017. “Anaphora.” In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2017 edition.
- Landau, Barbara, Smith, Linda, and Jones, Susan. 1988. “The importance of shape in early lexical learning.” *Cognitive development* 3:299–321.
- . 1998. “Object shape, object function, and object name.” *Journal of Memory and Language* 38:1–27.

- Larson, Richard and Segal, Gabriel. 1995. *Knowledge of Meaning: An Introduction to Semantic Theory*. MIT Press.
- Lederman, Harvey. forthcoming. "Uncommon knowledge." *Mind*.
- Lewis, David. 1975. "Languages and Language." In Keith Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*, 3–35. University of Minnesota Press.
- . 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8:339–359.
- . 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61:343–377.
- Ludlow, Peter. 1995. "Externalism, Self-Knowledge, and the Prevalence of Slow-Switching." *Analysis* 55:45–49.
- . 1997. "On the Relevance of Slow Switching." *Analysis* 57:285–86.
- Markman, Ellen M. 1990. "Constraints children place on word meanings." *Cognitive Science* 14:57–77.
- Markman, Ellen M and Hutchinson, Jean E. 1984. "Children's sensitivity to constraints on word meaning: Taxonomic versus thematic relations." *Cognitive psychology* 16:1–27.
- Pettit, Dean. 2002. "Why Knowledge is Unnecessary for Understanding Language." *Mind* 111:519–550.
- Pollock, Joey. 2015. "Social Externalism and the Problem of Communication." *Philosophical Studies* 172:3229–3251.
- Portner, Paul. 2005. *What is Meaning?: Fundamentals of Formal Semantics*. Blackwell.
- Putnam, Hilary. 1975. "The Meaning of 'Meaning'." *Minnesota Studies in the Philosophy of Science* 7:131–193.
- Schiffer, Stephen. 1972. *Meaning*. Oxford, Clarendon Press.
- . 1981a. "Indexicals and the Theory of Reference." *Synthese* 49:43–100.
- . 1981b. "Truth and the Theory of Content." In Herman Parret (ed.), *Meaning and Understanding*, 204–222. Berlin.
- Sider, Theodore. 2011. *Writing the Book of the World*. Oxford University Press.
- Stalnaker, Robert C. 1999a. "Assertion." In *Context and Content: Essays on Intentionality in Speech and Thought*, 78–95. Oxford: Oxford University Press.
- . 1999b. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford: Oxford University Press UK.
- . 1999c. "Pragmatic Presuppositions." In *Context and Content: Essays on Intentionality in Speech and Thought*, 47–62. Oxford: Oxford University Press.
- . 2009. "On Hawthorne and Magidor on Assertion, Context, and Epistemic Accessibility." *Mind* 118:399–409.

- Strawson, P. F. 1964. "Intention and Convention in Speech Acts." *Philosophical Review* 73:439–460.
- . 1970. *Meaning and Truth: An Inaugural Lecture Delivered Before the University of Oxford on 5 November 1969*. Oxford, Clarendon P.
- Sundell, Timothy. 2011. "Disagreements About Taste." *Philosophical Studies* 155:267–288.
- Wikforss, Asa Maria. 2001. "Social Externalism and Conceptual Errors." *Philosophical Quarterly* 51:217–31.
- Williamson, Timothy. 2007. *The Philosophy of Philosophy*. Blackwell.
- Yablo, Stephen. 2014. *Aboutness*. Princeton University Press.