

Semantic Variance

by

Martín Abreu Zavaleta

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Philosophy
New York University
September, 2018

James Pryor

©Martín Abreu Zavaleta
All Rights Reserved, 2018

Acknowledgements

Above all, I'd like to thank Jim Pryor, Cian Dorr, and Stephen Schiffer for their continuous encouragement and feedback. I want to thank them especially for always pushing me to go further, to be clearer, and overall to write better philosophy. If this dissertation has any merit, it is thanks to them.

One of my largest intellectual debts is due to Kit Fine and Stephen Yablo. Their work on truthmaker semantics has paved the way for many of the views I present in this dissertation. In many ways, I hope this dissertation contributes to showing how powerful their approach is, and how fruitful its application to semantic and philosophical problems can be.

Special thanks go also to the members of the intense but short-lived NYU Philosophy of Language Lab: Kyle Blumberg, Ben Holguín, Paul Horwich, and Richard Stillman. Thanks also to Dan Hoek and Ian Grubb for our many conversations on the nature of subject matters, and for being friends and sparring partners for many years during my stay at NYU.

I have benefited from conversations with members of various philosophical communities in the New York and Boston areas. For conversations about the topics of this dissertation, discussions in seminars or at NYU's thesis prep, thanks to Amanda Askill, Kent Bach, David Balcarras, Alan Barat, Max Barkhausen, Chris Barker, Nathaniel Baron-Schmitt, Aliosha Barranco, Sam Berstler, Harjit Bhogal, Justin Bledin, David Boylan, Michael Caie, Rosa Cao, David Chalmers, Lucas Champollion, Jeremy Dolan, Michelle Dyke, Sam Carter, Andreas Ditter, Rose Flinn, Vera Flocke, Carolina Flores, Naomi Francis, Laura Franklin-Hall, Jane Friedman, Melissa Fusco, Dmitri Gallow, Iliana Gioulatou,

Simon Goldstein, Camil Golub, Jeremy Goodman, Cosmo Grant, Yu Guo, Martin Hackl, Dan Harris, Sam Hesni, Arturo Javier, Justin Khoo, Arden Koehler, Allison Koslow, Harvey Lederman, Andrew Lee, Sam Lee, Robert Long, Adam Lovett, Tienmu Ma, Matt Mandelkern, Annette Martin, Carla Merino-Rajme, Ricardo Mena, Friederike Moltmann, Matt Moss, Olivia Murton, Sofía Ortiz, Asya Passinsky, Carlotta Pavese, Zee Perry, Milo Phillips-Brown, Rohan Prince, Hsueh Qu, Agustín Rayo, Mark Richard, Leo Rosenstein, Chelsea Rosenthal, Raúl Saucedo, Chris Scambler, Ginger Schulteis, Dan Sharp, Erica Shumener, Jack Spencer, Zeynep Soysal, Una Stojnic, Anna Szabolcsi, Melisa Vivanco, Jared Warren, Daniel Waxman, Mike Zhao. If I failed to mention someone, it is only due to my poor memory. Thanks also to audiences at the Institute for Advanced Studies in Paris, Massachusetts Institute of Technology, New York University, The New School for Social Research, the University of Turin, and Syracuse University for valuable feedback. Special thanks go to audiences at the New York Philosophy of Language Workshop.

Thanks to my friends for being great friends.

Thanks to Erica Shumener for being the best partner anyone could ever wish for (and, more generally, the best simpliciter). Not only has Erica offered me her unwavering support and helped me stay sane throughout the last few years; she has often been the first to hear many of the ideas in this dissertation, and every time she has offered her patient, attentive, and ever so insightful comments on them. I could not be happier.

Gracias a mi mamá, Luz del Carmen, y a Mamá Tina por la dedicación y cuidado que dedicaron a mi formación. A nadie debo más que a ellas.

Abstract

This dissertation argues for Semantic Variance, the thesis that for nearly every utterance and any two language users, there is no proposition that those two language users believe to be that utterance's truth-conditional content. I argue that Semantic Variance is problematic for standard theories concerning the nature of communication, the epistemic significance of ordinary disputes, and the semantics of speech reports. In response to the problems arising from the truth of Semantic Variance, I develop new accounts of the transmission of relevant information, ordinary disputes, and the semantics of speech reports using truthmaker semantics. Towards the end of the dissertation I outline a pluralistic account about the nature of communication and linguistic competence.

Contents

Acknowledgements	iii
Abstract	v
List of Figures	ix
1 The case for Semantic Variance	1
1.1 Introduction	1
1.2 The case for Variance	5
1.2.1 Context-dependent sentences	17
1.2.2 Context-independence	21
1.3 The case against Variance	23
1.3.1 Naturalness	26
1.3.2 Common ground	31
1.3.3 Social Externalism	36
1.4 Three kinds of communication-likeness	42
1.5 Conclusion	47
2 Transmission of relevant information	48
2.1 Introduction	48

2.2	A standard picture of communication	51
2.3	The question-based picture	55
2.3.1	Standard treatment of questions	56
2.3.2	Troubles with Variance and Question-variance	60
2.3.3	Potential responses, and why they won't do	62
2.4	The “split records” picture of communication	70
2.5	Subject matters in truthmaker semantics	76
2.5.1	Background: states of affairs	78
2.5.2	Subject matters	83
2.5.3	What is a proposition about?	85
2.5.4	Putting subject matters to use	91
2.6	Subject matters and the split-records picture	93
2.7	Further issues	97
2.7.1	Exploiting information in a conversational record	97
2.7.2	Troubles with <i>Uncertainty</i>	98
2.8	Conclusion	101
3	Partly factual disputes	103
3.1	Introduction	103
3.2	Epistemic conflicts	106
3.3	Material conflicts	111
3.4	Partly factual disputes	114
3.5	Are typical disputes partly factual?	121
3.6	Troubles with <i>Uncertainty</i>	127
3.7	Further work for partly factual disputes	132

3.7.1	Contextualism and the problem of lost disagreement	132
3.7.2	Purely verbal disputes	137
3.8	Conclusion	140
4	Homophonic speech reports	143
4.1	Introduction	143
4.2	Troubles with Variance	146
4.3	Diagnosing the problem	152
4.4	Constraints on Looseness	156
4.5	An implementation of Looseness	160
4.5.1	Matching subject matters	160
4.5.2	Truth-conditional similarity	166
4.5.3	Descriptive similarity	179
4.6	Knowledge of homophonic reports	180
4.7	Further issues	188
4.7.1	Homophonic reports of homophonic reports	188
4.7.2	A generalization of Descriptive Looseness	192
4.7.3	Troubles with truth-conditional nihilism	196
4.8	Conclusion	198
5	Variance, competence, and pluralism	200
	Appendix	206
	Bibliography	210

List of Figures

1.1	Some colors	43
2.1	Some colors	49
3.1	Some colors	113
3.2	Some areas of Brooklyn.	116
4.1	Some colors	158
4.2	(18) and (19) in the space generated by a 's quantity of cheese and corn. . .	175
4.3	Mapping from values for $\mu((\mathbf{P} \cap \neg \mathbf{Q}) \cup (\neg \mathbf{P} \cap \mathbf{Q}))$ to degrees of similarity, with $\tau = 20$	177

Chapter 1

The case for Semantic Variance

1.1 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.

¹ These assumptions follow from two claims: that knowing the meaning of a sentence (and so, of an utterance of that sentence) requires knowing that sentence’s truth-conditions, and that ordinary language users typically know the meaning of sentences in the language they are users of. The former is often stated in semantics textbooks. For instance, Heim and Kratzer (1998) 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).

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. For example, suppose John wants to know what kind of food Carla likes, and Anna wants to inform him that Carla likes quesadillas. According to the standard picture of communication, Anna will make an utterance with the truth-conditional content that Carla likes quesadillas. Normally, according to the standard picture, if communication is successful John will come to know that Anna wanted to inform him that Carla likes quesadillas. 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 quesadillas, and second, that Anna believes that her utterance has the truth-conditional content that Carla likes quesadillas.^{2, 3}

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

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 (1974). 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 quesadillas' as Anna used it is true if and only if Carla likes quesadillas, 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.

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

This thesis has significant consequences pertaining to the nature of communication, the epistemic significance of ordinary disputes, the semantics of indirect speech reports, and the nature of linguistic competence, among others:

- *The nature of communication.* 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 sentences' 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 quesadillas' 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 quesadillas'.
- *The epistemic significance of ordinary disputes.* Typically, part of what is at stake in a dispute over the truth of an utterance is knowledge of the facts the disputants take that utterance to be about. For example, if someone utters 'grass is green' and someone else replies 'no, it's not', it seems that part of what is at stake in their dispute is knowledge of whether grass is green. This is easy to explain if the disputants both know that the utterance of 'grass is green' has the truth-conditional content that grass is green. For if one of them knows that the utterance is true, then the other doesn't know that grass is not green, and if one of them knows that the utterance is false, then

the other doesn't know that grass is green. However, if Variance is true, then nearly every dispute over an utterance's truth is, in a sense, verbal: there is no proposition that all the participants in the dispute take to be the truth-conditional content of the utterance the dispute is about. However, verbal disputes are not in general the kind of disputes in which knowledge of extralinguistic facts is at stake. Thus, if Variance is true, we must find a way of explaining the epistemic significance of typical disputes while acknowledging that such disputes are, in a sense, verbal.

- *Indirect speech reports.* According to standard accounts of speech and attitude reports, a report of the form 'S said that ϕ ', where S is a denoting term and ϕ is a declarative sentence, is true only if the referent of S uttered something whose truth-conditional content entails ϕ 's truth-conditional content (in the context in which the report is made).⁴ If Variance is true, this requirement is all too strong. For example, suppose Anna utters 'Carla likes quesadillas'. John hears Anna's utterance, and wants to know whether Anna said that Carla likes quesadillas. As it happens, the sentence 'Carla likes quesadillas' as it occurs in the report has the truth-conditional content that Carla likes quesadillas, but John thinks that the truth-conditional content of Anna's utterance is the slightly different and logically independent proposition that Carla likes quesadillas*, where quesadillas* are extremely similar to quesadillas, but not exactly the same. Then, for all John knows, Anna didn't utter anything whose truth-conditional content entails that Carla likes quesadillas, and so, he doesn't know that the report 'Anna said that Carla likes quesadillas' is true. As we will see in chapter 4, if Variance is true, ordinary language users are typically in a situation like John's. Thus, on pain of claiming that we never know what other people say, we

⁴ See e.g. Hintikka (1969). Even less standard accounts of speech reports, such as the one I defend in Abreu Zavaleta (Forthcoming), endorse this commitment.

must give up the standard account of speech reports.

- *The nature of linguistic competence.* According to standard accounts of linguistic competence (see e.g. Heim and Kratzer 1998, Portner 2005), knowing the meaning of a sentence requires knowing the truth-conditional content of utterances of that sentence. Soon we will see that any given utterance has at most one truth-conditional content. If that is so and Variance is true, then nearly every utterance is such that at most one language user *knows* what its truth-conditional content is. Thus, provided that ordinary language users are generally competent in the use of sentences they are familiar with, that competence can't require knowing the truth-conditional content of utterances of that sentence.

I will discuss some of these issues more thoroughly throughout the rest of the dissertation, but my focus in this chapter will be to argue for Variance.

The structure of the chapter is as follows. 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 (chapter 3). I argue that those strategies fail. Towards the end of the chapter I will 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). Then I conclude (section 5).

1.2 The case for Variance

Recall Variance:

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

I will present the case for Variance in a moment, but it is worth clarifying the notion of truth-conditional content before proceeding to the argument. 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, 6}

Given this definition, and on pain of contradiction, at most one of any two propositions that are not necessarily equivalent can be the truth-conditional content of any given utterance. Hence, if P and Q are not necessarily equivalent, any language user can take at most one of those propositions to be U's truth-conditional content without having contradictory beliefs. In the definition, context is to be understood 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

⁵ Sentences here should be thought of as disambiguated sentences.

⁶ I say 'necessarily' because mere material equivalences between the truth of a sentence at a context and the truth of a proposition are not enough to capture the idea that an utterance's truth-conditional content is part of that sentence's meaning. For example, given that snow is white and grass is green, if 'grass is green' is true relative to *c* if and only if grass is green, it follows that 'grass is green' is true relative to *c* if and only if snow is white, yet we would hardly think it is part of the meaning of 'grass is green' that it is true if and only if snow is white. Necessary equivalences between the truth of a sentence at a context and the truth of a proposition are thus better suited to capture the link between meaning and truth-conditions. Note that the present definition remains neutral with respect to whether an utterance necessarily has the truth-conditional content it in fact has. On one hand, the definition is compatible with the claim that one and the same utterance could have occurred in a different context, or that it could have been an utterance of a sentence in a language other than the language it is actually in. As such, it is compatible with the present definition that one and the same utterance could have had a different truth-conditional content from the one it actually has. On the other hand, the present definition is also compatible with the claim that if an utterance is an utterance of a sentence in a certain language or it occurs in a certain context, that is necessarily so. If that is the case, it follows from the definition that any utterance has its actual truth-conditional content as a matter of necessity. Thanks to Cian Dorr and Jim Pryor for discussion.

contexts formally understood to propositions) of the sentence that utterance is an utterance of with the context the utterance occupies.

With this in mind, here is the main point in favor of Variance. For nearly every assertoric utterance, there are enormously many—sometimes, uncountably many—different truth-conditional contents any language user could easily have believed each of those utterances to have, none of which is more natural or intrinsically more eligible than the rest. Given the vast number of equally eligible truth-conditional content candidates each of those utterances has, and absent further explanation, it would be extremely unlikely for there to be a proposition that more than one language user believes to be the truth-conditional content of one of those utterances. Which is to say that Variance is extremely likely.⁷

For example, suppose Anna utters

(1) Carla likes quesadillas

in conversation with John, and that generally speaking Anna calls something a quesadilla just in case it is a folded tortilla filled with cheese. It is not enough for Anna to count something as a “quesadilla” as she used the term that it be filled with just *any* amount of cheese. For example, it would not suffice for Anna to count something as a quesadilla that it contained only a very tiny, almost imperceptible amount of cheese, nor that it contained so much cheese that it bursts the tortilla open: quesadillas, according to Anna, must have an amount of cheese within a certain range.

Suppose that, as a matter of fact, Anna believes that all and only folded tortillas with between 20.5g and 120.5g of cheese count as “quesadillas” as she used the term in her

⁷ See Dorr and Hawthorne (2014) for a related argument to the effect that, if the propositions expressed by most sentences 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 those arguments and the one I present here in future work.

utterance. As a result, she believes that her utterance has the truth-conditional content that Carla likes tortillas filled with between 20.5g and 120.5g of cheese. There is nothing special about having between 20.5g and 120.5g of cheese that makes it especially easy for Anna to believe that things count as “quesadillas” just in case they have a quantity of cheese in that range. Had Anna been slightly differently attuned to the relevant evidence, or had that evidence been slightly different, she could just as easily have believed that things count as “quesadillas” according to her use of the word just in case they have between 20g and 120g of cheese, just in case they have between 119.5g and 139.5g of cheese, just in case they have between 120g and 135g of cheese, and so on. As a result, there are many truth-conditional contents Anna could easily have believed her utterance to have. She could easily have believed that her utterance’s truth-conditional content is the proposition that Carla likes folded tortillas filled with between 20g and 120g of cheese; that it is the proposition that Carla likes folded tortillas filled with between 119.5g and 139.5g of cheese; etc. More generally, for a very large number of ranges of quantities of cheese, Anna could easily have believed that her utterance’s truth-conditional content is the proposition that Carla likes folded tortillas with a quantity of cheese in that range.

Given the large number of ranges of quantities of cheese such that Anna could easily have believed that her utterance’s truth-conditional content is the proposition that Carla likes tortillas filled with a quantity of cheese in that range, it would be extremely unlikely for John to also believe that the truth-conditional content of Anna’s utterance is the proposition that Carla likes folded tortillas filled with between 20.5g and 120.5g of cheese. Absent further explanation, there must not be any range of quantities of cheese such that both Anna and John believe that Anna’s utterance has the truth-conditional content that Carla likes tortillas filled with a quantity of cheese in that range.⁸

⁸ I am restricting my attention to the case in which the quantity of cheese someone believes to be required for an object to be called a ‘quesadilla’ in English is stated in grams. If we restrict our attention in this

In the example, the differences in the truth-conditional content Anna and John may believe Anna's utterance to have are determined by differences in the quantities of cheese they may require something to have in order to call it a 'quesadilla'. Nothing relies on that particular feature of the example. There are many different dimensions on which the application of the word 'quesadilla' depends, and any of them would have been just as good: differences in Anna and John's beliefs about the truth-conditional content of Anna's utterance could be determined by differences in the shape they require things to have in order to call them 'quesadillas', by differences in those things' sizes, etc.

Nor does the example rely on features specific to the word 'quesadilla'. Observations of the kind I just presented hold for (but are not limited to) any sentence involving terms whose application depends on the properties an object has along one or more sufficiently fine-grained dimensions. In order to see this, suppose for example that 'F' is a predicate of that kind and 'a' is a proper name. For any language user, there will be a huge number of extremely similar and equally natural (or otherwise eligible) properties—each corresponding to a slightly different cutoff point along one or more of the dimensions on which the application of 'F' depends—such that that language user could easily have believed that an utterance of 'a is F' has the truth-conditional content that *a* (the object denoted by 'a') has that property. If, for example, a language user believes that 'F' expresses the property corresponding to the cutoff point x along one of the dimensions relevant to the application of 'F', 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. Given the huge number of properties any language

way, round quantities of cheese in grams may seem more natural candidates. However, different speakers may adopt different measuring systems, and this will affect what candidates may seem more natural from their perspective. For instance, a different speaker's beliefs may concern the quantity of cheese required for something to be called a 'quesadilla' in ounces, pounds, or some other measure, with consequences to which propositions they find to be natural candidates for the truth-conditional content of Anna's utterance. Thanks to Jim Pryor for discussion.

user could easily have taken ‘F’ to express, it is extremely unlikely for any two language users to believe that an utterance of ‘a is F’ has the truth-conditional content that *a* has the exact same one of those properties.⁹ The same applies, *mutatis mutandis*, to terms that express relations rather than properties.

Some people may think that Variance can be resisted as follows. Surely, they may claim, if ‘F’ is a predicate and ‘a’ a proper name, every language user believes that an utterance of ‘a is F’ has the truth-conditional content that *a* is F. Since such beliefs are easy to come by, something in the case for Variance must be mistaken. After all, the sentences that express such beliefs are all instances of the schema *utterance U of ϕ has the truth-conditional content that ϕ* . Or so the thought goes.

I think this approach overestimates the power of disquotation. To begin with, the approach overgenerates when it comes to utterances of context-dependent sentences (utterances of sentences whose semantic content varies depending on the context of utterance). For example, suppose Anna utters ‘I like apples’; surely language users will not in general agree that Anna’s utterance of ‘I like apples’ has the truth-conditional content that I like apples. Yet if all language users believed every instance of the schema above, language users would in general believe that Anna’s utterance of ‘I like apples’ has the truth-conditional content that I like apples.

The approach fails even if we restrict to context-*independent* utterances. To see this, suppose there is a property ‘F’ semantically expresses in all contexts, say, F-ness. F-ness will itself be just one among the huge number of extremely similar and equally natural properties that could easily have figured in any language user’s beliefs about the truth-

⁹ All that is required 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. As the discussion so far illustrates, ‘quesadilla’ is an example of the kind of predicate I refer to in the main text.

conditional content of an utterance of ‘*a* is *F*’. For *F*-ness will itself correspond to some particular cutoff point in the dimensions relevant to the application of ‘*F*’, and there will be a huge number of properties whose cutoff points are extremely close to and just as natural as the one determined by *F*-ness, each of which could easily have figured in any language user’s beliefs about the truth-conditional contents of the utterance in question. Thus, it is no more likely that the participants in a conversation all believe that an utterance of ‘*a* is *F*’ has the truth-conditional content that *a* is *F* (i.e. has *F*-ness) than that they all believe it has the truth-conditional content that *a* is *F*^{*}, where *F*^{*} is a property that determines a slightly different cutoff point than *F*-ness along some dimension relevant to the application of ‘*F*’. But, as we have seen, it is in general very unlikely that more than one language user believes one and the same property like *F*^{*} to be the semantic content of ‘*F*’.

For example, suppose that there is a property the expression ‘is a quesadilla’ semantically expresses in all contexts, say, the property of being a quesadilla. That property will be just one among a huge number of extremely similar and equally natural properties any language user could easily have believed ‘is a quesadilla’ to express. For the property of being a quesadilla will itself correspond to some particular cutoff point in the dimensions relevant to the application of ‘is a quesadilla’—e.g. size, quantity of cheese, and so on—and there will be a huge number of properties whose cutoff points along those dimensions are extremely close to and just as natural as the one determined by the property of being a quesadilla, each of which any ordinary language user could easily have believed ‘is a quesadilla’ to express. Thus, it is no more likely that more than one language user believes that ‘is a quesadilla’ expresses the property of being a quesadilla than that more than one language user believes that ‘is a quesadilla’ expresses the slightly different property of being a quesadilla^{*}—where the property of being a quesadilla^{*} determines a slightly different cutoff point than the property of being a quesadilla along some dimension relevant to the

application of the word ‘quesadilla’. In turn, it is no more likely that more than one language user believes that an utterance of ‘q is a quesadilla’ (where ‘q’ is a proper name) has the truth-conditional content that q is a quesadilla than that they all believe that that utterance has the truth-conditional content that q is a quesadilla*. Nor is it more likely that Anna and John both believe that Anna’s utterance of ‘Carla likes quesadillas’ has the truth-conditional content that Carla likes quesadillas (i.e. things with the property of being a quesadilla) than that they believe that that utterance has the truth-conditional content that Carla likes quesadillas*, (i.e. things with the property of being a quesadilla*).¹⁰ I will soon discuss more examples illustrating the wide variety of utterances for which observations of this kind hold, but first I want to make two clarifications.

First. 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 (more or less) precise vocabulary (e.g. ‘between 20g and 120g of cheese’, etc.), 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

¹⁰ Considering sentences from a language other than English may help. For example, the property of being green determines a particular cutoff point along the dimensions relevant to the application of the Spanish word ‘verde’ (usually translated as ‘green’), and there are many other properties whose cutoff points are extremely close to (and just as natural as) the one determined by the property of being green. Given the huge number of such properties, it is no more likely that the participants in a conversation in Spanish all believe that an utterance of ‘el pasto es verde’ has the truth-conditional content that grass is green, than that they all believe that it has the truth-conditional content that grass is green*, where the property of being green* determines a slightly different cutoff point than the property of being green along some dimension relevant to the application of ‘verde’.

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.¹¹

Second. Call *propositions* of the form *U has truth-conditional content P*—where *U* is an assertoric utterance and *P* is a proposition—*content propositions* about *U*. Variance is the thesis that, for nearly every utterance, there is no content proposition about that utterance which more than one language user believes. As such, one way for Variance to be true is for any two people to believe different content propositions about nearly every utterance, and another is for them not to believe any content proposition about those utterances in the first place.

If Variance is true in the first of these ways, that is already problematic for the standard picture of communication. For that would entail that successful communication is not nor-

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. Thanks to Jim Pryor and Chris Scambler for discussion.

mally achieved thanks to shared beliefs in content-statements about the utterances made in the course of a conversation, and that we can't normally know what someone intended to communicate on the basis of our beliefs about the truth-conditional content of her utterances. Yet to think that Variance is true in this way may concede too much to advocates of the standard picture. As I explain in what follows, observations of the kind I presented above support the stronger thesis that people do not usually believe *any* content-statements about utterances like Anna's in the first place.

Take again Anna's utterance of (1)—'Carla likes quesadillas'. I said above that there is a huge number of content-statements about that utterance anyone could easily have believed: that Anna's utterance has the truth-conditional content that Carla likes folded tortillas filled with between 20 and 120g of cheese, that it has the truth-conditional content that that Carla likes folded tortillas filled with between 20.5 and 120.5g of cheese, and, more generally, for any quantity of cheese in a certain range, that Anna's utterance has the truth-conditional content that Carla likes folded tortillas filled with that quantity of cheese. Given how similar those content propositions are to one another, it is very unlikely that the evidence available to someone who hears Anna utter (1) will significantly support a belief in one of those content propositions over a belief in another. Most likely, anyone who hears Anna utter (1) will be uncertain as to which among an enormous number of content propositions about her utterance in fact holds. Such a person will give very similar credence to a huge number of content propositions about Anna's utterance, and her credence in each of those content propositions will be too low for that person to count as believing any one of them.

For example, suppose John is entertaining the following propositions:

- (2) Anna's utterance of (1) has the truth-conditional content that
Carla likes folded tortillas filled with between 119g and 139g

of cheese;¹²

- (3) Anna's utterance of (1) has the truth-conditional content that Carla likes folded tortillas filled with between 119.5g and 139.5g of cheese;
- (4) Anna's utterance of (1) has the truth-conditional content that Carla likes folded tortillas filled with between 120g and 140g of cheese;
- (5) Anna's utterance of (1) has the truth-conditional content that Carla likes folded tortillas filled with between 120.5g and 140.5g of cheese;
- (6) Anna's utterance of (1) has the truth-conditional content that Carla likes folded tortillas filled with between 121g and 141g of cheese.

Given how similar (2)–(6) are to one another, it is unlikely that John's evidence significantly supports a belief in one of them over a belief in another. Most likely, John's credence in some of those propositions will be higher than his credence in others, but even his highest credence in one of those propositions will be too low for him to count as believing that proposition. For example, John may have credence 0.25 in (4), credence 0.2 in each of (3) and (5), and credence 0.175 in each of (2) and (6). Though in such a situation John would be more confident in (4) than in any other of (2)–(6), he would still not be confident enough to count as believing (4). In such situation, John would not believe any of (2)–(6).

When it comes to content propositions about utterances like (1), people usually are in a position very much like John's. Given the huge number of extremely similar con-

¹² Sentences written in this font stand for propositions.

tent propositions about an utterance like Anna's that any person could have believed, it is very unlikely that that person's evidence will support a belief in any content proposition about that utterance over a belief in one of the others to a significant degree. Ordinarily, people will be undecided as to which of a huge number of very similar content propositions about a given utterance holds, which will make even their highest credence in one of those statements barely significant. So, provided that people don't usually believe content propositions they barely have evidence for, it is very unlikely they will believe any content propositions about ordinary utterances.

Call the thesis that the participants in a conversation rarely believe any content propositions about ordinary utterances *Uncertainty*. Uncertainty entails Variance, but the converse does not hold. I believe that Uncertainty is true, but at many points throughout this dissertation I will make the simplifying assumption that people ordinarily believe content propositions about ordinary utterances. As we will see in chapters 2, 3, and 4, the points I make about communication, disagreement, and speech reports, continue to hold once we drop that simplifying assumption. I will omit discussion of Uncertainty for the rest of this chapter, but this thesis will often come up in chapters 2–4.¹³

That ends the clarifications. I said above that, for nearly every utterance, there is an enormous number of truth-conditional contents any language user could easily have believed that utterance to have. I used this kind of consideration to support my claim that it is very unlikely that any two language users believe any one of those utterances to have the same truth-conditional content. Call considerations of that kind *considerations about multiple candidates*. The rest of this section illustrates the wide variety of utterances for which considerations about multiple candidates hold, including utterances of context-dependent and context-independent sentences. Variance gets support from the fact that considerations

¹³ Thanks to Cian Dorr, Ian Grubb, and Stephen Schiffer for helpful discussion of these issues.

about multiple candidates hold for utterances of those kinds.

1.2.1 Context-dependent sentences

Many sentences are widely recognized to be context-dependent, in that the truth-conditions of utterances of those sentences depend on contextual factors. Examples include sentences like ‘Carla is here’, ‘John is rich’, and ‘Anna is tall’. From a purely formal perspective, the linguistic meaning of such sentences is often represented as a function from a context of utterance to a possible-worlds proposition (Kaplan 1989). In turn, contexts of utterance are represented as sequences of values for the various parameters a sentence’s content—i.e. the proposition it expresses—could depend on. The most common parameters are an agent, a possible world, and a time (Kaplan 1989). Other parameters philosophers and linguists postulate in the formal analysis of natural language constructions include: domain restrictions (Stanley and Szabó 2000), standards of knowledge or justification (DeRose 1992), optional complements for adjectives like ‘ready’ (Bach 1994), etc. An utterance’s truth-conditional content is the proposition resulting from providing the linguistic meaning of the sentence that utterance is an utterance of with the context the utterance occupies.

Contexts of utterance understood as sequences of values for certain parameters should not be conflated with the concrete circumstances in which a sentence is uttered.¹⁴ In most concrete circumstances in which an utterance takes place, there is a huge number of equally natural or plausible possible values for the different parameters that determine the semantic content of a context-dependent sentence, each of which corresponds to a different context of

¹⁴ Though Kaplan (1989) introduced a notion of a proper context that some may confuse with a concrete circumstance of utterance, Kaplan himself distinguished contexts of utterance understood as sequences of values from concrete circumstances in which a sentence is uttered. See MacFarlane (2014, 2005), Predelli (2013) for elaboration on the distinction between contexts of utterance understood as sequences of parameters and concrete circumstances in which a sentence is uttered. See Vision (1985), Predelli (1998), Cohen (2013), Michaelson (2014) for discussion and objections to Kaplan’s notion of a proper context.

utterance formally understood, and each of which determines a different truth-conditional content for a given utterance of that sentence when given as input to the sentence's linguistic meaning. Thus, in most circumstances, there will be a huge number of non-equivalent, equally natural or plausible truth-conditional contents any language user could easily believe an utterance of that sentence to have. Given the vast number of such truth-conditional contents, it would be extremely unlikely for any two speakers to believe the utterance to have exactly the same one.

The following examples illustrate this point. What the examples show is that there is a wide variety of utterances of context-dependent sentences for which considerations about multiple candidates hold. In turn, the variety of context-dependence sentences for which considerations about multiple candidates hold supports Variance.

Location adverbs. The contextual parameter that determines a denotation for a location adverb (e.g. 'here' or 'there') is usually taken to be a location—i.e. a spatial region. For most utterances in which those adverbs occur, there are enormously many (perhaps uncountably many) different spatial regions that could plausibly be given as values to the location parameter of the uttered sentence. Together with that sentence's linguistic meaning, those different spatial regions determine different truth-conditional contents for an utterance of that sentence. Given the vast number of equally plausible truth-conditional contents there are for an utterance of that sentence, it is extremely unlikely that any two language users agree on that utterance's truth-conditional content.

For example, suppose Anna is talking to John on the phone and hears him say 'Carla is here'. Because earlier Anna agreed to meet with John outside of Great Jones Cafe, Anna believes that John's utterance has the truth-conditional content that Carla is somewhere outside the entrance to Great Jones. Now, there are many areas that could count as outside the entrance to Great Jones: the area that extends two meters outside the entrance, the area

that extends three meters outside the entrance, the area that extends four meters outside the entrance, and anywhere in between (among others). Since none of those regions is a more natural or plausible referent for ‘here’ than the others, it is unlikely that there will be some region such that both Anna and John believe that John’s utterance has the truth-conditional content that Carla is in *that* region.

Gradable adjectives. One fruitful line of research in linguistics analyzes the positive form of gradable adjectives (e.g. ‘tall’, as opposed to ‘taller’) as 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 ‘tall’ is taken to express the property of having a degree of tallness that is at least as great as a contextually determined standard of tallness; the predicate ‘expensive’ is taken to express the property of having a degree of cost that is at least as great as a contextually determined standard of cost; etc. For any gradable adjective, there are many cutoff points along the scale of the property they measure, such that objects that have the property to at least as great a degree as the cutoff count as instances of the predicate.¹⁵ Given the vast number of contextual standards that could plausibly determine the extension of any gradable adjective in a concrete situation in which that adjective is used, it is unlikely that any two language users coincide in taking the exact same standard to determine the truth-conditional content of an

¹⁵ According to Kennedy (2007, section 3), this is not true of so-called “absolute” gradable adjectives, which Kennedy takes to come with fixed standards. For example, 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’m 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.

utterance involving that adjective. For example, it is unlikely that there is any standard of tallness such that both Anna and John believe that a certain utterance of ‘Carla is tall’ has the truth-conditional content that Carla has a degree of tallness above that standard.

Other context-dependent predicates. Consider the predicate ‘is ready’. The truth-conditional contents of utterances of sentences in which ‘is ready’ occurs seem to depend on some contextually-determined parameter, such as an event or an action that the predicate’s subject is ready for. Or take the verb ‘arrives’. The truth-conditional contents of utterances of sentences in which ‘arrives’ occurs seem to depend on a contextually-determined location at which the verb’s subject arrives.¹⁶ Other examples are ‘finished’ and ‘prefers’. Partee (2004) claims that ‘local’, and ‘approaches’ are also context-dependent, since the truth-conditional contents of utterances involving them depend on a contextually determined location. Prior (1985) and Choi (2008) argue that dispositional predicates such as ‘is fragile’ or ‘is soluble’ are context-dependent, since the truth-conditional contents of utterances in which they occur depend on contextually determined background-conditions under which the subjects of the dispositional predicate would exhibit the disposition’s characteristic manifestation—for example, the contextually-determined background conditions relevant to the truth-conditions of ‘glass is fragile’ may include the temperature (or range thereof) at which hitting glass with a light force would break it.

In my discussion of location adverbs I explained how in most cases there will be many different locations relevant to the truth of an utterance involving such adverbs, and that the huge number of those locations makes it unlikely that two language users would take the exact same location to play a role in determining the truth-conditional content of utterances in which location adverbs occur. Similar considerations apply to other contextually deter-

¹⁶ Predicates like ‘is ready’ and ‘arrives’ are sometimes called *incomplete*. As far as I know, the first author to call these predicates incomplete was Bach (1994). I think the name ‘incomplete predicate’ seems too general to capture what is special (if anything) about these predicates, but I will continue to refer to those predicates in that way.

mined elements on which the truth-conditional content of an utterance involving predicates like ‘is ready’ or ‘local’ may depend.

Here is my conclusion from the discussion so far. I offered various examples of utterances of context-dependent sentences for which considerations about multiple candidates hold. This supports my claim that considerations about multiple candidates hold for a great variety of utterances, which in turn supports Variance. Before continuing, I should note that even if considerations of multiple candidates held only for utterances of one of the three kinds of sentences I discussed in this subsection (say, sentences involving gradable adjectives), that would be enough to put pressure on the standard picture of communication. For it would show that there is a great number of seemingly normal interactions (e.g. interactions involving utterances of sentences involving gradable adjectives) in which communication is not achieved by way of having shared beliefs about the truth-conditional content of the utterances made in the course of the interaction.

1.2.2 Context-independence

The case for Variance is not limited to utterances involving context-dependent expressions. Assuming that a word like ‘quesadilla’ is not context-dependent, Anna’s utterance of ‘Carla likes quesadillas’ already illustrates this, but I want to offer two more examples before moving on.

Suppose Anna and John are discussing Carla’s hobbies and Anna utters:

(7) Carla has run or will run at least once in her lifetime.

The verb ‘to run’ is not usually taken to be context-sensitive; nevertheless, there are many similar, equally natural, truth-conditional contents any language user could easily have taken Anna’s utterance to have. For example, 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.¹⁷ 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 huge 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 is very unlikely that any two language users believe that utterance to have exactly the same truth-conditional content.

Here is the last example. There are many different, equally natural, ways to draw the line between things that are called ‘dogs’ and those that are not. For example, someone may believe that some of the dog-like creatures first domesticated around 36,000 years ago are called ‘dogs’; others may refuse to call such creatures ‘dogs’, but believe that the dog-like domestic creatures from around 14,000 years ago are called ‘dogs’; etc. Because of these different ways of drawing the line between things that are called ‘dogs’ and things that are not, there are many different truth-conditional contents any language user could easily have believed an utterance of ‘Carla has a dog’ to have. A language user could easily have believed that said utterance has the truth-conditional content that Carla has a dog-like creature of the same species as the creatures first-domesticated around 36,000 years ago; that it has the truth-conditional content that Carla has a dog-like creature of the same species as the doog-like domestic creatures from around 14,000 years ago; etc. Due to the huge number of different beliefs about the utterance’s truth-conditional content any

¹⁷ See JBiz et al. (2010) for an example of this kind of disagreement about the definition of the verb ‘to run’.

language user could easily have had, it is very unlikely that any two language users have exactly the same such belief.

Considerations about the great number of different beliefs about an utterance's truth-conditional content any language user may have had are not limited to these examples. As I explained above (p. 11), those considerations generally apply to utterances involving predicates whose application depends on the properties an object has along one or more sufficiently fine-grained dimensions. If the examples above really are examples of context-dependent sentences, they illustrate the variety of context-*independent* terms for which those considerations hold. But even if those examples turn out to be context-dependent, the case for Variance will apply to any context-*independent* utterances whose truth depends on an object's properties along one or more sufficiently fine-grained dimensions. In turn, such instances support the claim that even conversations involving utterances of a wide variety of context-*independent* sentences exemplify Variance.

1.3 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., it assumes 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 the lack of perfect correlation in people's beliefs I have assumed in stating the case for Variance, consider the following example. Suppose there are two 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 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 correla-

tion 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.¹⁸ 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 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.¹⁹

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

¹⁸ 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 responsibility, but it shouldn't be difficult to see that the same remarks apply to implementations using the causal notion of responsibility. 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.

¹⁹ Thanks to Cian Dorr and Ian Grubb for discussion.

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, in the majority of cases, the participants in typical conversations can come to have the same beliefs about the truth-conditional content of the utterances they make.

1.3.1 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 *Carla likes tortillas filled with between 19 and 139g of cheese*, *Carla likes tortillas filled with between 19.5 and 139.5g of cheese*, and so on, is a more natural candidate for being the truth-conditional content of Anna's ut-

terance of ‘Carla likes quesadillas’.

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.²⁰ 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

(8) there is water in the glass

than that it has the truth-conditional content

(9) 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 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 properties²¹ It is very unlikely, for example, that any one of (15), (16) and (17) 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

²⁰ See e.g. Lewis (1983), Sider (2011), among others.

²¹ Terms 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.

will take to be the truth-conditional content of Anna's utterance of 'Carla likes quesadillas'. More generally, an appeal to 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

(10) there is a dog

than that it is

(11) 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

(12) there is a dog-shaped thing

than that it is

(13) 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 (10) is true than that it is

(14) 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 (10) or (12) than that it has the truth-conditional content (11), (13), or (14), (10) and (12) are psychologically more natural than (11)–(14).

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 pertain 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 quesadillas' to have. For example, the propositions

(15) Carla likes tortillas filled with between 20 and 120g of cheese

(16) Carla likes tortillas filled with between 18.5 and 137.875g of cheese

(17) Carla likes tortillas filled with between 21 and 141g of cheese

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 (15)–(17) (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*.

1.3.2 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.²² That common ground is itself supposed to determine the content of an assertion,²³ 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.²⁴

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 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 be-

²² 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 (1974) and subsequent work.

²³ 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)

²⁴ 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.

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.²⁵

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.²⁶

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-dependence sentences they utter, but I will set that worry aside throughout my discussion.²⁷ What I want to point out is that, for most

²⁵ See Stalnaker (1999) 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." (1999, p.10)

²⁶ 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.

²⁷ 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-

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.²⁸ Just like there is an enormous number of equally natural candidates for the truth-conditional content of most utterances, there is an enormous number of 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

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.

²⁸ I say 'if' because, for reasons related to the discussion of *Uncertainty* (see above, pp. 13–16), it is unlikely that ordinary speakers make presuppositions rich enough to determine such truth-conditional contents.

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 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 sentences involving gradable adjectives. I believe similar remarks apply to the other cases of context-sensitivity I discussed in section 2.1.

The remarks 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.²⁹

1.3.3 Social Externalism

Social externalism is the view that differences in an individual's social environment may produce differences in the contents of that individual's thoughts, beliefs, and so on.³⁰ For instance, according to social externalism, people with the same internal states—e.g. people with exactly the same brain states—may have thoughts with different contents on account of belonging to different communities. To use one of Burge's famous examples (see 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 somewhat different community, one in which the word 'arthritis' denoted *tharthritis* rather than arthritis, her thought would have been about tharthritis and not about arthritis.

²⁹ See Lederman (forthcoming) for further challenges to the notion of common knowledge.

³⁰ The most prominent defender of social externalism is Burge (1979, 1986), but the general idea may be traced back to Putnam (1975). Williamson (2007) constantly appeals to social externalism in his attacks against analyticity. See Ludlow (1995, 1997), Pollock (2015), Wikforss (2001) for criticism of social externalism.

Some opponents of Variance 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. Those opponents' line of thought may go 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, the opponents of Variance may claim, if two people belong to the same linguistic community, they are very likely to have the same beliefs about the truth-conditional content of any given utterance 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.

This line of reasoning overestimates the extent to which social environment can determine what we believe, even if social externalism is true. If social externalism is true, the linguistic community a person belongs to plays a role in determining her beliefs by determining the content of her internal states (e.g. brain states); however, even if social externalism is true, different members of the same linguistic community may have different beliefs on account of having different internal states (e.g. brain states). For example, even if Anna and John belong to the same community, they may have different beliefs about whether arthritis is curable (i.e. one may think that it is curable, and the other that it is not) 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, even if social externalism is true, the fact that two people belong to the same linguistic community does not by itself make it likely that they have the same beliefs, 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 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 .³¹ 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 she has arthritis in her thigh. Accordingly, social environment plays a role in determining someone's beliefs about an utterance's truth-conditional content by determining the content of the mentalese sentences about that utterance's truth-conditional content tokened in that person's 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 anyone's belief-box. Yet, for any natural-language utterance, there are as many mentalese sentences expressing different propositions about that utterance's truth-conditional content as possible truth-conditional contents for the natural-language utterance; given the huge number of such mentalese sentences, it is unlikely that any two people will token exactly the same one in their respective belief-boxes. So 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. Individ-

³¹ See Field (1978) and Fodor (1987).

uate 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). Finally, 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.

For example, suppose that Anna and John are members of the same linguistic community, that Anna utters (7)—‘Carla runs’—in the course of a conversation with John, that Anna associates the English words ‘Carla’ and ‘runs’ with the mentalese words ‘C’ and ‘r’,³² respectively, and that John associates the same English words with the mentalese words ‘C’ and ‘r’, respectively. According to this view, Anna will token the mentalese sentence

(18) Anna’s utterance of ‘Carla runs’ has the truth-conditional content that C r
in her belief-box, whereas John will token the mentalese sentence

(19) Anna’s utterance of ‘Carla runs’ has the truth-conditional content that C’r’.

³² Expressions written in this font are expressions of mentalese.

Since Anna and John belong to the same linguistic community and the mentalese sentence types ‘Cr’ and ‘C’r’ are both counterparts of the English sentence ‘Carla runs’, tokens of (18) and (19) in Anna and John’s respective belief boxes will express the same proposition. In turn, this means that Anna and John will have the same beliefs about the truth-conditional content of Anna’s utterance of (7).³³

The present strategy may seem appealing when it comes to context-*independent* sentences, but it lacks the resources to resist the case for Variance when it comes to context-dependence. We can state the problem 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 ‘Carla is tall’ (regardless of the circumstances in which the utterances are made) as tokens of the same type, the present view is at odds with context-dependence. In particular, 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 expression-types of natural language are individuated so as to count utterances of ‘Carla is tall’ with different truth-conditional contents as tokens of different types, the present strategy does little to resist considerations from multiple candidates. Say for example that we individuate sentence-types so as to count tokens of ‘Carla is tall’ evaluated with respect to different Kaplanian contexts as tokens of different types.

³³ 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).

Given the enormous number of different sentence-types ‘Carla is tall’ could be disambiguated between, it is unlikely that Anna and John token content statements in their respective belief-boxes linking an utterance of ‘Carla is tall’ with mentalese counterparts of the exact same disambiguation. The same holds for the other context-dependence sentences susceptible to multiple-candidate considerations, e.g. sentences involving gradable adjectives.

To make the point more concrete, suppose there are standards of tallness $t_0, t_1, \dots, t_{10000}$, each of which determines different truth-conditional contents for an utterance U of ‘Carla is tall’ 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 ‘Carla is tall’ could be disambiguated between, each with different truth-conditions. Given the number of sentence-types ‘Carla is tall’ 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

(20) Utterance U of ‘Carla is tall’ 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 ‘Carla is tall’. The same goes for other context-dependent sentences.

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 and other popular views (see introduction to this chapter). 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.

1.4 Three kinds of communication-likeness

I said in the introduction that, 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 sentences they utter. 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 chapter, 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 sentences they utter)—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. Chapters 2 and 3 offer more detailed discussions of the first two of those ways for a conversation to be communication-like.

Consider the following example. Anna invited John to her house, but John doesn't have the exact address; however, he does know the house's block. As it happens, there are only two houses in Anna's block: Anna's house, which is color 4, and another house, which is color 1. Anna and John both believe that the word 'green' is context-*independent*, but 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 4.1). When John asks Anna how he can identify her house, she utters:

(1) It is green.

For the sake of simplicity, suppose that Anna and John both believe that ‘It’ refers to *H*, Anna’s house.³⁴ Because of their respective beliefs about the meaning of the word ‘green’, Anna and John disagree about the truth-conditional content of Anna’s utterance. Anna believes that her utterance has the truth-conditional content that *H* is one of colors 2–5, whereas John believes that it has the truth-conditional content that *H* is one of colors 3–6.



Figure 1.1: Some colors

Despite this difference in Anna’s and John’s beliefs about the truth-conditional content of Anna’s utterance, Anna and John’s conversation resembles ideal cases of successful communication in the following ways. First, through her utterance of (1), Anna transmitted to John information about the color of her house: that it is neither color 1 nor color 7, that it 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 (1).

Second, given the circumstances in which Anna and John’s conversation takes place, they can have agreements and disputes about the truth of Anna’s utterance that are not purely verbal. Such agreements and disputes are not purely verbal insofar as the same fact is responsible for the truth of the utterance by both Anna’s and John’s lights: by both of

³⁴ 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.

their lights, the utterance is true in virtue of the fact that Anna's house is color 4. This is evidenced by the fact that if John disputed the truth of Anna's utterance, Anna could settle the issue by pointing to the color of the house in the circumstances in which the interaction takes place as the fact responsible for her utterance's truth. In that case, John could not reasonably reject Anna's defense by replying something like 'Wait a moment! The fact that the house is that color does not make what you said true!' or 'Wait a moment! The fact that the house is that color doesn't make it green!'.

Third, given the circumstances in which the conversation takes place, that interaction puts John in a position to identify Anna's house. This is not a mere coincidence: 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 the utterance has the truth-conditional content that Anna's house is one of colors 3–6 or that it has the truth-conditional content that Anna's house is one of colors 2–5. Furthermore, 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 a conversation to be communication-like can all be explained in terms of similarity between propositions. According to this view, it is because the propositions

(2) *H* is one of colors 2–5 and

(3) *H* is one of colors 3–6

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

I find this view dissatisfying. (2) and (3) 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, (2) and (3) 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 (1). In that scenario, however, (2) would have been true and (3) and false. Thus, at least in principle, Anna could have known that the proposition she took to be her utterance's truth-conditional content was true, whereas John could have known that the proposition *he* took to be the truth-conditional content of Anna's utterance was false. This suggests that Anna and John's dispute would have been merely verbal in the new scenario: Anna and John could each have known the facts they took Anna's utterance to be about. Hence, since (2) and (3) 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's dispute about the truth of Anna's utterance not to be merely verbal.

Nor does the similarity between (2) and (3) 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 Anna's utterance of (1) 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 Anna's utterance is true will not put John in a position to reach Anna's house (given his beliefs about the utterance's truth-conditional content): if John believes that Anna's utterance is true, he will go to the house color 6, which is not Anna's house. Unlike in the original example, the differences between (2) and (3) *do* matter for the purposes of reaching Anna's house in

the new scenario. Since Anna's house is in fact color 2, if John had believed that Anna's utterance has the truth-conditional content that H is one of colors 2–5, believing that that utterance is true would have put him in a position to reach Anna's house. Yet (2) and (3) are as similar in this new scenario as they were in the original one.

In future chapters I will discuss the first two of these ways for a conversation to be communication-like in more detail. For the time being, I want to observe that the fact that the three communication-like features I just introduced 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-conditional content 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 sentences they utter, conversations between them will tend to have all or most of the communication-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.

I will discuss pluralism towards the end of the dissertation, in chapter 5. For the time being, it is worth thinking of pluralism as the overarching picture that ties the discussion in the next three chapters together.

1.5 Conclusion

In this chapter I argued for *Variance*, the thesis that for nearly every utterance and any two language users, there is no proposition those two language users believe to be that utterance's truth-conditional content. The case for *Variance* relied on the observation that, for most utterances, there is an enormous number of different truth-conditional contents any language user could easily have believed that utterance to have. This observation makes it extremely unlikely that any two language users believe a given utterance to have the same truth-conditional content. I supported this line of reasoning with examples of utterances of context-dependent and context-independent sentences.

In the introduction I mentioned some of the consequences of *Variance* to debates about the nature of communication and indirect speech reports. The rest of my dissertation discusses those consequences. Chapter 2 offers an account of the transmission of relevant information, given the truth of *Variance*. Chapter 3 offers an account of the epistemic significance of ordinary disputes capable of accommodating *Variance*. Chapter 4 offers an account of speech reports that explains how we can know which speech reports are true despite the truth of *Variance*. Chapter 5 concludes by relating those various accounts to a more general picture of the nature of communication and linguistic competence.

Chapter 2

Transmission of relevant information

2.1 Introduction

One of the main aims of linguistic interaction is the transmission of relevant information. For example, if I ask you whether your house is green and you reply ‘it’s green’, you will have given me a piece of relevant information: that your house is green. If, on the other hand, you reply ‘it’s at least 10 ft tall’, you may have given me some information, but that information is irrelevant to what I wanted to know. This chapter accounts for the transmission of relevant information, given the truth of Variance.

We can get a better idea of what’s at issue by comparing two examples. Consider first COLOR:

COLOR. Anna has just bought a new toy and decided to call it ‘Charlie’. Wanting to inform John of Charlie’s color, she utters:

(1) Charlie is green

As it happens, Anna and John disagree about the meaning of the word ‘green’:

Anna believes that it applies to things of colors 2–5, whereas John believes that it applies to things of colors 3–6 (see figure 4.1).



Figure 2.1: Some colors

Because of their disagreement about the meaning of ‘green’, Anna and John also disagree about the truth-conditional content of Anna’s utterance. Anna believes that her utterance’s truth-conditional content is

(2) *Charlie is one of colors 2–5,*¹

whereas John believes that it is

(3) *Charlie is one of colors 3–6*

instead. Despite their disagreement about (1)’s truth-conditional content, Anna could still have transmitted information to John about Charlie’s color. For example, Anna could have transmitted to John the information that Charlie is not red, that he is not colors 1 or 7, etc. Though this information does not exhaust the information Anna originally intended to transmit, it is entirely about the subject matter Anna intended to inform John about—i.e. Charlie’s color.

Compare this example with HEIGHT:

HEIGHT. Anna has just bought a new toy and decided to call it ‘Charlie’.

Wishing to inform *Carla* of Charlie’s color, she utters

(1) *Charlie is green*

¹ I will use expressions in *this* font as names for semantic contents, such as propositions, properties, etc.

As it happens, Anna and Carla disagree about the meaning of the word ‘green’. Anna believes that it applies to things of colors 2–5, whereas, due to an idiosyncratic upbringing, Carla believes that it applies to things that are at least half a meter tall.

As a result of their beliefs about the meaning of the word ‘green’, Anna and Carla disagree about the truth-conditional content of Anna’s utterance. Anna believes that her utterance’s truth-conditional content is (2), whereas Carla believes that it is:

(4) Charlie is at least half a meter tall.

Because of this difference, Anna did not transmit any information about Charlie’s color to Carla through her utterance of (1): by Carla’s lights, Anna’s utterance is entirely about Charlie’s height, and not at all about Charlie’s color.

Pretheoretically, we can describe the difference between the two cases as follows. (2) and (3) both exclude some of the same possibilities pertaining to Charlie’s color: that Charlie is not color 1, that it is not color 7, and so on. The information about Charlie’s color that Anna transmitted to John is just the information that those possibilities do not obtain. On the other hand, (4) does not exclude any possibility pertaining to Charlie’s color: for all the truth of (4) requires, Charlie can be color 1, color 7, or any other color. Thus, there is no possibility pertaining exclusively to Charlie’s color that both (4) and (2) exclude; in that sense, there is no information exclusively about Charlie’s color that they both entail.

I will discuss these examples in more detail later on. For the time being, I only want to make the point that an adequate account of the transmission of relevant information should distinguish between cases like COLOR and cases like HEIGHT. The aim of the present chapter is to offer such an account. The structure of the chapter is as follows. I start

by arguing that a standard, broadly Gricean account of communication can't explain the phenomenon at hand, nor can the account be straightforwardly modified so as to explain it (sections 2 and 3). Then I present a new account (sections 4, 5, and 6), discuss further issues and potential objections (section 7), and conclude (section 8).

2.2 A standard picture of communication

I said in the previous chapter that, according to the standard picture, communication is usually achieved thanks to our shared knowledge of the truth-conditional content of the utterances we make. This section elaborates on the problems Variance raises for the standard picture of communication. To do so, it will be helpful to focus on a popular implementation of the standard picture: Stalnaker's common ground picture.

According to Stalnaker (1978, 1974), in every conversation there is a *common ground*: a set of propositions all the participants presuppose for the purposes of the conversation, that they all presuppose that they presuppose, that they all presuppose that they presuppose that they presuppose, and so on. According to this picture, by making an assertoric utterance the speaker makes a proposal to add information (i.e. a proposition) to the common ground. If all the participants in the conversation accept that proposal, the information is added to the common ground, and the conversation can go on. On this picture, what it is for a speaker to transmit information to her audience through a given utterance is for there to be a proposition which (a) she proposes adding to the common ground, (b) is not entailed by the information already in the common ground, (i.e. is genuinely new information) and (c) her audience accepts adding to the common ground. Typically, the information a speaker proposes to add to the common ground through a literal assertoric utterance is the

proposition that speaker takes to be the truth-conditional content of her utterance.²

According to Stalnaker, the common ground will include substantive information about the world (e.g. that California is part of the United States, or that water makes things wet), but it should also contain information that puts the participants in a conversation in a position to know what information a speaker intends to add to the common ground through an utterance.³ In cases of literal assertoric utterances, that information will be the information that the utterance's truth-conditional content is such and such—or, at the very least, information that puts the participants in the conversation in a position to know that the utterance's truth-conditional content is such and such.

Here is an example. Forget about Variance for a moment and suppose Anna utters (1) ('Charlie is green') wishing to inform John about Charlie's color. According to Stalnaker, by uttering (1) Anna made a proposal to add the proposition that Charlie is green to the common ground. John will know that Anna made this proposal thanks to its being common ground that: (a) it is common ground that (1)'s truth-conditional content is the proposition that Charlie is green, and (b) Anna is speaking literally. If John accepts Anna's proposal, the proposition that Charlie is green will be added to the common ground and, in turn, Anna will have informed John that Charlie is green.

Let's see how this picture interacts with Variance. Recall that Variance is the thesis that for almost every utterance and any two language users, there is no proposition that they both believe to be that utterance's truth-conditional content. Absent such a proposition, the

² According to standard implementations of this picture, if two propositions are equivalent modulo the common ground, then one proposes adding one iff one proposes adding the other.

³ 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)

participants in a conversation will often not be in a position to tell what information the speaker is proposing to add to the common ground through a literal assertoric utterance.

For example, consider once more Anna and John's conversation in COLOR. In that conversation, Anna uttered (1) thinking that her utterance's truth-conditional content was (2), whereas John thought the truth-conditional content of Anna's utterance was (3).⁴ Due to this discrepancy, by uttering (1) Anna is proposing to add to the common ground the proposition that Charlie is one of colors 2–5, whereas John takes Anna to propose adding to the common ground the different proposition that Charlie is one of colors 3–6.

Despite this discrepancy, standard theorists may still account for the fact that Anna transmitted *some* information to John.⁵ Standard theorists will claim that proposing and accepting to add information to the common ground is closed under entailment: if someone proposes or accepts to add a proposition to the common ground, she must thereby propose or accept to add that proposition's logical consequences which are not yet part of the common ground. Since the disjunction of (2) and (3) is a logical consequence of both (2) and (3), both Anna and John are committed to thinking that Anna proposed to add said disjunction to the common ground. Thus, if John accepts what he takes to be Anna's proposal, and the disjunction of (2) and (3) is not entailed by the propositions already in the common ground, Anna and John will both believe that the proposition that Charlie is either one of colors 2–5 or one of colors 3–6 has been added to the common ground. That,

⁴ Repeated here:

- (1) Charlie is green
- (2) Charlie is one of colors 2–5
- (3) Charlie is one of colors 3–6

⁵ It might be worth noting that the discrepancy in what Anna and John take to be Anna's proposal will itself be worrisome for defenders of the standard view such as Stalnaker, who insists that the participants in a conversation must all agree on what particular proposal is being made through an utterance (see footnote 3). As COLOR illustrates, if Variance is true, this requirement is rarely satisfied.

according to the present picture, will be the information Anna transmitted to John through her utterance of (1).

More generally, according to the standard view, if a speaker proposes to add the proposition P to the common ground, her audience takes her to propose to add the different proposition Q , and the audience accepts what she takes to be the speaker's proposal, the speaker will have transmitted to her audience the information that $P \vee Q$ —provided that this proposition is not already entailed by the information in the common ground. This is so regardless of how intuitively different P and Q are. Thus, despite the truth of Variance, the standard picture can still explain the transmission of information.

The problem for the standard view is that it doesn't distinguish between cases like COLOR and cases like HEIGHT. According to the standard view, both conversations are cases of information transmission. Yet, as I claimed in the introduction, the information Anna transmitted to John through her utterance of (1) is entirely about Charlie's color, whereas the information Anna transmitted to *Carla* through a similar utterance is not. Pretheoretically, the difference is that, whereas the disjunction of (2) and (3) excludes certain possibilities pertaining exclusively to Charlie's color—e.g. that Charlie is color 1 or color 7—the disjunction of (2) and (4) can be true regardless of Charlie's color.

In a way, the fact that the standard picture fails to distinguish between COLOR and HEIGHT should not come as a surprise. Even if we set Variance aside, the standard picture does not distinguish between transmission of *relevant* information and transmission of *irrelevant* information. For example, let's set Variance aside for a moment and suppose John prompts Anna's utterance of (1) by asking:

(5) Is Charlie green?

If Anna replies to John's question by uttering (1) ('Charlie is green'), she will have given John information relevant to what John wanted to know. If, on the other hand, Anna had

replied

(6) Charlie is more than a meter tall,

she may have given John some information (e.g. that Charlie is more than a meter tall), but that information would have been irrelevant to their conversation. The standard account doesn't distinguish between the two cases: it counts both conversations as cases of information transmission, but lacks the resources to explain why the information Anna transmitted to John through her utterance of (6) is irrelevant. The next section discusses a popular solution to this problem.

2.3 The question-based picture

In the contemporary literature, the standard way of addressing the problem of relevance without departing too much from the standard picture is through the notion of a *question under discussion* (See e.g. Roberts 2012, Groenendijk and Roelofsen 2009, Ciardelli et al. 2013). Very generally, the idea is that, in addition to a body of information that the participants in a conversation presuppose (and presuppose that they presuppose) for the purposes of the conversation, there is also a “stack” of questions under discussion; i.e., a series of questions all the participants in the conversation intend to answer, and of which it is common ground that they intend to answer them. According to the present approach, if assertoric utterances are proposals to add information to the common ground, interrogative ones are proposals to add a new question to the stack of questions under discussion. Typically, according to the present view, the participants in a conversation will be in a position to know what question someone proposes to discuss through an interrogative utterance due to its being common ground that that question is the utterance's semantic content.

Defenders of the present approach take a proposition to be relevant in the context of a conversation just in case it answers a question under discussion, at least in part.⁶ In turn, according to this view, what it is for a speaker to transmit *relevant* information to her audience through an assertoric utterance is for there to be a proposition which (a) she proposes to add to the common ground, (b) is not entailed by the information already in the common ground, (c) is at least a partial answer to a question under discussion in the conversation, and (d) her audience accepts to add to the common ground. For example, according to this view, if Anna utters (1) in response to John's interrogative utterance of (5), what makes the information Anna transmitted to John through her utterance of (1) *relevant* is the fact that that information answers a question under discussion—whether Charlie is green. In contrast, the information Anna transmitted to John by uttering (6) in response to (5) does not answer any question under discussion, which is why it is *irrelevant*.

Call the standard picture plus the present approach to relevance the *question-based picture* of communication. Soon I will argue that, given Variance, the question-based picture can't distinguish between the transmission of relevant information and the transmission of irrelevant information. In order to do that, it is important to present some of the details concerning the implementation of the question-based picture.

2.3.1 Standard treatment of questions

Popular approaches treat questions as partitions of logical space. Intuitively, each of the cells in a partition corresponds to a possible complete answer to the question; that is, a proposition that determines exactly how things are insofar as that question is concerned.

⁶ Roberts (2012) would claim something stronger. She would claim that a proposition P is relevant in the context of a conversation just in case it answers the last question to be added to the stack of questions under discussion at the time P was proposed as an addition to the common ground. The difference between the stronger and the weaker constraints on relevance does not matter for present purposes, since I will argue that even the weaker constraint is almost never satisfied.

For example, suppose Anna and John are the only people there are, and consider the question *who walks?*. The possible *complete answers* to this question are: that both Anna and John walk; that Anna walks but John doesn't, that Anna doesn't walk but John does, and that neither Anna nor John walk. Each of these propositions determines exactly how things are with respect to who walks, and corresponds to a cell in the partition of logical space that puts any two worlds in the same cell just in case they are equivalent with respect to who walks. A partial (and incomplete) answer to the same question is the proposition that Anna walks, and another is the proposition that John doesn't walk; neither of these propositions fully determines how things are with respect to who walks, but each *partly* determines who walks.

Roughly following Roberts (2012), we can determine what question an interrogative expresses by first considering what *alternatives* it makes salient in the context of a given conversation. For example, according to Roberts, in a conversation in which the only people in the contextually determined domain are Anna and John, an utterance of

(7) Who walks?

will make salient the proposition that Anna walks and the proposition that John walks. Those two propositions—that Anna walks and that John walks—are the alternatives an utterance of (7) makes salient in the context of that conversation. Given that set of alternatives, the question expressed by an utterance of (7) is the coarsest-grained partition of logical space each of whose cells determines a truth-value (i.e. truth or falsity) for each alternative (7) makes salient. That is, the partition whose cells are: the set of possible worlds in which both Anna and John walk, the set of possible worlds in which Anna walks but John doesn't, the set of possible worlds in which Anna doesn't walk but John does, and the set of possible worlds in which neither Anna nor John walks.⁷

⁷ The present approach differs from Roberts' in the following way. Roberts takes the question expressed

More generally, the set of alternatives for a *wh*-interrogative is the set of propositions obtained by replacing the question's *wh*-element—e.g. 'who' or 'what'—with an element in the contextually determined domain of the same type.⁸ That is, where '*wh*' stands for a *wh*-element (e.g. 'who', 'what'), β is an expression with semantic content of type τ , σ is a type that can combine with τ to form a proposition, and D is a contextually determined domain, the set of alternatives of *wh* β ? is the set of propositions:

$$(8) \text{ALT}(\text{wh } \beta?) = \{p \mid \exists u \in D_{\sigma}. p = \llbracket \beta \rrbracket(u)\}$$

In turn, the question expressed by *wh* β ? is the partition of logical space each of whose cells determines a truth-value for each proposition in $\text{ALT}(\text{wh } \beta?)$.⁹

Polar questions—questions whose answers are yes or no—are determined in a similar way. Where β is a sentence, the only alternative the polar question $\beta?$ raises is the proposition semantically expressed by β . For example, the only alternative raised by the polar question 'Is Charlie green?' is the proposition that Charlie is green. Accordingly, the question expressed by 'Is Charlie green' is the partition of logical space into two sets: the set of possible worlds in which Charlie is green, and the set of possible worlds in which

by an interrogative utterance to be the set of alternatives that interrogative makes salient, whereas, following Groenendijk and Stokhof (1989, 1982), I take the question expressed by an interrogative utterance to be a *partition* of logical space each of whose cells determines a truth-value for each alternative the interrogative makes salient. The reason for this choice is twofold: on the one hand, taking questions to be partitions of logical space allows for a more straightforward definition of answerhood; on the other, the use of alternatives will make it easier to see how certain theoretical options are available for question-based theorists in response to objections from Variance (see below, section 3.3) The treatment of questions as sets of propositions goes back to Hamblin (1973).

⁸ A domain here will include objects of various types. In particular, it will include objects of types e (i.e. individuals), type t (i.e. truth-values), s (i.e. possible worlds) and, in general, where σ and τ are types, objects of type $\langle \sigma, \tau \rangle$. For example, in addition to objects of type e like Anna and John, a standard domain will include properties, which we can think of as objects of type $\langle s, \langle e, t \rangle \rangle$.

⁹ See Roberts (2012, pp. 6:9–13) for details. See Groenendijk and Stokhof (1982) for an equivalent way of determining the question expressed by an interrogative, using equivalence relations rather than alternatives as the basis. See Groenendijk and Roelofsen (2009), Ciardelli et al. (2013) for recent work on questions that takes questions to be sets of downward closed propositions, rather than partitions of logical space. The differences between the downward-entailment approach and the partition approach don't matter for present purposes.

Charlie is not green.

Given this way of understanding questions, we can define partial answerhood as follows:

Partial answerhood: Where P is a possible-worlds proposition, and $Q?$ is a question, P is a partial answer to $Q?$ just in case P is the union of the sets in a non-empty *proper* subset of $Q?$.¹⁰

For example, assuming that Anna and John are the only individuals in the domain, the proposition that Anna walks is a partial answer to *who walks?* because it is equivalent to the union of two of the possible complete answers to that question: that Anna and John both walk, and that Anna walks but John doesn't. On the other hand, the proposition that either Anna or John walk, or neither does, is not a partial answer to the question: though that proposition is the union of a subset of *who walks?*, that subset is *who walks?* itself (hence, not a *proper* subset of *who walks?*).

Setting Variance aside, in a case in which John asks (5) ('Is Charlie green?') and Anna replies by uttering (1), Anna gives relevant information to John because the proposition Anna proposed to add to the common ground (i.e. *Charlie is green*) is a complete answer to the question John asked (i.e. *Is Charlie green*), which is now a question under discussion. In contrast, if John asks (5) and Anna replies by uttering 'Charlie is tall', the information Anna gave to John is irrelevant because the proposition Anna proposed to add to the common ground (i.e. *Charlie is tall*) does not answer John's question at all—not even partially.

I will now argue that, given the considerations that support Variance, the question-based picture cannot distinguish between the transmission of relevant and irrelevant information.

¹⁰ Note that, if P is a complete answer to $Q?$ and $Q?$ has more than one cell, then P is a partial answer to $Q?$.

2.3.2 Troubles with Variance and Question-variance

I said above that a question under discussion is one which all the participants in the conversation intend to answer, and believe that every participant in the conversation intends to answer. According to the view I've been discussing, the participants in a conversation can agree on what questions to discuss due to their common knowledge of the questions expressed through interrogative utterances. Yet the same observations that support Variance support *Question-variance*, the further thesis that, for most interrogative utterances, there is no question which more than one language user takes that utterance to express. If this is so, it's difficult to see how all the participants in a conversation could come to agree on what questions to discuss. And, as I will explain now, if there is no question under discussion in the context of a conversation, the present account predicts that no relevant information is transmitted.

We can illustrate the problem through the following example. Consider again a case like COLOR, but this time suppose John prompts Anna's utterance of (1) by asking

(5) Is Charlie green?

There is a huge number of different areas of the color space Anna and John could easily have taken to be denoted by the word 'green'. Given the huge number of such areas, it would be extremely unlikely that Anna and John took 'green' to denote the exact same one. And if Anna and John disagree as to what area of color space is denoted by 'green', they will also disagree as to what question (5) expresses.

If, as in COLOR, Anna takes 'green' to apply to things of colors 2–5 and John takes it to apply to things of colors 3–6, then John will take himself to have asked the question

(9) Is Charlie one of colors 3–6?,

whereas Anna will take him to have asked

(10) Is Charlie one of colors 2-5?

Since (9) and (10) are different questions—i.e. they correspond to different partitions of logical space—there is no question both Anna and John take the latter to have proposed to discuss. Absent further explanation as to how there could be some question Anna and John both intend to answer (and believe they intend to answer), there is no such question.

Now, I said above that, according to the present picture, what it is for a speaker to transmit *relevant* information through her audience through an assertoric utterance is for there to be a proposition which (a) she proposes to add to the common ground, (b) is not entailed by the information already in the common ground, (c) is at least a partial answer to a question under discussion in the conversation, and (d) her audience accepts to add to the common ground. If there is no question under discussion in a given conversation, clause (c) of the present definition fails; accordingly, the present view predicts that the information Anna transmitted to John is irrelevant, no matter what information that is.¹¹

This prediction is not limited to cases involving polar questions. For instance, suppose John asks Anna

(11) Which of your toys are green?

Because of the differences in Anna and John's beliefs about the meaning of the word 'green', Anna and John will disagree about what question John asked: Anna will take John to have asked *which of Anna's toys are one of colors 2-5?*, whereas John will take himself to have asked *which of Anna's toys are one of colors 3-6?*. Absent further explanation of how Anna and John could have agreed on what question to discuss, the present view predicts that Anna wouldn't have transmitted any relevant information to John by uttering (1) in the context of the present conversation.

¹¹ Here I'm restricting my attention to relatively specific questions, i.e. questions other than the question 'what are things like?'. I will consider the possibility of appealing to non-substantive questions in the next subsection.

More generally, whenever two people disagree about the semantic content of one of the expressions occurring in an interrogative utterance, they will also disagree as to what question the utterance expresses. Absent further explanation of how there could be some question both speaker and audience agree to discuss, there is no such question. And without a question under discussion, the question-based picture predicts that any transmitted information is irrelevant. Since, given the truth of Question-variance, the participants in a conversation will often fail to agree on what question to discuss, the question-based picture predicts that no relevant information is transmitted in a great number of seemingly normal conversations. This, I take it, is the wrong result.

I will now consider several responses on behalf of the question-based picture, and argue that they fail. Readers already convinced that the question-based picture fails may go straight to section 4, where I start developing my positive proposal.

2.3.3 Potential responses, and why they won't do

I am going to consider three potential responses to the present objection, all of which attempt to determine what question is under discussion in conversations whose participants disagree about what question an interrogative utterance expresses. Respectively, those responses appeal to the notion of a “big question”, to a relation of entailment between questions, and to what Roberts (2012) calls “focal alternatives”. I will argue that none of these lines of thought put the question-based picture in a position to distinguish between cases like COLOR and cases like HEIGHT.

The big question. According to Roberts (2012), in every conversation there is at least one question under discussion, even if that question is just the “big question”—the question ‘what are things like?’, or ‘what is the world like?’. The big question corresponds to the

partition of logical space each of whose cells contains exactly one possible world. The problem is that the big question is too general to distinguish between relevant and irrelevant information.

Suppose for example that Anna is having a conversation with *Carla* rather than with John, and recall that Carla believes that the word ‘green’ expresses the property of being at least half a meter tall. Wishing to know Charlie’s height, she asks Anna

(12) Is Charlie green?

to which Anna replies by uttering (1)—‘Charlie is green’. Given Carla’s beliefs about the meaning of the word ‘green’, Carla will take Anna to propose to add the proposition

(13) Charlie is at least half a meter tall,

whereas Anna will take herself to have proposed adding (2)—the proposition *Charlie is one of colors 2–5*.

According to the present account, through her utterance of (1), Anna transmitted to Carla the information that either Charlie is one of colors 2–5 or at least half a meter tall. That proposition can be true regardless of Charlie’s height (e.g. if it is true in virtue of Charlie being one of colors 2–5), and in that sense it is irrelevant to what Carla wanted to know—i.e. whether Charlie is at least half a meter tall. However, the proposition that Charlie is one of colors 2–5 or at least half a meter tall is a partial answer to the big question—since it is equivalent to the union of a proper subset of the cells in the big question, namely, the cells in which Charlie is one of colors 2–5 or at least half a meter tall. Thus, if the big question is a question under discussion, the present view predicts that Anna transmitted general information to Carla through her utterance of (1).

More generally, if the big question is always a question under discussion, the present view predicts that the speaker always transmits relevant information to her audience, no

matter how grossly the latter misinterprets the former's utterance. Thus, the present view does not make any substantive distinction between relevant and irrelevant information.

Question entailment. In section 2, when I introduced the standard picture of information transmission, I said that proposals to add information to the common ground are closed under logical consequence: anyone who proposes to add P to the common ground thereby proposes to add its logical consequences as well. Motivated by this idea, defenders of the question-based approach may claim that something analogous holds of proposals to discuss a question: anyone who proposes to discuss question $Q?$ thereby proposes any other question $Q?$ entails. Defenders of the present view may claim that, even if speaker and audience disagree as to what question an interrogative utterance expresses, there may still be questions they are both committed to discussing if the utterance is accepted; i.e. the questions entailed by both the question the speaker takes herself to have proposed to add to the questions under discussion and the question the audience takes the speaker to have proposed to add.

The most popular definition of question entailment goes as follows. A question entails another just in case every complete answer to the former entails a complete answer to the latter.¹² For example, the question `Which objects are one of colors 2-5?` entails the question `Is Charlie one of colors 2-5?`, since every proposition that determines exactly which objects are one of colors 2-5 thereby determines whether Charlie in particular is one of colors 2-5.

The present amendment to the question-based picture won't take us very far. Call a question substantive just in case it is a partition of logical space with more than one cell.

¹² As far as I know, this definition of question entailment was first proposed by Groenendijk and Stokhof (1982, p.426). Others, such as Roberts (2012, 6:7), have also endorsed this definition. See Groenendijk and Roelofsen (2009) for a recent discussion of entailment between questions. As far as I can tell, the adoption of the latter's approach to questions and answers does not make a difference for present purposes. In any case, I will leave it to defenders of that view to argue otherwise.

Only substantive questions have partial answers, so only substantive questions can be the basis for distinguishing between relevant and irrelevant information.¹³ Now, two questions $P?$ and $Q?$ have a substantive entailment in common only if there is some proposition that partially answers both $P?$ and $Q?$.¹⁴ The problem is that, for any two different polar questions, there is no proposition that partially answers them both,¹⁵ so no two polar questions have a substantive entailment in common.¹⁶ Thus, in a case in which John proposes to discuss the question `Is Charlie one of colors 2-5?` but Anna takes him to propose discussing the question `Is Charlie one of colors 3-6?`, there is no substantive question they are both committed to discussing (i.e. no substantive question entailed by both of the questions just mentioned).

More generally, if speaker and audience disagree about the meaning of a polar interrogative utterance, they won't be able to agree on what question to add to the stack of questions under discussion on the basis of that utterance. And, as I stated above, without a question under discussion, the present view predicts that whatever information the speaker transmits to her audience in response to a polar interrogative utterance, that information will be irrelevant in the context of their conversation. On the other hand, if the big question is a question under discussion, the present view cannot distinguish cases like `COLOR` from

¹³ Recall that P is a partial answer to $Q?$ just in case P is the union of the sets in a non-empty *proper* subset of $Q?$. Thus, since non-substantive questions don't have proper non-empty subsets, they don't have partial answers either.

¹⁴ This follows from a general fact about partitions: if A and B are partitions of the same set, and every cell in A is a subset of a cell in B (i.e. A is a refinement of B), then every cell in B is the union of a subset of A . Given this fact, if $P?$ and $Q?$ both entail $S?$, and $S?$ is substantive, every complete answer to $S?$ will be the union of a proper subset of $P?$ (hence a partial answer to $P?$) and also the union of a proper subset of $Q?$ (hence a partial answer to $Q?$).

¹⁵ Recall that polar questions divide logical space into exactly two sets. Given that a proposition is a partial answer to a given question just in case it is equivalent to the union of some but not all of that question's complete answers, it follows that the only partial answers to a polar question are its complete answers. Since no complete answer to a polar question is also a complete answer for a different polar question, there is no proposition that partially answers two different polar questions.

¹⁶ Questions here are individuated by their members, so that, e.g. $P?$ and $\neg P?$ are the same question: the partition of logical space that divides it into the set of worlds in which P is true and the set of worlds in which it isn't.

cases like HEIGHT. Needless to say, the problem is not exclusive to polar questions, but I take it that the fact it arises for polar questions is enough evidence against the present view.

Focal alternatives and question congruence. Focus is a kind of linguistic marking of elements within an utterance which convey important information within a conversation.¹⁷ In English, focus is realized by stressing certain elements in an utterance. For example, there are two ways of pronouncing (5), depending on which of ‘Charlie’ and ‘green’ are stressed:

(14) Is CHARLIE green?

(15) Is Charlie GREEN?

(14) and (15) bring to mind different alternatives, and a speaker’s choice of pronunciation will depend on which of those alternatives she wants to make prominent for the purposes of the conversation. For instance, (14) seems to ask whether *Charlie*, as opposed to some other object, is green. On the other hand, (15) seems to ask whether Charlie is *green*, as opposed to, say, blue.

This difference between the alternatives (14) and (15) make salient is cashed out as a difference in their *focal alternatives*: the propositions obtained by replacing the focused element with a variable and then “interpreting the result relative to each member of the set of all assignment functions which vary at most in the values they assign to those variables.” (see Roberts 2012, 6:33). For example, assuming that Charlie, Anna, and John are all the individuals in the contextually determined domain, (14)’s focal alternatives are the propositions: *Charlie is green, Anna is green, and John is green.* In contrast,

¹⁷ See Stevens (2017) for an overview of research on focus. See Rooth (1992) for an influential approach to the semantics of focus.

assuming that the property of being green, the property of being blue, and the property of being at least half a meter tall are the only properties in the domain, (15)'s focal alternatives are the propositions: *Charlie is green, Charlie is blue, and Charlie is at least half a meter tall.*¹⁸

According to Roberts, the use of focus can help determine what question is under discussion in a conversation at the time of an utterance. In particular, she claims that in order for an utterance to be felicitous, its *focal alternatives* must be exactly the same as the alternatives the question under discussion makes salient.¹⁹ Thus, if an utterance is felicitous, we can use its focal alternatives to determine what question is under discussion. Or so the thought goes.

For example. Recall that Anna and John disagreed about what question (5) expressed: Anna thought it expressed the question of whether Charlie is one of colors 2–5, whereas John thought it expressed the question of whether Charlie is one of colors 3–6. The hope is that, despite this disagreement, Anna and John will agree about what (5)'s focal alternatives are, in which case they will agree about what the question under discussion was at the time John uttered (5).

Unfortunately for the present view, Anna and John are no more likely to agree about what (5)'s focal alternatives are than they are to agree about what question (5) expresses. Suppose that John pronounces (5) as (15)—‘Is Charlie GREEN?’. Because the word ‘green’ is focused, the focal alternatives of John’s utterance are the propositions obtained by substituting a property in the contextually determined domain for F in

¹⁸ See above, fn. 8. As of this writing, there is no standard account of how a contextual domain is restricted. See von Stechow (1994), Stanley and Szabó (2000), Schwarz (2012), Barwise and Perry (1981, 1983) for discussion of various issues pertaining to domain restrictions.

¹⁹ Note that a question’s focal alternatives and the alternatives that question makes salient are not the same. The alternatives the polar question *P?* makes salient are just P and \neg P. However, depending on intonation, its focal alternatives may include propositions, e.g. Q or R.

(16) Charlie is F.

The kind of considerations that support Variance cast doubt on the idea that Anna and John will agree on what properties there are in the contextually determined domain. For there are a huge number of different sets of candidate properties, given which it would be extremely unlikely for Anna and John to agree that the exact same set is the set of contextually determined properties. Suppose for instance that the contextually determined domain has only color properties, and that each color property corresponds to a region of the color space. There is a huge number of ways of dividing the color space into different regions. Since each of those ways of dividing the color space corresponds to a different set of properties that could be the set of properties in the contextually determined domain, it would be extremely unlikely for Anna and John to agree on which of those sets of properties is the set of properties in the contextually determined domain. And if Anna and John disagree about what properties are in the contextually determined domain, they will disagree about what the focal alternatives of John's utterance are. Consequently, they will disagree about what question was under discussion at the time of John's utterance.

General questions about an object's properties. As a last resort, defenders of the question-based picture may claim that in a case like COLOR, the question under discussion is something like *what maximally specific color is Charlie?*. This question corresponds to the partition of logical space each of whose cells has all the worlds that are the same with respect to Charlie's maximally specific color (e.g. one cell will contain exactly the worlds in which Charlie is color 1, another the worlds in which Charlie is color 2, and so on). Since both (2) and (3) are partial answers to that question, if that question is under discussion in COLOR, the view will predict that relevant information is transmitted. In contrast, even if that question was under discussion in HEIGHT, (4) doesn't entail a

partial answer to it, so the view would predict that no relevant information is transmitted.

Present implementations of the question-based view lack the resources to explain why a question like `what maximally specific color is Charlie?` would be under discussion in a case like `COLOR`. The main way in which a question can be added to the stack of questions under discussion is through explicit interrogative utterances, and we have seen that, due to Question Variance, it would be unlikely for different people to agree on which question the speaker proposed adding to the stack through an interrogative utterance.

In fact, even if someone explicitly asked ‘what color is Charlie?’ it would be unlikely for two people to agree on what question she proposed adding to the common ground. Given (8), the alternatives that question makes salient are the propositions that result from substituting a property in the contextually determined domain for `F` in

(17) `F` is a color and Charlie is `F`.

As we saw above, there is a huge number of ways of dividing the color space into different regions. Since each of those ways of dividing the color space corresponds to a different set of properties that could be the set of properties in the contextually determined domain, it would be extremely unlikely for two people to agree on which of those sets of properties is the set of properties in the contextually determined domain. And if Anna and John disagree about what properties are in the contextually determined domain, they will disagree about what alternatives an utterance of ‘what color is Charlie?’ makes salient. Consequently, they will disagree about what question that utterance expresses.

Defenders of the present view may claim that, perhaps “by default”, language users will take the color properties in the domain to be maximally specific—e.g. that they take one such property to be that of being this or that maximally specific shade of red, this or that maximally specific shade of green, and so on. If that is the case, ordinary language users

would be in a position to agree about the question expressed by ‘what color is Charlie?’. Yet I see no reason to believe that ordinary language users make that kind of assumption.^{20, 21} That said, I believe appealing to an object’s maximally specific properties along certain dimensions can take us a long way towards distinguishing cases like COLOR from cases like HEIGHT. I will come back to this idea in section 2.5, where I will offer a notion of relevance that comes apart from partial answerhood.

To summarize, in this section I discussed the question-based picture of communication. According to that view, the information a speaker transmits to her audience is relevant to the conversation just in case it answers a question under discussion. I have argued that, given the truth of Question-variance, the participants in a conversation will rarely agree as to what question to answer. Absent such a question, the view predicts that relevant information is rarely exchanged in the context of ordinary conversations. The next three sections develop an alternative account of the transmission of relevant information, starting with an alternative picture of communication.

2.4 The “split records” picture of communication

In the last two sections I argued against the standard and the question-based pictures of communication. The arguments show that neither picture explains what it is for a speaker

²⁰ Perhaps another way of adding a question to the stack of questions under discussion is by accommodation. Defenders of the present view may claim that, even if two people rarely agree on what question is expressed by ‘what color is Charlie?’, in cases like COLOR in which no question is explicitly asked, they may accommodate so as to both assume that *What is Charlie’s maximally specific color?*. Defenders of this strategy may appeal to the use of focal alternatives discussed a few paragraphs above. However, as we saw there, appeal to focal alternatives won’t take us very far. If an utterance’s focal alternatives depend on a contextually determined domain, there is little chance that different people will accommodate by adding the same question to the stack of questions under discussion, since there is little chance that they will agree on what the contextually relevant domain is.

²¹ Thanks to Cian Dorr for pressing me on these issues.

to transmit relevant information to her audience through an assertoric utterance. However, this doesn't mean we should discard every element of those pictures. At some level of abstraction, it seems right to think that the participants in a conversation keep a record of what is being presupposed, as well as of what each participant in the conversation has said or asked. Because of Variance and Question-variance, it is very unlikely that any two participants' records of a given conversation are exactly the same, but the idea that each participant keeps her own personal record of the conversation is something we can build upon.

What I want to call a conversational record is not very different from what Stalnaker or Roberts call the common ground, or from what Lewis (1979) calls a conversational score. It is a body of information (i.e. a set of propositions) somebody presupposes for the purposes of the conversation, together with a stack of questions she intends to be answered over the course of the conversation, and whatever other items the participants in a conversation usually keep track of (e.g. discourse referents). The only difference between the common ground and the personal conversational records belonging to the various participants in a conversation is that the latter need not be “common” or “shared”: though two participants in a conversation may happen to have the same conversational record (i.e. to presuppose the same information and to intend that the same questions be answered), their conversational records may in principle be very different. My claim is that the participants in typical conversations each have their own conversational record, and that those records need not be shared or common.

The claim that the participants in a conversation need not share a conversational record raises various questions. For example, in Stalnaker's original picture, to make an assertion is to make a proposal to add certain information to the common ground. But if there is no common ground, one may wonder what kind of proposals (if any) we make by mak-

ing assertions. Or take interrogative utterances. According to the question-based picture, by making an interrogative utterance the speaker proposes that all the participants in the conversation intend to answer a certain question (i.e. she proposes to add that question to the stack of questions under discussion). But if there is no shared stack of questions under discussion, it may not be clear what the speaker proposes in making an interrogative utterance.

Here is my proposal. We can think of assertoric and interrogative utterances as proposals to add information or questions to the audience’s record of the conversation. In the case of literal assertoric utterances, the speaker will propose to add the proposition she takes to be her utterance’s truth-conditional content to the audience’s conversational record, but she may also propose to add other propositions in addition (e.g. implicatures).²² Upon hearing the speaker’s utterance, the audience will form a belief as to what proposition(s) the speaker proposed the audience to add to her conversational record. If the audience accepts what she takes to be the speaker’s proposal, she will add the proposition(s) she takes the speaker to have proposed to add to her conversational record. In turn, if the speaker believes that the audience accepts the proposal she takes the speaker to have made, the speaker will add the proposition(s) she takes herself to have proposed to add to the audience’s conversational record to her own conversational record.²³ The same goes, *mutatis mutandis*, for interrogative utterances. Call this the “split-records” picture of communication.²⁴

²² Here and throughout the next two sections I will assume that language users have determinate beliefs about the truth-conditional content of utterances made in the conversations they take part in. This assumption is in conflict with Uncertainty, but later (section 2.7.2) I will modify the present view so as to accommodate it.

²³ This description of a conversation takes for granted that speakers are not aware of the truth of Variance. Speakers aware of the truth of Variance may have somewhat different attitudes. For instance, rather than proposing to add a conversation to the conversational record, a speaker A will make utterance U expecting that her audience B will add the proposition she takes to be U’s truth-conditional content to B’s conversational record. If A believes that B has satisfied this expectation, A will add the proposition she believes to be U’s truth-conditional content to her own conversational record.

²⁴ Perhaps this is close to what Stalnaker has in mind when he talks about defective contexts in his (1978). According to Stalnaker, a context is defective if the participants in the conversation don’t make all the same

For example, consider again Anna and John’s conversation in COLOR. Through her utterance, Anna proposed to add (2) to John’s conversational record, but John took her to have proposed to add (3) instead. According to the present picture, if John accepts what he takes Anna’s proposal, he will add (3) to his conversational record, and if Anna believes that John has accepted her proposal, she will add (2) to her own conversational record. The same goes, *mutatis mutandis*, for Anna and Carla’s conversation in HEIGHT.

Because of Variance and Question-variance, speaker and audience may not agree about what proposal the speaker in fact made or on what conversational record accurately tracks the conversation. But, as COLOR and HEIGHT illustrate, not all conversations in which that happens are the same. Some of those conversations (e.g. Anna and John’s conversation in COLOR) will allow for the transmission of relevant information, and other’s won’t (e.g. Anna and Carla’s conversation in HEIGHT). By itself, the split-records picture does not capture this difference, but we can capture it as follows. A speaker transmits the relevant information P to her audience through an assertoric utterance just in case: (a) through her utterance, the speaker proposes to add to her audience’s record a proposition that entails P ; (b) P is not entailed by information already in the speaker’s conversational record or in the audience’s conversational record (i.e., it is informative by both speaker’s and audience’s lights); (c) the proposition the audience takes the speaker to propose adding entails P ; (d) the audience accepts what she takes to be the speaker’s proposal; and (e) P is about presuppositions. He adds:

A context is CLOSE ENOUGH to being nondefective if the divergences do not affect the issues that actually arise in the course of the conversation. Suppose for example that you know that Jones won the election, believe mistakenly that I know it as well, and are prepared to take the truth of this proposition for granted if the occasion should arise, say by using it as a suppressed premise in an argument, or by using the description *the man who won the election* to refer to Jones. [...] Since I do not know that Jones won the election, I do NOT presuppose it, and so the context is defective. But the defect may be harmless. (p. 85)

As far as I know, there is no developed account of this suggestion. Nor has Stalnaker’s account been developed so as to account for communication in defective contexts. Proponents of the common ground picture may think of the split-records picture as such a development.

a subject matter under discussion in the conversation. This last requirement is the part of the definition that is meant to capture the difference between relevant and irrelevant information, and which in turn will allow us to distinguish between cases like COLOR and cases like HEIGHT.

In a way, the differences between this definition of information transmission and the question-based theorist’s definition may seem superficial. In principle, the question-based theorist who adopted the split-records picture could accept all of (a)–(e) as long as subject matters are taken to be questions and aboutness is understood in terms of partial answerhood. In that case, requirement (e) will be that the transmitted information *P* be a partial answer to a question under discussion. Thus, the crux of the issue is not so much whether to accept (a)–(e), but what subject matters are, what it is for a proposition to be about a subject matter, and what it takes for a given subject matter to be under discussion in a given conversation. In the next section I will develop a conception of subject matter that allows us to answer these questions, but before proceeding it will be useful to see what such a conception is meant to accomplish.

Above all, and in combination with (a)–(e), an adequate conception of subject matters should put us in a position to distinguish between cases like COLOR and cases like HEIGHT. Furthermore, we should be able to accomplish that task while avoiding the problems that arise from thinking of subject matters as questions. We have already seen the first of those problems: if Question-variance is true and subject matters are questions, then it is difficult to see how the participants in a conversation could ever agree about what substantive questions to discuss. To improve upon the question-based picture, an adequate conception of subject matter should be such that we can explain how a subject matter could be under discussion in a given conversation despite the truth of Variance and Question-variance.

The second problem with the question-based conception of subject matters, which I

haven't yet discussed, is this. If the subject matter of whether Charlie is one of colors 3–6 is the partition that divides logical space into the set of worlds in which Charlie is one of colors 3–6 and the set of worlds in which he isn't, the proposition that Charlie is one of colors 2–5 turns out to be irrelevant to whether Charlie is one of colors 3–6. Yet there is a clear sense in which the proposition that Charlie is one of colors 2–5 is relevant to whether Charlie is one of colors 3–6: if Charlie is one of colors 2–5, that excludes some possibilities that affect whether Charlie is one of colors 3–6, which takes us some ways towards determining whether Charlie is one of colors 3–6. In contrast, the proposition that Charlie is at least half a meter tall is entirely irrelevant to whether Charlie is one of colors 3–6: for all the truth of the proposition that Charlie is at least half a meter tall requires, Charlie could be any color whatsoever. In this sense, the fact that Charlie is half a meter tall excludes no possibilities that would contribute to determining whether Charlie is one of colors 3–6.

This fact—that Charlie's being one of colors 2–5 is relevant to whether Charlie is one of colors 3–6—suggests we would be better off distinguishing between questions and subject matters. Question-based theorists would claim that one subject matter is whether Charlie is one of colors 2–5, another is whether Charlie is one of colors 3–6, and yet another Charlie's color. But I'd rather distinguish questions and subject matters: even if some questions may correspond to subject matters in some more or less natural way, not all of them do, nor does every subject matter correspond to a question. From this perspective, the proposition that Charlie is one of colors 2–5 is relevant to whether Charlie is one of colors 3–6 because the subject matter the latter is exactly about (Charlie's color) encompasses and, in this particular case, is identical to the subject matter the former is exactly about (i.e. Charlie's color). In contrast, the proposition that Charlie is at least half a meter tall is irrelevant to the question of whether Charlie is one of colors 3–6 because the former's exact subject matter

(Charlie's color) is orthogonal to (and so, not encompassed by) the latter's exact subject matter (Charlie's height).

It is one thing to distinguish between questions and subject matters, and another to offer an account of subject matter that adequately captures this distinction. The next section offers such an account of subject matters using truthmaker semantics. In combination with the split-records picture I introduced in this section, that notion will put us in a position to distinguish cases like COLOR from cases like HEIGHT while avoiding the problems arising from identifying subject matters with questions.

2.5 Subject matters in truthmaker semantics

Consider again COLOR and HEIGHT. In the introduction I claimed that in COLOR Anna transmitted information about Charlie's color to John through her utterance of (1); namely, that Charlie is not color 1 or color 7, that he is not red, and so on. I would like to say that this is so because both the proposition that Anna took herself to have expressed and the proposition John took her to express entail that Charlie is not color 1 or color 7, that he is not red, and so on. I also claimed that in HEIGHT Anna transmitted no information about Charlie's color to Carla through her utterance of (1). I would like to say that this is so because the proposition Carla took Anna to express (i.e. that Charlie is at least half a meter tall) entails no information about Charlie's color. What I will do now is offer a conception of subject matters that makes sense of this way of characterizing the difference between the two examples. Later, in section 6, I will explain how this conception of subject matters combines with the definition of information transmission I proposed in section 4.

Let's start with a bit of history. The treatment of questions as partitions of logical space was popularized by Groenendijk and Stokhof in the 80s. Around the same time,

Lewis (1988) proposed thinking about subject matters as sets of possibilities: the set of all possible ways for the world to be insofar as that subject matter is concerned. Lewis treated possibilities as sets of possible worlds, which led him to think of subject matters as partitions of logical space: the partition whose cells contain exactly the worlds that are equivalent with respect to the subject matter in question. For example, consider the subject matter of Charlie's color. According to Lewis, this subject matter should correspond to the partition of logical space whose cells contain exactly the worlds that are identical with respect to Charlie's color. If one way for the world to be with respect to Charlie's color is for Charlie to be color 1, and another for him to be color 2, then the subject matter of Charlie's color will contain both a possibility in which Charlie is color 1 and one in which he is color 2. Given this way of thinking about subject matters, it is natural to think of relevance in more or less the way question-based theorists think about it. That is, a proposition is about a given subject matter just in case it is equivalent to the union of a set of cells in that subject matter.

I am in principle sympathetic to Lewis' idea that a subject matter should be thought of as a set of ways for things to be with respect to that subject matter. But, as I said above, in a possible-worlds framework this comes down to thinking of subject matters as sets of sets of possible worlds. And if we think of propositions as sets of possible worlds, and of questions as sets of propositions, we are once more identifying subject matters with questions—and we have already seen the ways in which this is problematic. In order to keep the difference between questions and subject matters clear, I'd like to flesh out Lewis' idea in a truthmaker framework instead, and propose thinking of subject matters as sets of *possible states of affairs*: the sets of states of affairs that determine exactly how things are with respect to the subject matter in question.

Here is the plan: I will start by presenting some background on states of affairs; then

I will define the notion of a subject matter using that background, as well as some useful relations between subject matters; then I will define what it is for a proposition to be about a subject matter. In the next section I will show how the resulting notions of subject matter and relevance combine with the split-records picture to distinguish cases like COLOR from cases like HEIGHT while avoiding the problems of the question-based conception of subject matters.

2.5.1 Background: states of affairs

The main difference between possible worlds and possible states of affairs is that possible worlds settle the truth-value of absolutely all propositions, whereas most possible states settle the truth-value of only *some* propositions. For instance, a possible state of affairs that settles the truth-value of *it rains in Nepal* (e.g. a state of affairs in which that proposition is true, or a state of affairs in which that proposition is false) need not settle the truth-value of *Anna is 1.7m tall*. In fact, that state's obtaining may be completely irrelevant to the truth of the latter.²⁵

Possible states, unlike possible worlds, can stand in parthood relations—where parthood is understood as a reflexive, transitive, and antisymmetric relation. For example, a possible state in which it rains lightly in Nepal is part of a state in which it rains lightly in Nepal *and* Anna is 1.7m tall, and the latter is part of a state in which it rains lightly in Nepal, Anna is 1.7m tall, and John weighs 100kg. If two possible states are compatible

²⁵ The use of possible state of affairs or situations can be traced back at least to Barwise and Perry (1981, 1983), but the present implementation differs from theirs in significant ways. For example, Barwise and Perry require that states (what they call situations) have particular spatiotemporal locations, whereas, following Fine, I will impose no such constraint. See Fine (MSc,M,M) for recent work within truthmaker semantics, and Kratzer (2012, 1989, 2016) for recent work within a different situationist framework. Using Yablo (2014)'s version of truthmaker semantics, we could in principle think of possible states as sets of possible worlds. The reason I don't adopt that framework here is just that I'd like to keep the difference between subject matters and questions as clear as possible.


with each other—i.e. if they can jointly obtain—their *fusion* or mereological sum, defined below, is also a possible state of affairs. Taking the notion of a possible state as basic, we can think of possible worlds as maximal states: states of affairs that contain (i.e. have as parts) every state they are compatible with.

In a possible-worlds framework, it is customary to think of logical space as the set of all possible worlds. In truthmaker semantics, we can think of logical space as what Fine calls a *state space* (cf. Fine MSc,M,M, 2016). A state space is an ordered triple $\langle S, \leq, C \rangle$, where S is a non-empty set of possible states, \leq is a partial order on S (a *parthood* relation on S), C is a set of non-empty subsets of S (the set of exactly the subsets of S all of whose members are compatible with one another) such that:

- (a) For every $c \in C$, there is a state $s \in S$ such that every member of c is part of s , and every part of s has a part in common with some member of c . s will be the *fusion* of c , symbolised $c_1 + c_2 + \dots$ (where c_1, c_2, \dots are the members of c).^{26, 27}
- (b) For all states $s, s' \in S$, if there is no $c \in C$ such that $s, s' \in c$, then there is no $s'' \in S$ such that $s \leq s''$ and $s' \leq s''$.²⁸

²⁶ In other words, $\forall x(x \in c \rightarrow x \leq s) \ \& \ \forall y(y \leq s \rightarrow \exists z(z \in c \ \& \ \exists x'(x' \leq y \ \& \ x' \leq z))$. Hovda (2009) attributes this definition of fusion to Tarski (1956) and Lewis (1991).

²⁷ Fine takes the fusion of a set of states to be the set's least upper bound—the state containing every member of the set and which is included in any state which also contains every member of the set. Without further constraints on the space, Fine's definition allows for state spaces like this one, which the definition of fusion given by (a) excludes:



In the figure, the lines represent holdings of a parthood relation from a lower point to a higher point; \leq is the transitive and reflexive closure of that parthood relation. According to Fine's definition of fusion, w is the fusion of x and z , but according to my preferred definition, it isn't, since w has a part which doesn't overlap at all with x or z , namely, y . Thanks to Kit Fine for discussion.

²⁸ Fine imposes no such constraint, since he is happy to admit the existence of impossible states. In prin-

(c) For all $c \in C$, every subset of c is also a member of C .

(a) defines the fusion of a set of compatible states, and requires that the fusion of every set of compatible states in the state space be itself a member of the state space.

As I said above, by a set of compatible states I mean a set of states that can jointly obtain. For example, a state in which it rains in Nepal is compatible with a state in which Anna is 1.7m tall, but the latter is not compatible with a state in which Anna is 1.9m tall. The second requirement, (b), states that there can't be impossible states (i.e. states that have incompatible states as parts). As I also mentioned, we can think of possible worlds as maximal states of affairs: states that have every state they are compatible with as parts.

In order to get a better grasp of the notion of a state of affairs, it may be useful to think of them as *partial* functions from propositions to truth-values. For example, if P and Q are different propositions, one possible state of affairs will be the function that assigns truth to P, another the function that assigns falsehood to Q, and yet another the function that assigns truth to P and falsehood to Q.²⁹ If we think of states of affairs in this way, when I say that s is a state in which P, what I mean is just that $s(P) = T$.

If we think of states of affairs as partial functions from propositions to truth-values, we can define the parthood and compatibility relations as follows. First, if s and s' are possible states of affairs, then $s \leq s'$ just in case, for all propositions P, if $s(P)$ is defined, then $s(P) = s'(P)$. Second, two possible states of affairs s and s' are compatible just in case it's possible that all the propositions that any of them map to T are true and all the propositions that any of them map to F are false.

Given our characterization of possible states, there is a sense in which most possible

principle, I don't oppose to the use of impossible states, but they have no role to play in the present discussion. Another important difference is that, while Fine assumes the existence of a null state—i.e. a state which is part of every possible state—I make no such assumption.

²⁹ Here and elsewhere, I will assume that propositions are sufficiently finely individuated so that, if P and Q are logically independent propositions, then $P \vee \neg P$ is different from $Q \vee \neg Q$.

states are incomplete: there are propositions to which they do not assign truth-values. On the other hand, a crucial assumption in the context of the present discussion is that possible states are fully determinate in the following sense:

- (i) If a state of affairs determines that an object x has property F , then it also determines all of x 's maximally determinate properties along whatever dimensions—e.g. size, weight, color, degree of similarity to paradigmatic members of a certain class, etc.—on which the instantiation of F depends;
- (ii) If a state of affairs assigns truth to a disjunctive proposition $P \vee Q$, then it must also assign truth to at least one of P and Q .

Call the conjunction of (i) and (ii) the *determinacy assumption*.

For example, according to the determinacy assumption, any state of affairs in which Anna is between 1.5 and 1.6m tall is either a state of affairs in which Anna is 1.5m tall, a state of affairs in which Anna is 1.51m tall, or so on for some height between 1.5 and 1.6m. In other words, in order for a state of affairs to be one in which Anna is between 1.5 and 1.6m tall, that state of affairs must be one in which Anna is some specific height between 1.5 and 1.6m. Or suppose we want to talk about color. According to the determinacy assumption, there is no state of affairs in which Charlie is green simpliciter, or blue simpliciter. Instead, there are states of affairs in which Charlie is this or that maximally specific shade of green, this or that maximally specific shade of blue, and so on. Finally, according to the determinacy assumption, every state of affairs in which it is true that either Charlie is green or at least 1.5m tall must be a state in which Charlie is some maximally specific shade of green, or some maximally specific height above 1.5m tall.

Here is how to understand the determinacy assumption if we think of possible states of affairs as partial functions from propositions to truth-values. If, for example, in order

for a partial function that assigns truth to the proposition *Anna is between 1.5m and 1.6m tall* to count as a state of affairs, it must also assign truth to some proposition that states a maximally determinate height for Anna; that is, it must assign truth to one of the propositions *Anna is 1.51m tall*, *Anna is 1.57m tall*, or so on for some height between 1.5m and 1.6m. Or, to use another example, if a partial function assigns truth to the proposition *Charlie is green*, it must also assign truth to the proposition that Charlie is this or that maximally specific shade of green in order to count as a state of affairs.

The determinacy assumption will play a crucial role in what follows. As we will see, it is thanks to the determinacy assumption that, in conjunction with the account of information transmission from p. 74, my preferred account of subject matters distinguishes between cases like *COLOR* and cases like *HEIGHT*. The importance of the determinacy assumption will become clearer in section 2.5.4, when I apply the present framework to our discussion. For the time being, it is worth noting that, because of the determinacy assumption, it would be a mistake to think that every possible state of affairs corresponds to a set of possible worlds. For a disjunction may be true at a set of worlds (i.e. true in all the members of that set) without either disjunct being true at that set, yet, given the determinacy assumption, there is no possible state in which a disjunction is true but neither of its disjuncts is.³⁰

³⁰ Item (ii) of the determinacy assumption is very much in line with Fine and Yablo's respective versions of truthmaker semantics. For, according to Fine, the truthmakers for a disjunction are just the truthmakers for its disjuncts. According to Yablo, the kind of things that are suitable for playing the role of truthmakers are partial models for subsets of a language's atomic sentences. On the other hand, though neither Fine nor Yablo make explicit mention of item (i) of the determinacy assumption, they may find it more controversial. For example, Yablo (2014) insists that what counts as a possible truthmaker for a given sentence depends largely on context; as such, we may be willing to think that more or less determinate states of affairs can play the role of truthmakers depending on the context. I hope to explore the idea that the notion of truthmaking may itself be context-sensitive in future work, and to what extent this idea is compatible with the treatment I develop here.

2.5.2 Subject matters

As I said above, a subject matter is a set of states of affairs: the set of states of affairs that determine a way for things to be exclusively insofar as that subject matter is concerned.³¹ For example, the subject matter of Anna's height corresponds to the set of states of affairs that determine Anna's height and nothing else (other than whatever needs to be determined in order to determine Anna's height). That set will contain a possible state in which Anna is 1m tall and nothing else happens, a possible state in which she is 2m tall and nothing else happens, etc. It will not contain possible states of affairs in which it rains lightly in Nepal, or in which Anna is 1m tall *and* it rains lightly in Nepal. Or take the subject matter of Charlie's color. That subject matter corresponds to the set of states of affairs that determine Charlie's color and nothing else (other than whatever must be determined in order to determine Charlie's color). That set will contain, for example, a state in which Charlie is this or that shade of green, a state in which Charlie is this or that shade of red, a state of affairs in which Charlie is colorless, etc. It will not contain a state in which Charlie is this or that shade of green *and* it rains lightly in Nepal, or a state in which Charlie is this or that shade of green *and* Anna is 2m tall.

For the sake of simplicity, here I will suppose that every set of states of affairs is a subject matter, but there is a substantive question as to whether every set of states of affairs really does correspond to a subject matter. For instance, if every set of possible states of affairs is a subject matter, then we will count the singleton set of a state in which Charlie is 1.6m tall (say) as a subject matter, but this may not correspond to any subject matter

³¹ This is not the only available possibility. Fine treats subject matters as (often impossible) states of affairs, and Yablo treats subject matters as *divisions* of logical space. One reason not to adopt Fine's treatment is its commitment to the existence of impossible states of affairs. One reason not to adopt Yablo's treatment is the close relation between Yablian subject matters and ordinary language questions. As I have argued in the previous section, given the truth of Variance, we are better off not drawing such close connections between subject matters and questions. Another reason to distinguish subject matters from divisions of logical space has to do with the determinacy assumption, discussed in the previous subsection.

that we have a name for. As of this writing, my view is that such sets should be counted as subject matters, though they may not be especially natural. A more natural (though not entirely natural) class of subject matters is the class of subject matters such that, for any state s in that subject matter, that subject matter contains a state s' different from s and all of whose parts are incompatible with s . An even more natural class of subject matters is the class of subject matters M such that: (i) every part of every state $s \in M$ is incompatible with some other state in M , and (ii) for any state $s \in M$, every state s' that s is incompatible with and which satisfies condition (i) is also in M . Though all the subject matters I will discuss below are natural by at least one of these characterizations, the notion of naturalness itself plays no role in distinguishing cases like COLOR from cases like HEIGHT.³²

For our purposes, it will be useful to define a relation of *encompassment* between subject matters. When I say that a subject matter encompasses another, I mean that any way for things to be with respect to the latter goes some ways towards determining how things are with respect to the former. More precisely, where M and N are subject matters, M encompasses N just in case every state of affairs in N is part of a state of affairs in M —i.e., for all $s \in N$, there is some $s' \in M$ such that $s \leq s'$.

For example, take the subject matter of Anna's and John's respective heights, which corresponds to the set of states of affairs that determines how things are with respect to both Anna's and John's height. One of the states in this subject matter will be a state in which Anna and John are 1m tall, another will be one in which Anna is 1m tall and John is 1.4m tall, another will be one in which Anna is 1m tall but John is 1.5m tall, etc. This subject matter encompasses the subject matter of Anna's height, since every state that determines how things are with respect to Anna's height is part a of a state that determines how things are with respect to Anna's and John's heights. Accordingly, every way for

³² Thanks to Jim Pryor and Cian Dorr for discussion.

things to be with respect to Anna's height goes some ways towards determining how things are with respect to Anna *and* John's heights. In section 6 I will use the notion of subject matter encompassment to define what it is for a subject matter to be under discussion in a given conversation.

2.5.3 What is a proposition about?

Within the present framework, we can think of a proposition as an ordered pair of sets of states: the first member of the pair is the set of that proposition's possible exact truthmakers (possible truthmakers, for short), and the second the set of that proposition's possible exact falsitymakers (possible falsitymakers, for short). A proposition's *exact* subject matter is the set of all of that proposition's possible exact *truthmakers* and *falsitymakers*. Roughly speaking, a possible state is a possible exact truthmaker for a given proposition just in case it guarantees the proposition's truth if it obtains, and all of its parts play a role in guaranteeing the proposition's truth. A possible state is a possible exact falsitymaker for a proposition just in case it guarantees the proposition's falsity if it obtains, and all of its parts play a role in guaranteeing the proposition's falsity. For example, a state in which Anna is 1.7m tall is a truthmaker for both the proposition *Anna is 1.7m tall* and the proposition *Anna is at least 1.6m tall*, and a falsitymaker for both the proposition *Anna is 1.9m tall* and the proposition *Anna is at most 1.6m tall*, but a state in which Anna is 1.7m tall *and* it rains lightly in Nepal is neither a truthmaker nor a falsitymaker for any of those propositions. This is so because that state contains a part—the part in which it rains in Nepal—that plays no role in guaranteeing the truth of *Anna is 1.7m tall* or *Anna is at least 1.6m tall*, or in guaranteeing the falsity of *Anna is 1.9m tall* or *Anna is at most 1.6m tall*.

The question here is how to associate sentences with propositions. Here I will restrict

my attention to simple predications: i.e. sentences of the form *x is F*, where *x* is a denoting expression and *F* is a predicate and Boolean compounds thereof. Examples of simple predications are ‘Anna is tall’, ‘Charlie is green’, etc. A complete implementation of truthmaker semantics for natural language falls out of the scope of this dissertation.

To begin with, say that the truth of the proposition expressed by a simple predication *Fx* depends exactly on *x*’s properties along dimensions d_1, d_2, \dots, d_n —e.g. size, weight, degree of similarity to members of a certain class, etc.—just in case there is a set of points in the space generated by those dimensions such that (i) necessarily, *Fx* is true if and only if *x*’s properties fall in one of those points, and (ii) for every set *E* of dimensions of which the former condition is true, $\{d_1, d_2, \dots, d_n\} \subseteq E$. For example, the truth of ‘Anna is at most 1.6m tall’ Anna is at most 1.6m tall depends exactly on Anna’s height. This is so because there is a set of points along the dimension *Anna’s height*—namely, the open interval of Anna’s heights $[1.6, \infty)$ —such that (i) necessarily, ‘Anna is at most 1.6m tall’ is true if and only if Anna’s height falls in one of those points, and (ii) any set of dimensions *E* of which (i) holds—e.g. the set containing Anna’s height *and* weight, the set containing Anna’s height *and* shape, etc.—is a superset of the set containing only the dimension of Anna’s height.

My proposal is that if the truth of a simple predication depends exactly on an object *x*’s properties along dimensions d_1, d_2, \dots, d_n , then that sentence expresses a proposition whose possible truthmakers and falsitymakers are the minimal states of affairs that specify all of *x*’s maximally specific properties along those dimensions. In particular, that proposition’s possible truthmakers will be the minimal such states in which the proposition is true, and its possible falsitymakers will be the minimal such states in which the proposition is false. Here, *s* is a minimal state in which *P* is the case if and only if *P* is the case in *s* and, for no proper part of *s*, is *P* the case in it.

Keeping in mind that states of affairs should be maximally determinate, if we think of states of affairs as partial functions from propositions to truth-values, we can think of the possible truthmakers (falsitymakers) for the proposition expressed by a simple predication as *minimal* functions that assign truth (falsity) to that proposition. In other words, if P is the proposition expressed by a simple predication, a possible state s is a possible truthmaker for P just in case: $s(P) = T$, and for no s' such that $s' \leq s$, is it the case that $s'(P) = T$. Similarly, s is a possible falsitymaker for P just in case: $s(P) = F$, and for no s' such that $s' \leq s$, is it the case that $s'(P) = F$.³³

Given possible truthmakers and falsitymakers for the propositions expressed by atomic sentences, we can define possible truthmakers and falsitymakers for the propositions expressed by Boolean compounds as follows. Roughly following Fine (2012), where ‘ \Vdash ’ symbolizes the truthmaking relation between a state and the proposition a sentence expresses, and ‘ $\dashv\vdash$ ’ symbolizes the falsitymaking relation,³⁴

- $s \Vdash \neg\phi$ iff $s \dashv\vdash \phi$
- $s \dashv\vdash \neg\phi$ iff $s \Vdash \phi$
- $s \Vdash \phi \wedge \psi$ iff for some states v, v' , is it the case that $v \Vdash \phi$, $v' \Vdash \psi$, and $s = v + v'$

³³ This definition of truthmakers corresponds to Kratzer (2012)’s definition of what it is for a situation to exemplify a proposition. As Fine (MSa) and Yablo (forthcoming) note, the present definitions do not really capture the intuitive idea of a truthmaker or a falsitymaker. For example, take the proposition that there are infinitely many apples. Presumably, there is no minimal state in which this proposition is true, but that proposition may still have a truthmaker. Or take the proposition $P \vee (P \wedge Q)$; presumably, the fusion of a truthmaker for P and a truthmaker for Q is a truthmaker for this proposition, yet it is not a minimal state in which the proposition is true. Partly for these reasons, I only intend the claim about minimality to hold for the truthmakers and falsitymakers of propositions expressed by simple predications, though even in this case the present definition may not be fully satisfying. For philosophical elaboration on what truthmakers and falsitymakers are, see Yablo (forthcoming).

³⁴ The present clauses correspond to Fine’s notions of *exact* verification and falsification, which correspond to what I call truth and falsitymaking. The truthmaking clause for disjunction and falsitymaking clause for conjunction correspond to Fine’s “inclusive” version of verification and falsification, according to which a fusion of verifiers for a disjunction is itself a verifier for that disjunction.

- $s \dashv\vdash \phi \wedge \psi$ iff $s \dashv\vdash \phi$, $s \dashv\vdash \psi$, or for some states v, v' , is it the case that $v \dashv\vdash \phi$, $v' \dashv\vdash \psi$, and $s = v + v'$.
- $s \Vdash \phi \vee \psi$ iff $s \Vdash \phi$, $s \Vdash \psi$, or for some states v, v' is it the case that $v \Vdash \phi$, $v' \Vdash \psi$, and $s = v + v'$
- $s \dashv\vdash \phi \vee \psi$ iff for some states v, v' , is it the case that $v \dashv\vdash \phi$, $v' \dashv\vdash \psi$, and $s = v + v'$

I will elaborate on subject matters by means of an example below. Before that, it is worth introducing a new notion of entailment we will use later on, when I discuss the application of the present framework to the split-records picture. Within the present framework, we can define a new relation of *exact* entailment as follows: P exactly entails Q just in case every possible truthmaker for P has a possible truthmaker for Q as a part. So, for example, if p^+ is the only truthmaker for P, q^+ is the only truthmaker for Q, q^- is the only falsitymaker for Q, and neither q^+ nor q^- are parts of p^+ , P will not exactly entail $Q \vee \neg Q$. On the other hand, P will exactly entail $P \vee Q$, and $P \wedge Q$ will entail both P and Q.³⁵

Let's get back to subject matters. An example comparing the subject matters of two different propositions expressed by simple predications will make things clearer. Suppose that the only maximally determinate colors an object can be are 1–7, and that the only heights an object can be are 0.5 and 1m tall. Suppose also that the only possible states of affairs are $C_1, C_2, \dots, C_7, H_1$, and H_2 (plus their compatible fusions), where C_n is a state of affairs in which Charlie is color n and nothing else happens (other than whatever must happen for Charlie to be color n), H_1 is a state in which Charlie is 0.5m tall and nothing else happens, and H_2 is a state in which Charlie is 1m tall and nothing else happens. Finally,

³⁵ The nomenclature here may be confusing to readers familiar with Fine's work. On Fine's use of "exact entailment", P exactly entails Q just in case every truthmaker for P is itself a truthmaker for Q. This is not what I mean by "exact entailment". As far as I know, the notion of exact entailment I define here plays no role in Fine's framework.

suppose that, for no i, j , does C_i overlap with H_j , and that for every i, j , C_i is compatible with H_j .

In this case, we can represent the subject matter *Charlie's color* as the set whose members are all of C_1 – C_7 . Now, take for instance the sentence ‘Charlie is one of colors 2–5’, and say it expresses the proposition (2), `Charlie is one of colors 2–5`. Given our present suppositions, (2)’s possible truthmakers are the states C_2, C_3, C_4, C_5 , and its possible falsitymakers are C_1, C_6, C_7 ; thus, the set of all of (2)’s possible truthmakers and falsitymakers is the set containing exactly the states C_1 – C_7 . Accordingly, given our present definition, (2)’s exact subject matter is Charlie’s color. The proposition whose truthmakers are the states C_3, C_4, C_5, C_6 , and whose possible falsitymakers are the states C_1, C_2, C_7 (i.e. `Charlie is one of colors 3–6`) is also exactly about Charlie’s color, since the set of all its possible truthmakers and falsitymakers is the set containing exactly the states C_1 – C_7 .

Besides the subject matter a proposition is *exactly* about, a proposition may be about many other subject matters. In particular, a proposition is about every subject matter that encompasses its exact subject matter—i.e. a proposition is about every subject matter such that all of that proposition’s possible truthmakers and falsitymakers are part of a state in that subject matter. For example, consider the subject matter of Charlie’s color *and* height, which, given the present assumptions, we can represent as the set $\{s + s' \mid s \in \{C_1, \dots, C_7\}, s' \in \{H_1, H_2\}\}$. This subject matter contains a state in which Charlie is color 1 and 0.5m tall, a state in which Charlie is color 1 and 1m tall, etc. Since all of (2)’s truthmakers and falsitymakers (i.e. C_1 – C_7) are parts of states in the set $\{s + s' \mid s \in \{C_1, \dots, C_7\}, s' \in \{H_1, H_2\}\}$, (2) is also about Charlie’s color *and* height.

This is not to say that (2) is about Charlie’s height. Given the present assumptions, we can represent the subject matter of Charlie’s height as the set containing exactly H_1 and

H_2 . And, provided that none of the states C_1 – C_7 is part of H_1 or H_2 , no state in (2)'s exact subject matter (i.e. the set containing exactly C_1 – C_7) is part of a state in the subject matter of Charlie's height. The reason (2) tells us something about Charlie's color *and* height is just that it tells us something about Charlie's color, not that it tells us something about Charlie's height.

We can define the subject matters of questions in a similar way. Following the literature on the semantics of interrogatives, I will take questions to be sets of propositions (i.e. the set of the question's possible complete answers). The only difference is that current literature identifies a question's complete possible answers with cells in a partition of logical space, whereas I will do no such thing. On my preferred approach, where $ALT-\alpha$ is a set of alternatives determined by a given interrogative, the semantic value of the interrogative *wh*- α ? is the set of all the maximally consistent conjunctions of propositions in $ALT-\alpha$ or their negations. So, for example, if $ALT-\alpha$ is the set containing P and Q , the semantic content of *wh* α ? is the set containing the propositions: $P \wedge Q$, $\neg P \wedge Q$, $P \wedge \neg Q$, $\neg P \wedge \neg Q$.

I want to propose that a question's exact subject matter is the union of the exact subject matters of its complete answers. A question is also about every subject matter that encompasses its exact subject matter; i.e. it is about every subject matter that any of its complete answers are about. For example, the question `Is Charlie one of colors 3-6?` is both about *Charlie's color* and about *Charlie's color and height*. This is so because that question's complete answers are the propositions `Charlie is one of colors 3-6` and `Charlie is not one of colors 3-6`, which are themselves about Charlie's color and about Charlie's color *and* height. On the other hand, the question `Is Charlie one of colors 3-6?` is not about Charlie's height alone, since none of its complete answers is about Charlie's height alone.

2.5.4 Putting subject matters to use

By the beginning of this chapter I said that one difference between COLOR and HEIGHT is the following. In COLOR Anna transmitted information about Charlie's color to John through her utterance of (1): that Charlie is not color 1 or color 7, that he is not red, and so on. I said that this is so because both the proposition that Anna took herself to have expressed and the proposition John took her to express entail that Charlie is not color 1 or color 7, that he is not red, and so on. In HEIGHT Anna transmitted no information about Charlie's color to Carla through her utterance of (1). I said that this is so because the proposition Carla took Anna to express (i.e. that Charlie is at least half a meter tall) entails no information about Charlie's color. The present conception of subject matter allows us to make sense of these claims.

Take first the claim that, in COLOR, Anna transmitted information to John about Charlie's color, such as the information that Charlie is not color 1 or color 7. Given the present assumptions, the possible truthmakers for the proposition that Charlie is not color 1 or color 7 are C_2 – C_6 , and its possible falsitymakers are C_1 and C_7 . Thus, the proposition that Charlie is not color 1 or color 7 is exactly about Charlie's color (i.e. the set containing exactly C_1 – C_7); furthermore, since its truth excludes at least some ways for the world to be with respect to that subject matter (i.e. not just any state in the subject matter *Charlie's color* is a possible truthmaker for it), that proposition is informative about Charlie's color.

Now take the claim that, in a case like HEIGHT, Anna did not transmit any information to Carla about Charlie's color. Given the present framework, this is so because (4)—Charlie is at least half a meter tall—entails no informative proposition about Charlie's color.³⁶ This is so because every possible truthmaker for (4) is compatible with all of C_1 – C_7 , which is to say that (4)'s truth excludes no possibility pertaining

³⁶ I'm assuming that tautologies are not informative, since they exclude no possibilities whatsoever.

exclusively to Charlie's color. And since (4) entails no informative proposition about Charlie's color, there is no informative proposition about Charlie's color that both (4) and (2) entail.³⁷

Using the present framework, we can also distinguish between the case in which *John* prompts Anna's utterance of (1) by asking (5) ('Is Charlie green?') and the case in which *Carla* does it. By asking (5), John asked Anna whether Charlie is one of colors 3–6. Since that question and the information Anna transmitted to John through her utterance (e.g. that Charlie is neither color 1 nor color 7) are both exactly about Charlie's color, the present framework allows us to capture the idea that that information was relevant to John's question even though it does not constitute an answer to it. In contrast, by asking (5) *Carla* asked Anna whether Charlie is at least half a meter tall. That question's exact subject matter is Charlie's height, which, given present assumptions, corresponds to the set containing exactly H_1 and H_2 . Since (2) entails no information about that question's exact subject matter, Anna could not have transmitted to Carla any information relevant to her question.

Before concluding this section, I would like to call attention to an assumption I made in section 5.1. There, I said that possible states of affairs must be maximally determinate—e.g. that any possible state in which, say, Charlie is green, must be a state in which Charlie must be this or that maximally determinate shade of green. It is thanks to this assumption that we can say that various questions and propositions may have the same exact subject matter, which in turns explains why a proposition like (2) can be relevant to a question like (5) even though it doesn't answer it.

As the discussion illustrates, the present conception of subject matter allows us to dis-

³⁷ Note that, assuming that the truthmakers for a disjunction include the truthmakers for each of its disjuncts, the proposition $(2) \vee (4)$ is not about Charlie's color because one of its truthmakers—namely, the state in which Charlie is 1m tall—is not part of a state in that subject matter.

tinguish between information that is relevant and information that is irrelevant to a given subject matter. This conception of subject matter allows us to point to a difference between cases like COLOR and cases like HEIGHT, and to capture the idea that a proposition may be relevant to a given question even if the former is not a partial answer to the latter. However, the framework by itself remains silent on whether a piece of information is relevant in the context of a conversation. For it remains silent on what subject matter, if any, is at issue in the context of a conversation. The next section remedies this by embedding the present account of relevance into the split-records picture of communication.

2.6 Subject matters and the split-records picture

In section 4 I defined what it is for a speaker to transmit relevant information to her audience in the context of a conversation. I said that a speaker transmits the relevant information P to her audience through an assertoric utterance just in case: (a) through her utterance, the speaker proposes to add to her audience's record a proposition that entails P ; (b) P is not entailed by information already in the speaker's conversational record or in the audience's conversational record (i.e., it is informative by both speaker's and audience's lights); (c) the proposition the audience takes the speaker to propose adding entails P ; (d) the audience accepts what she takes to be the speaker's proposal; and (e) P is about a subject matter under discussion in the conversation. Here, the notion of entailment should be thought of as *exact* entailment (see above, p. 68).

With the conception of subject matter I developed in the previous section, we are now in a position to cash out (e). I want to propose that, by adding a proposition to the set of propositions in her conversational record, or a question to the stack of questions she intends to discuss, a participant in a conversation commits herself to adopting certain subject

matters as subject matters under discussion in the conversation. For example, if I take my interlocutor to propose adding P to my conversational record, I could not reasonably add P to my conversational record without taking myself to have added information about P's exact subject matter or without taking my interlocutor to have said something about it. In that way, by accepting to add P to my conversational record, I commit myself to adopting P's exact subject matter as a subject matter under discussion in the conversation, along with any subject matter P's exact subject matter encompasses.³⁸

More generally, if someone adds a proposition or a question to her conversational record as a result of hearing somebody's utterance, she thereby commits herself to adopting that proposition or question's *exact* subject matter as a subject matter under discussion in the conversation, along with whatever subject matters that subject matter encompasses. A subject matter is under discussion for the purposes of a conversation if all the participants in the conversation are committed to discussing subject matters that *encompass* it.

For example, if, after Anna's utterance of (1) ('Charlie is green') in COLOR John adds (3) (Charlie is one of colors 3-6) to his conversational record, and Anna adds (2) (Charlie is one of colors 2-5) to hers, they thereby commit themselves to adopting *Charlie's color* as a subject matter under discussion—i.e. John could not reasonably add (3) to his conversational record without taking himself to have added information about Charlie's color, and the same goes for Anna with respect to (2). And since Anna and John are so committed, and the subject matter of Charlie's color encompasses itself, that subject matter will be among the subject matters under discussion for the purposes of Anna and John's conversation. The relevant information Anna will have transmitted to John through her utterance of (1) is the information about Charlie's color entailed by both (2) and (3).

³⁸ This is not to say that I *must* accept any proposal regardless of its subject matter. For example, I may reject proposals on the basis that they don't address subject matters already under discussion. The point is just that, *if* I accept a proposal, then I commit myself to thinking that its exact subject matter is under discussion.

Compare this case with HEIGHT, in which Carla took Anna's utterance of (1) to have the truth-conditional content (4) (*Charlie is at least half a meter tall*). If Anna adds (2) to her conversational record, she will be committed to discussing Charlie's *color*; in contrast, by adding (4) to her conversational record, Carla is committed to discussing Charlie's *height*. Given our assumptions on these two subject matters from the previous section, there is no subject matter that those two subject matters both encompass. Accordingly, there is no subject matter under discussion in Anna and Carla's conversation, and whatever information Anna transmits to Carla will count as irrelevant to the conversation. The present view can thus distinguish between cases like COLOR and HEIGHT.³⁹

Here is another, more radical example. Suppose that Anna and John disagree about the range of application of 'green' to an even greater extent than we have assumed. Anna believes that 'green' applies to individuals whose color is within a certain range of hues (2–5) and saturation (say, 60-100%), but that lightness is irrelevant to whether something ought to be called 'green'. John believes that 'green' applies to individuals whose color is within a different range of hues (3–6) and a specific range of lightness (say, 60-100%), but he also believes that saturation is irrelevant to whether something ought to be called 'green'. Thus, when Anna utters 'Charlie is green', the truth-conditional contents Anna and John (respectively) associate with Anna's utterance have different exact subject matters. By Anna's lights, her utterance is exactly about Charlie's hue and saturation; in contrast, by John's lights, Anna's utterance is exactly about hue and lightness. The subject matter of

³⁹ Note that the determinacy assumption from section 63 plays a crucial role in delivering the result that there is no substantive subject matter under discussion in a case like HEIGHT. Without the determinacy assumption, there could be a state in which Charlie is at least half a meter tall or one of colors 2–5 true without being any determinate height or color. Similarly, there could be a state in which Charlie is less than half a meter tall or color 1, 6, or 7 without being any more determinate height or color. Call these states x and y , respectively. Given that a state s is part of s' just in case s' agrees with s on every proposition for which s is defined, x and y would both be part of truthmakers or falsitymakers for the proposition that Charlie is one of colors 2–5 and the proposition that Charlie is at least half a meter tall. In that case, the present view would predict that there is a substantive subject matter under discussion in a case like HEIGHT, namely, the subject matter containing x and y .

Charlie's hue is encompassed by both of the previous subject matters, so if Anna and John add the proposition they each take to be the truth-conditional content of Anna's utterance, the subject matter of Charlie's hue is under discussion in the context of their conversation.

To summarize, using the conception of subject matters developed in the previous section, we can flesh out what it is for a subject matter to be under discussion in the context of a conversation. Together with my proposal on what it is for someone to transmit relevant information to her audience through an assertoric utterance, that allows us to fully distinguish between cases like COLOR and cases like HEIGHT.

In section 4 I also said that an adequate conception of subject matter should (unlike the question-based conception) put us in a position to explain how to determine what subject matters are under discussion in a given conversation given the truth of Variance and Question-variance, and that it should explain why a proposition can be relevant to a given question even if it does not partially answer it. As I showed in this section, the truthmaker conception of subject matter satisfies these two requirements. As I explained in this section, the subject matters under discussion in a given conversation are determined by the subject matters the participants in the conversation are committed to adopting. In turn, the participants in a conversation are committed to adopting subject matters in virtue of the propositions and questions they accept to add to their respective conversational records. Given this account, and despite the truth of Variance and Question-variance, it is not a mystery what subject matters are under discussion in ordinary conversations. As I showed in the previous section, the present conception of subject matter also explains why a proposition can be relevant to a given question even if it does not partially answer it. In this way, the split-records picture in conjunction with my proposed characterization of what it is for a piece of information to be relevant in the context of a conversation and the truthmaker conception of subject matters is highly preferable to the standard common-ground

and question-based pictures. The next section discusses two further issues arising for the present view.

2.7 Further issues

2.7.1 Exploiting information in a conversational record

Above I proposed that a speaker transmits *relevant* information P to her audience through an assertoric utterance U just in case: (a) through her utterance, the speaker proposes to add to her audience's record a proposition that entails P; (b) P is not entailed by information already in the speaker's conversational record or in the audience's conversational record; (c) the proposition the audience takes the speaker to propose adding entails P; (d) the audience accepts what she takes to be the speaker's proposal; and (e) P is about a subject matter under discussion.

For simplicity, I have focused on cases in which the information a speaker proposes to add to other conversationalists' records is just the proposition she believes to be the truth-conditional content of her utterance. This may not always be so. Oftentimes, the speaker will attempt to exploit some of the information she presupposes for the purposes of the conversation. For example, suppose Anna assumes (and assumes that Anna assumes) for the purposes of the conversation that all bachelors are available for dating. Accordingly, Anna utters 'John is a bachelor', expecting Carla to infer that John is available for dating on the basis of that utterance plus information in Carla's conversational record.

The present picture can accommodate cases like this without further modification. For in making a single utterance, a speaker may make several different proposals to update her conversational partner's conversational records. For instance, in the case I just presented, Anna may take herself to propose adding what she takes to be her utterance's truth-

conditional content to Carla's conversational record, but, in addition, she may take herself to propose adding the proposition that John is available for dating. Anna will have transmitted the information that John is available for dating to Carla if: Carla doesn't already assume that John is available for dating, Carla takes Anna to propose adding a proposition to her conversational record that entails that John is available for dating, and Anna accepts that proposal. The information that John is available for dating will be relevant in the context of Anna and Carla's conversation just in case it is about a subject matter under discussion in the context of that conversation.

2.7.2 Troubles with *Uncertainty*

In chapter 1 I said that the kind of considerations that support Variance also support *Uncertainty*—the thesis that, for almost every assertoric utterance and every language user, there is no proposition which that language user believes to be the utterance's truth-conditional content. Uncertainty is problematic for the split-records picture of communication I presented above. As I presented it, that picture presupposes that the participants in a conversation have determinate beliefs as to what proposition(s) a speaker proposes to add to a conversational record through an assertoric utterance. Throughout the discussion, I assumed that the participants in a conversation formed such beliefs partly on the basis of their beliefs about an utterance's truth-conditional content. However, if *Uncertainty* is true, ordinary language users have no such beliefs, which makes it difficult to see how they could form determinate beliefs about what proposition(s) a speaker proposes to add to a conversational record through an assertoric utterance. Similar problems arise with respect to questions and interrogative utterances.

The current picture can be modified so as to accommodate *Uncertainty*. Suppose for example that Anna makes an utterance and is undecided as to which of propositions P_1

to P_5 are that utterance's truth-conditional content: she believes that her utterance's truth-conditional content is one of those propositions, but she is just as certain that it is P_1 as she is that it is P_2 , that it is P_3 , and so on. Carla hears Anna's utterance and is undecided as to which of propositions P_3 to P_8 is that utterance's truth-conditional content. In such a case, though neither Anna nor Carla would have definite beliefs about the utterance's truth-conditional content, Anna could still have transmitted to Carla the information that all of P_1 – P_8 entail and is in neither Anna's nor Carla's conversational records. So, at a minimum, we could say that the information Anna transmits to Carla through her utterance U is the information which is (a) entailed by all the propositions Anna takes to be live candidates for U 's truth-conditional content, and (b) entailed by all the propositions Carla takes to be live candidates for the same utterance, and (c) not entailed by information already in Anna's or Carla's conversational records.

Now, just as Anna and Carla may be undecided as to which of many different propositions is an utterance's truth-conditional content, they may also be undecided as to which conversational record accurately tracks the course of the conversation. For example, Carla may be undecided as to whether the information presupposed for the purposes of the conversation is the set of propositions Γ_1 or Γ_2 . Accordingly, if Carla accepts Anna's utterance, she will be undecided as to which of the sets $\Gamma_1 \cup \{P_3\}$, $\Gamma_2 \cup \{P_3\}$, $\Gamma_1 \cup \{P_4\}$, and so on, accurately tracks the conversation once Carla accepts Anna's utterance. More generally, suppose Carla is undecided as to which of the sets of propositions in a set A accurately tracks the information presupposed for the purposes of the conversation, and she is undecided as to which of the propositions in a set B her interlocutor is proposing to add to her conversational record. If Carla accepts her interlocutor's proposal, she will be undecided as to which of the conversational records in the set $\{a \cup \{b\} \mid a \in A, b \in B\}$ accurately tracks the conversation. As the conversation progresses, Carla will update the set of records she

is undecided between in the same way.⁴⁰

Here is my proposal. What it is for a speaker to transmit relevant information to her audience through an asertoric utterance is for there to be a proposition P such that: (a) P is entailed by every proposition the speaker considers to be a candidate for the information she proposes to add to her interlocutor's conversational record; (b) P is not entailed by information already in any of the speaker's or the audience's candidate conversational records; (c) P is entailed by every proposition the audience considers to be a candidate for the information the speaker proposes to add; (d) the audience updates the set of her candidate records in the way proposed above with a candidate proposal; and (e) P is about a subject matter under discussion. The subject matter under discussion in the context of the conversation is determined in a similar way: the subject matter someone is committed to by accepting an utterance is the *maximal* subject matter encompassed by the exact subject matters of all the candidate contents of that utterance—i.e. the subject matter M encompassed by the exact subject matters of all the candidate contents of that utterance, and such that any subject matter N encompassed by the exact subject matters of all the candidate contents of that utterance are encompassed by M . In turn, a subject matter is under discussion in the context of a conversation just in case every participant in the conversation is

⁴⁰ I am assuming for the sake of simplicity that Carla assigns equal weight to each conversational record she takes to be a live possibility and, similarly, that she assigns equal weight to each proposition she takes to be a live candidate for being the truth-conditional content of Anna's utterance. However, we can accommodate the idea that Carla may assign different weights to different candidates as follows. If Anna has credence c that conversational record a accurately tracks the conversation, and credence c' that proposition P is the truth-conditional content of Anna's utterance, her credence that $a \cup \{P\}$ accurately tracks the conversation should be $c \cdot c'$. Thanks to Jim Pryor for discussion. See footnote 42, below, for a brief discussion of how this affects the present view of information transmission.

committed to discussing a subject matter that encompasses it.^{41, 42}

2.8 Conclusion

This section offered an account of the transmission of relevant information, given the truth of Variance. As I argued, given Variance and the related thesis Question-variance, the standard and question-based pictures of communication cannot explain that phenomenon. Before moving on to the next chapter, I would like to discuss two interesting consequences of the present approach.

The first consequence concerns the debate between contextualists and invariantists in philosophy of language. As I explained in the previous chapter, invariantists have claimed that contextualists are not in a position to explain cross-contextual communication. According to contextualists, speakers who find themselves in different contexts are prone to have different beliefs about the truth-conditions of the same utterances, which, under the standard picture of communication, makes it difficult to see how there could be cross-contextual communication. As I said in the previous chapter, given the truth of Variance, there is no special problem of cross-contextual communication, and so, no special objection

⁴¹ If it turns out that language users are usually uncertain as to which truth-conditional content candidates they are uncertain between, we can iterate the present strategy as many times as necessary. For example, suppose that S is uncertain as to which of sets $\Gamma_1, \Gamma_2, \dots, \Gamma_n$ is such that S is undecided as to which of its members is U's truth-conditional content. Then the set of propositions that should play a role in determining what subject matters are under discussions in a conversation involving U should be the union of $\Gamma_1, \Gamma_2, \dots, \Gamma_n$. Alternatively, we could just take the set Γ such that S is uncertain as to which of the propositions in Γ is U's truth-conditional content, and it is definitely the case that S is uncertain as to which of the propositions in Γ is U's truth-conditional content, and it is definitely the case that . . . , and so on, ad infinitum. I hope to explore these options in more detail in future work, but an initial worry with the first of these options is that it may lead to triviality if every truth-conditional content whatsoever can be reached by a long enough chain of "uncertain whether"s. Thanks to Cian Dorr and Stephen Schiffer for pressing me on these issues.

⁴² If a language user assigns different weights to different potential truth-conditional content candidates for a given utterance, the live candidates will be those for which the speaker has a sufficiently high credence that they are the actual truth-conditional content of the utterance in question. Similarly, the live conversational record candidates will be those for which the speaker has a sufficiently high credence that they accurately track the conversation. Thanks to Jim Pryor for discussion.

to contextualism. But, more importantly, the account of relevance and the general picture of communication I proposed in this chapter can be used to explain the transmission of relevant information across contexts and in the same context alike. As we saw in chapter 1, the transmission of relevant information is only one of the features of typical cases of successful communication; nevertheless, this is progress.

The second consequence concerns the role of the common ground and the notion of a question under discussion in explanations of typical pragmatic phenomena. Philosophers like Stalnaker (1978) and linguists like Roberts (2012) have made heavy use of those notions to explain the felicity of certain utterances: informative identity sentences in the case of Stalnaker, and felicitous discourse structure in the case of Roberts. Those explanations are, at bottom, explanations of certain ways of reasoning in the context of a conversation, and are entirely compatible with the split-records picture of communication. The availability of the split-records picture thus casts doubt on the theoretical usefulness of the notion of a common ground.

Chapter 3

Partly factual disputes

3.1 Introduction

Think of disputes as linguistic interactions (as opposed to, say, conflicts in people's attitudes).¹ If we think of disputes in this way, it is enough for two people to have a dispute over the truth of a given utterance that one of them makes that utterance and the other says 'that's not true', 'that's false', 'you're wrong', or something along those lines.

Beyond whatever satisfaction or prestige one may get from winning a dispute, it seems that part of what is at stake in typical disputes is knowledge of certain extralinguistic facts.² In particular, whenever two people have a dispute over the truth of an utterance U, at most one of them knows the facts they take the utterance under dispute to be about. For example,

¹ Cf. Cappelen and Hawthorne's (2009) distinction between disagreement as an activity and disagreement as a state.

² Here and elsewhere in the chapter I call the facts the disputants take an utterance to be about *extralinguistic*, but in a way this is a misnomer. For the facts the participants in a dispute take the disputed utterance to be about may themselves be linguistic. For example, two people may have a dispute about the properties of a certain word or construction. Throughout this chapter, when I talk about extralinguistic facts, I am merely talking about whatever facts the participants in a dispute take an utterance to be about, even when, strictly speaking, those facts may be entirely linguistic. In this respect, my use of 'extralinguistic' is very idiosyncratic.

if Anna utters ‘grass is green’ and John says ‘no, it isn’t’, and they are both competent speakers of English, it seems at most one of Anna and John knows whether grass is green; in this way, Anna and John’s dispute exhibits an epistemic conflict.

I will state an improved definition later on, but for now let’s say that a dispute over the truth of an utterance exhibits an *epistemic conflict* just in case at most one of the disputants knows the truth-value of the proposition she believes to be the truth-conditional content of the utterance under dispute. Given certain assumptions, the fact that Anna and John’s imagined dispute over the truth of Anna’s utterance of ‘grass is green’ exhibits an epistemic conflict may not seem surprising. After all, one may think that if Anna and John are both competent speakers of English, they will both know that the truth-conditional content of Anna’s utterance is the proposition that grass is green. And if that is so, then if Anna knows that her utterance is true, John doesn’t know that grass is green; and if John knows that Anna’s utterance is false, then Anna doesn’t know that grass is not green. More generally, according to this kind of explanation, typical disputes exhibit epistemic conflicts because they are *fully factual*: in typical disputes over an utterance’s truth, there is a proposition all the parties of the dispute take to be that utterance’s truth-conditional content.

The truth of Variance raises a problem for this kind of seemingly straightforward explanation. If Variance is true then nearly every dispute over an utterance’s truth is verbal—i.e. there is no proposition that all the participants in the dispute take to be the truth-conditional content of the utterance the dispute is about³—and verbal disputes are not in general the

³ This kind of characterization of verbal disputes is very popular. See e.g. Hirsch (1993, p.181), who states a similar characterization in terms of believed equivalences between sentences rather than in terms of beliefs about an utterance’s truth-conditions. Sidelle (2007, p.89) takes a similar characterization to be the standard picture of verbal disputes. (Sider 2006, p.76) uses a similar definition, albeit in terms of what an utterance means for the different parties of a dispute. See Vermeulen (2018) for a similar characterization, albeit in terms of speaker meaning. Versions of this characterization of verbal disputes also play crucial roles in standard arguments from “lost disagreement”, according to which if two parties to a dispute believe compatible propositions, they do not (really) disagree (See e.g. MacFarlane 2007, 2014, Kolbel 2004, Egan 2010, Laserohn 2005, Stephenson 2007, Huvenes 2012). As Cian Dorr has pointed out in personal communication, the present characterization of verbal disputes only seems to apply to disputes over an utterance’s truth, and it is

kind of disputes that exhibit epistemic conflicts.

For example, suppose that Anna and John have a mutual friend, Carla, who was supposed to leave a kilo of cocaine in a certain locker and failed to do so: though she kept the cocaine with her at all times, she lost the key to the locker. When John asks Anna why Carla didn't leave the cocaine in the locker, Anna replies:

(1) Carla lost the key,

to which John replies

(2) No, she didn't! She had it last time I saw her!

As it happens, Anna and John have different beliefs about the meaning of the word 'key' as it was used in Anna's utterance. Anna takes the word 'key' to denote a lock-opening artifact, whereas John takes it to denote a kilo of cocaine.⁴ As a result, Anna and John disagree over the truth-conditional content of Anna's utterance. Anna believes that her utterance's truth-conditional content is the proposition^{5,6}

(3) Carla lost the door-opening artifact that would open locker L,
whereas John believes that that utterance's truth-conditional content is

(4) Carla lost the kilo of cocaine she was to deposit in locker L

Given that (3) is true and (4) is false, even if Anna knows that (3) is true, John may still know that (4) is false. In this way, Anna and John's dispute does not exhibit an epistemic conflict: it is entirely possible that, despite their dispute over the truth of (1), Anna knows

not obvious how to extend it to other kinds of disputes. I will abstract away from this issue in what follows.

⁴ The drug-denoting use of 'key' is usually spelled 'ki'. See paul (2004).

⁵ Here and elsewhere I will be taking some liberties with respect to the semantics of definite descriptions. See e.g. Fara (2001), Heim (1982), Schoubye (2013) for discussion.

⁶ As in previous chapters, sentences in this font stand for propositions.

that Carla lost the lock-opening artifact that would open a certain locker, while John knows that Carla did not lose the kilo of cocaine she was to deposit in the same locker.

The aim of this chapter is to explain why typical disputes exhibit epistemic conflicts while acknowledging that those disputes are, in a sense, verbal. In particular, I will identify a class of disputes which are not fully factual but nevertheless exhibit epistemic conflicts, and argue that typical disputes belong in that class. The structure of this chapter is as follows. I start by refining the notion of an epistemic conflict (section 2). Then I characterize the notion of a material conflict and show that, though all material conflicts exhibit epistemic conflicts, the notion of a material conflict doesn't capture the special way in which ordinary epistemic conflicts are typically grounded (section 3). Then I identify the class of partly factual disputes, show that they all exhibit epistemic conflicts, and show that partly factual disputes also capture the special way in which ordinary epistemic conflicts are typically grounded (section 4). After arguing that typical disputes are partly factual (section 5), I discuss some problems arising from the thesis I called Uncertainty, discuss further roles that the notion of a partly factual dispute can play, and conclude (sections 6–8).

3.2 Epistemic conflicts

I just said that a dispute exhibits an epistemic conflict just in case at most one of the disputants knows the facts she takes the utterance under dispute to be about. We can refine this characterization as follows:

Epistemic conflicts: A dispute between parties A and B over the truth of utterance U exhibits an epistemic conflict just in case, where P is the proposition A associates with U, and Q is the proposition B associates with U:

- (i) It's not the case that A knows that P is true and B knows that Q is false, and

- (ii) It's not the case that A knows that P is false and B knows that Q is true.

In this definition and throughout the chapter, when I say that a speaker *associates* a given proposition with a given utterance, what I mean is that that speaker believes that that proposition is that utterance's truth-conditional content.

Below I will say more about the importance of epistemic conflicts, and elaborate on how we could explain why typical disputes exhibit epistemic conflicts by assuming that they are fully factual. Before that, it is important to make four clarifications. The first concerns other phenomena one may intuitively think are involved in epistemic conflicts. For example, one may think that disputes which exhibit epistemic conflicts must involve a difference in the justification the participants in the dispute base their beliefs about the truth-value of the utterance under dispute. Or one may think that if a dispute exhibits an epistemic conflict, it must be a dispute worth having. My characterization of epistemic conflict is not meant to capture these other features a dispute may have.

According to the present characterization, a dispute may exhibit an epistemic conflict even though the participants in the dispute have no justification for their respective positions, or even if having that dispute would be unproductive. For example, take again the case in which Anna utters 'grass is green' and John replies 'no, it's not'. Even if John disputes the truth of Anna's utterance only to spite her, and even if neither Anna nor John have any justification whatsoever for their respective positions over the truth of Anna's utterance, their dispute may still exhibit an epistemic conflict. It will exhibit that kind of conflict as long as at most one of Anna and John knows the truth-value of the proposition she/he associates with Anna's utterance.⁷

The second clarification concerns a presupposition I made in characterizing epistemic conflict. My definition of epistemic conflict presupposes that typical disputants have deter-

⁷ Thanks to Jim Pryor for discussion.

minate beliefs about the truth-conditional content of the utterances their disputes are about. Thus, in conjunction with Uncertainty—the thesis that, for nearly every utterance, there is no proposition any language user takes to be that utterance’s truth-conditional content—the present definition would yield the result that no ordinary dispute exhibits an epistemic conflict. This is an important issue, and I will address it in section 3.6. For the time being, and merely for the sake of simplicity, I will continue to presuppose that typical disputants do have determinate beliefs about the truth-conditional content of the utterances their disputes are about.

The third clarification concerns some examples of disputes that exhibit epistemic conflicts which I will set aside for the purposes of the present discussion. An example of those cases is the following. Suppose that there is some proposition *P* such that Anna associates with utterance *U* the proposition *John knows that P is true*, and John associates *P* with *U*. Then if Anna claims that *U* is true and John claims that it is not, their dispute will exhibit an epistemic conflict. Or suppose that Anna associates *P* with *U* and her belief that *P* is true is based on her belief that *Q* is true, whereas John associates *Q* with *U* and his belief that *Q* is false is based on his belief that *P* is false. Then if Anna claims that *U* is true and John claims that it is not, their dispute will (presumably) exhibit an epistemic conflict even in a scenario in which *P* and *Q* have different truth-values.

What characterizes these cases is that the propositions the disputants associate with the utterance their dispute is about stand in special kinds of epistemic dependence to one another. In the first case, that relation of epistemic dependence is grounded in a logical relation of entailment (since the claim that John knows that *P* is true entails that *P* is true); in the second, it is grounded in a relation of “epistemic basing” (John’s belief that *Q* is false is based on his belief that *P* is false, and Anna’s belief that *P* is true is based on her belief that *Q* is true). Though nothing in the discussion that follows requires us to exclude these

kinds of epistemic conflicts, I will ignore them for the sake of simplicity. Our task here is not to explain why those kinds of epistemic conflicts arise, but only why typical disputes exhibit epistemic conflicts.⁸

The fourth clarification concerns cases that the present characterization counts as epistemic conflicts, but which some people may not want to count as such. For example, suppose that Anna makes utterance U, and John replies ‘that’s not true’. As it happens, Anna and John associate radically different propositions with U: Anna associates with it the proposition that a certain object x is at least half a meter tall, whereas John associates with U the proposition that a different object y weighs 5kg. As it happens, Anna doesn’t know that x is at least half a meter tall, and John doesn’t know that y doesn’t weigh 5kg. According to the present characterization, Anna and John’s dispute exhibits an epistemic conflict, but this isn’t due to any kind of conflict between Anna and John’s beliefs (not even a conflict in their actual truth-values).⁹ For present purposes, I am happy to admit that the present characterization may be too weak to accord with whatever intuitive idea we may have of such conflicts. For, presumably, typical disputes really do exhibit even the weaker kind of conflict characterized by the present definition, and it is important to explain why this is so despite the truth of Variance. In this way, though some features of epistemic conflicts intuitively understood will be left unexplained, the present discussion still constitutes progress insofar as it will explain some of those conflicts’ features. With these four clarifications in mind, let’s move on.¹⁰

Exhibiting epistemic conflicts is an important feature of typical disputes and, as such,

⁸ Thanks to Jim Pryor for discussion.

⁹ Thanks to Jim Pryor for bringing up this kind of case to my attention.

¹⁰ To some extent, the clarifications so far may suggest that the label ‘epistemic conflict’ is a misnomer. That is fine, since it is not crucial to the present discussion that the property of disputes I have characterized really corresponds to whatever intuitive notion we have of epistemic conflicts. What matters is that typical disputes really do exhibit epistemic conflicts as I have characterized them, and that exhibiting such conflicts is an important feature.

it is important to explain it. If ordinary disputes did not exhibit epistemic conflicts, there would seem to be no point in engaging in those disputes with the aim of discovering the truth about the facts we take our utterances to be about. For example, consider again Anna and John's dispute over the truth of Anna's utterance of "Carla lost the key". Since Anna associates the proposition (3) with her utterance, John associates (4) with the same utterance, (3) is true, and (4) is false, it may be that Anna knows that (3) is true all while John knows that (4) is false. In this way, it would be entirely pointless (even in principle) for Anna to engage in her dispute with John over the truth of her utterance with the aim of determining whether (3) is really true. It would be similarly pointless for John to engage in that dispute with the aim of determining whether (4) is really false. Explaining why typical disputes exhibit epistemic conflicts would thus amount to an explanation of why it would make sense to engage in those disputes with the aim of discovering the truth about certain facts (provided that the participants in those disputes are being sincere, have reasonably good evidence, and so on).^{11, 12}

As I said in the introduction, if Variance is true, typical disputes are not fully factual—

¹¹ Note that there may be many disputes which exhibit epistemic conflicts and may not be worth engaging in. I am not claiming that exhibiting an epistemic conflict is a sufficient condition for a dispute to be worth engaging in with the purpose of discovering truth. Rather, I only claim that exhibiting epistemic conflicts is a necessary condition for it to be worth in principle to engage in that dispute with such purposes. Thanks to Stephen Schiffer for discussion.

¹² Another role of the notion of an epistemic conflict concerns the role they may play in justifying the importance of certain debates in epistemology. For example, consider the debate in the epistemology of disagreement between equal-weight theorists and steadfasters. Equal-weight theorists (e.g. Christensen 2007 and Elga 2007) claim, among other things, that "learning that a peer disagrees with you about P gives you a reason to believe you are mistaken about P." Frances and Matheson (2018), while steadfasters (e.g. van Inwagen 1996 and Huemer 2011) deny it. The notion of an epistemic conflict can help explain why the debate between equal-weight theorists and steadfasters would be of interest when it comes to typical everyday disputes. Take for example Anna and John's dispute over the truth of Anna's utterance of (1). If Anna and John's dispute exhibits an epistemic conflict—i.e. if at most one of Anna and John could know the truth-value of the proposition they respectively associate with Anna's utterance of (1)—and Anna is as knowledgeable in matters pertaining to the truth of (3) as John is in matters pertaining to the truth of (4), then the question arises whether the fact that they have a dispute gives them reason to modify their beliefs about those propositions—i.e. whether the dispute gives Anna reason to modify her belief that (3) is true, and John reason to modify his belief that (4) is false. But if, on the other hand, Anna can know that (3) is true while John knows that (4) is false, it is not obvious why the fact of their disagreement would give them reason to revise their beliefs.

since, given Variance, the participants in typical disputes don't agree on what proposition is the truth-conditional content of the utterance under dispute. Though such explanation of epistemic conflicts is not available, it is easy to see its appeal. For, necessarily, if a dispute is fully factual, it exhibits an epistemic conflict.¹³ In this way, the claim that typical disputes are fully factual would afford us an especially robust explanation of the fact that they exhibit epistemic conflicts.

In the next section I will consider an alternative explanation of the fact that typical disputes exhibit epistemic conflicts. That explanation does not assume that typical disputes are fully factual but, as we will see, it is too weak to capture other important features of typical disputes.

3.3 Material conflicts

I just said that, if Variance is true, then typical disputes are not fully factual. If this is so, we can't explain why typical disputes are cases of epistemic conflicts by appealing to the fact that those disputes are fully factual. One potential explanation of the phenomenon at hand may appeal to the notion of a *material conflict*:

Material conflicts: A dispute between A and B over the truth of utterance U exhibits a material conflict just in case, where P is the proposition A associates with U and Q is the proposition B associates with U, P and Q are materially equivalent.

¹³ To see why, note first that if a dispute is fully factual, then there is a proposition all of its participants associate with the utterance the dispute is about. By bivalence and the factivity of knowledge, if one of the parties to the dispute knows that that proposition is true, then the other doesn't know that it is false, and if one of the parties knows that that proposition is false, then the other doesn't know that it is true. Since the argument does not rely on any special features of the utterance under dispute, the proposition the participants in the dispute associate with that utterance, or the participants in the dispute themselves, it follows that, necessarily, if a dispute is fully factual then it is epistemically significant.

Necessarily, if a dispute exhibits a material conflict, then it exhibits an epistemic conflict. For suppose A and B have a dispute over the truth of U, A associates P with U, B associates Q with U, and P and Q are materially equivalent. Then either P and Q are both true, in which case B doesn't know that Q is false and A doesn't know that P is false, or they are both false, in which case A doesn't know that P is true and B doesn't know that Q is true. Thus, if typical disputes exhibit material conflicts, that suffices to explain why they exhibit epistemic conflicts.

The problem with the present approach is not that it fails to explain the fact that typical disputes exhibit epistemic conflicts. In that respect, the appeal to the notion of material conflicts offers as good an explanation as any. Rather, the problem is that the fact that typical disputes exhibit material conflicts fails to explain another, related feature of those disputes. In order to characterize that feature, let me start with an example.

Suppose Anna and Carla are talking about Anna's new house, and Carla asks Anna to describe it. Referring to her house, Anna utters:

(5) It is green.

Let's suppose for simplicity that Anna and Carla both believe that 'It' as it occurred in Anna's utterance refers to *H*, Anna's house. Suppose further that Anna believes that the word 'green' applies to things of colors 2–5 (see figure 4.1), but, due to her idiosyncratic upbringing, Carla believes that it applies to things that are at least 10m tall. Due to their respective beliefs about the meaning of the word 'green', Anna and Carla associate different propositions with Anna's utterance. Anna associates the proposition (6) with her utterance, whereas Carla associates the proposition (7) with the same utterance:

(6) *H* is one of colors 2–5,

(7) *H* is at least 10m tall.

Carla, who feels skeptical that Anna's house is at least 10m tall, tells Anna:

(8) I don't believe it!

As it happens, Anna's house is color 4 and is 11m tall, so (7) and (6) are both true.



Figure 3.1: Some colors

Since (7) and (6) are both true, Anna and Carla's dispute exhibits a material conflict. And since every dispute that exhibits a material conflict exhibits an epistemic conflict, Anna and Carla's dispute also exhibits an epistemic conflict. That is, since (7) is true, Carla doesn't know that it is false.

Despite the fact that Anna and Carla's dispute exhibits an epistemic conflict, something about that dispute seems defective. In particular, even though (6) and (7) are both true, the facts in virtue of which (6) is true are completely orthogonal to the facts in virtue of which (7). For (6) is true in virtue of the fact that H is color 4, whereas (9) is true in virtue of the fact that H is 11m tall, and the fact that H is color 4 is completely unrelated to the fact that H is 11m tall. In this way, even though Anna and Carla's dispute exhibits an epistemic conflict, Anna and Carla are not talking about the same thing.

This defect has further repercussions for how Anna and Carla may go about resolving their dispute. Because (9) is true in virtue of the fact that H is 11m tall, and this fact is independent of facts about H's color, Carla would not change her mind about the truth-value of (9) no matter how much information Anna offered her about H's color, or how convincingly Anna proved that H is in fact color 4. In other words, given the description of the example, Anna and Carla's dispute could not be resolved by appealing to facts about the subject matter Anna takes her utterance to be about. The same is true the other way

around: because Anna doesn't take her utterance to be about H's height, Carla would not change Anna's mind no matter how many facts about H's height she appealed to.

Presumably, typical disputes are not like Anna and Carla's. One would expect that, in typical disputes, epistemic conflicts are grounded on facts all the participants to the dispute deem relevant to the truth of the utterance the dispute is about. The problem with the explanation of epistemic conflicts in terms of material conflicts is that it does not capture this further feature of typical disputes. The next section offers an alternative explanation of epistemic conflicts which captures that feature through the notion of a *partly factual dispute*.

3.4 Partly factual disputes

Consider the following case. 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 4.1, above). Anna just got a new house and is describing it to John. Referring to her house, she utters:

(5) It is green.

Suppose for simplicity that Anna and John both believe that 'It' as it occurred in Anna's utterance refers to *H*, Anna's house. Then Anna will associate (6) with her utterance, whereas John will associate (9) with the same utterance:

(6) *H* is one of colors 2–5,

(9) *H* is one of colors 3–6.

Though (6) and the negation of (9) are logically compatible, there are certain circumstances in which Anna and John could have a dispute over the truth of Anna's utterance which would not strike us as merely verbal. For example, suppose that Anna's house is in

fact color 4. Shortly after the conversation takes place John disputes the truth of Anna's utterance by saying

(10) That's not what I've heard!

and that, pointing to a photograph that clearly shows *H* is color 4, Anna replies,

(11) Well, you've heard wrong; my house is *this* color!

In such a case, assuming John has normal vision and given the proposition John associates with Anna's utterance, John could not reasonably reply something like

(12) 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 *merely* verbal. Though the dispute is certainly verbal in the sense characterized in the introduction, it is not *merely* verbal in that there is a fact responsible for the truth of both (6) and (9)—i.e. that *H* is color 4. Accordingly, keeping fixed the propositions Anna and John respectively associate with Anna's utterance, if Anna and John both knew that *H* is color 4, they should both believe that Anna's utterance is true.

Here is another example. Anna and John use 'Dumbo' to refer to different areas of Brooklyn. Anna uses 'Dumbo' to refer to the area constituted by areas A and B, and John uses 'Dumbo' to refer to the area constituted by B and C (see figure 3.2). John wants to meet with Anna, so he calls her to ask where she is. Anna replies:

(13) I'm in Dumbo.

Because of the difference in their beliefs about what area is denoted by 'Dumbo', Anna and John associate different propositions with Anna's utterance. Anna associates (14) with her utterance, whereas John associates (15) with it:




Figure 3.2: Some areas of Brooklyn.

(14) Anna is in area A or area B,

(15) Anna is in area B or area C.

As it happens, just as Anna utters (13), John sees her on the corner of Gold St. and York St., which Anna and John both know to be in area D. Surprised, John tells Anna over the phone:

(16) No, you're not! You're at Gold and York, I'm looking at you!

Though the truth of (14) is in principle compatible with the falsity of (15), in the present scenario there is a fact responsible for the falsity of both. In particular, both of those propositions are false in virtue of the fact that Anna is at the corner of Gold and York. Accordingly, once John points out that Anna is at the corner of Gold and York, Anna should modify her belief in the truth of (14). In this sense, Anna and John's dispute is not merely verbal.

For the final example, suppose Anna and John's beliefs about the meaning of the words 'green' and 'Dumbo' are as above. Anna has given John instructions to get to her house, which John has followed. Thinking he might be at the wrong place, John calls Anna to ask for some easy way of identifying her house, to which Anna replies:

(17) It's green, and it's in Dumbo.

John is sure he has followed Anna's instructions to the letter, but the house in front of him is in the corner of Gold and York and is color 2. Confused, John replies:

(18) Well, either it is not green, or it is not in Dumbo.

Suppose that both Anna and John take ‘it’ as it occurs in John’s utterance to refer to H, Anna’s house. Due to the difference in their beliefs about the meaning of ‘green’ and ‘Dumbo’, Anna associates (19) with Anna’s utterance, whereas John associates (20) with it:

(19) H is one of colors 2-5 and it is in area A or B,

(20) H is one of colors 3-6 and it is in area B or C.

The truth of (19) is logically compatible with the falsity of (20), but in the present scenario there is a fact responsible for the falsity of both. In particular, both propositions are false in virtue of the fact that H is in the corner of Gold and York. Accordingly, once John calls Anna’s attention to this fact, she should modify her belief in the truth of (19). In this sense, Anna and John’s dispute is not merely verbal.

Note that, unlike in the previous examples, not every fact responsible for the falsity of (20) is also responsible for the falsity of (19). In particular, (20) is false both because Anna’s house is color 2 and because H is in the corner of Gold and York, whereas (19) is false only because H is in the corner of Gold and York. Nevertheless, the present dispute is not purely verbal: there is a fact responsible for the falsity of the propositions the parties to the dispute associate with the utterance their dispute is about.

These three disputes are examples of what I will call *partly factual disputes*. What characterizes those disputes is that there is a fact responsible for the truth (falsity) of all the propositions the different parties to the dispute associate with the utterance the dispute is about. More precisely, using the framework of truthmaker semantics introduced in chapter 2, we can define partly factual disputes as follows:

Partly factual disputes: A dispute between parties A and B over the truth of an utterance U is partly factual just in case, where P is the proposition A associates with U and Q is the proposition B associates with U, P and Q have at least one truthmaker in common if one of them is true, or at least one falsitymaker in common if one of them is false.

Importantly, truthmakers and falsitymakers in this definition are meant to be *actual*, as opposed to merely possible. That is, a dispute over the truth of U is partly factual relative to a given possible world just in case some fact in that world which is a truthmaker for the propositions the parties to the dispute (respectively) associate with U, or a falsitymaker for them. For instance, Anna and John's dispute over Anna's utterance of (5) is partly factual because the fact that H is color 4 is a truthmaker for both (6) and (9). In contrast, in a possible world in which H was color 2, a dispute like Anna and John's would not be partly factual; in that scenario, the fact that H is color 2 would be the only *truthmaker* for (6) and the only *falsitymaker* for (9).

We can get a better grasp of partly factual disputes by comparing them with fully factual ones. The first difference of note is that every fully factual dispute is partly factual, but the converse is not true. Fully factual disputes are partly factual because if two people associate the same proposition with a given utterance, then, *a fortiori*, the propositions those two people associate with that utterance have the exact same truthmakers if they are true, or the same falsitymakers if they are false (since those propositions are one and the same).¹⁴ But, as Anna and John's disputes in this section illustrate, there can be partly factual disputes that are not fully factual.

The second important difference is this. If in a given possible world *w* two people

¹⁴ I'm assuming that propositions are finer-grained than mere sets of possible worlds. For example, on this assumption, the proposition that it rains or it doesn't is not identical to the proposition that either pandas are cute or they aren't.

have a *fully factual* dispute over a given utterance's truth, then in every possible world in which those two people have a dispute over that utterance's truth and associate the same propositions with that utterance as they did in w , their dispute will be fully factual.¹⁵ The same is not true of *partly factual* disputes. Even if we keep fixed the propositions the parties to a dispute associate with a given utterance, their dispute may count as partly factual in some possible worlds, but not in others. For example, as I just mentioned, in a possible world in which H is in fact color 2, the fact that H is color 2 is a truthmaker for (6) and a falsitymaker for (9); accordingly, in such a world, a dispute like Anna and John's dispute over Anna's utterance of (5) would not be partly factual.

On the other hand, an important similarity between fully factual disputes and partly factual ones concerns their relationship to epistemic conflicts. In section 3.2 I said that, necessarily, if a dispute is *fully* factual, it exhibits an epistemic conflict. The same is true of *partly* factual ones: necessarily, if a dispute is partly factual, then it exhibits an epistemic conflict.

To see why this is so, it suffices to note that, necessarily, if a dispute is partly factual, it exhibits a material conflict. For, necessarily, if two propositions have at least one actual truthmaker in common, they are both true, and if they have at least one actual falsitymaker in common, they are both false—since every possible truthmaker necessitates the truth of the propositions it is a possible truthmaker for if it obtains, and every possible falsitymaker necessitates the falsity of the propositions it is a falsitymaker for if it obtains. And, as we saw above (p. 111, it is necessarily the case that if a dispute exhibits a material conflict, it exhibits an epistemic conflict.

¹⁵ I'm being somewhat casual about the identity conditions of utterances. If utterances are "world-bound", the present condition should instead read: if in a given possible world w two people have a fully factual dispute over an utterance U , then in every possible world in which those two people have a dispute over the truth of an utterance U' , and associate with U' the same propositions that they associated with U in w , their dispute will be fully factual. The difference between assuming that utterances are world-bound and assuming that they are not does not matter for present purposes.

In this way, if typical disputes are partly factual, that would (via the fact that partly factual disputes exhibit material conflicts) explain why they exhibit epistemic conflicts. Furthermore, for any utterance *U*, if a dispute over *U*'s truth is partly factual, then the fact that the propositions the parties to the dispute associate with *U* have the same truth-value is grounded in at least some facts that all the participants in the dispute would deem relevant to determining *U*'s truth-value. Thus, if typical disputes are partly factual, that would also explain why those disputes can be resolved—at least in principle—by appealing to facts all of the participants deem relevant to the truth-value of the utterance their dispute is about.

These observations take us one step closer to answering the challenges from Variance. The first of those challenges was to explain why ordinary disputes exhibit epistemic conflicts while accepting that they are, in a sense, verbal. We saw that this challenge is easily met through the claim that typical disputes exhibit material conflicts, but we also saw that material conflicts don't seem to capture all the interesting features of typical disputes. This gave rise to a second challenge: to explain why typical disputes can be resolved, at least in principle, by appealing to facts all the participants in the dispute deem relevant to the truth-value of the utterance their dispute is about.

If typical disputes are partly factual, we can answer the second challenge as well. Since epistemic conflicts arising from partly factual disputes are grounded in facts all the participant to the dispute deem relevant, if typical disputes are partly factual, we could explain why appealing to such facts would put the participants in a dispute in a position to solve it. The next section completes this explanation by arguing that everyday disputes are typically partly factual.

3.5 Are typical disputes partly factual?

This section argues for two claims. First, that Variance is entirely compatible with the claim that ordinary disputes are partly factual. Second, that, typical disputes are likely to be partly factual. I start with the first of these claims.

Recall the case for Variance. In chapter 1 I said that, for nearly every utterance, there is an enormous number of equally natural and extremely similar propositions any language user could easily have associated with that utterance. Given the huge number of such propositions, even small differences in the evidence that determines what proposition to associate with the utterance, in the way someone is attuned to that evidence, and so on, are likely to produce differences in what proposition that person associates with that utterance. And since different language users rarely have the exact same evidence or are attuned to the evidence in exactly the same way, it is extremely unlikely that they associate the same proposition with the utterance in question.

Take for instance Anna's utterance of (5)—'it is green'. For the sake of simplicity, in the previous section I wrote as if there were only seven colors (see fig. 1), but the actual situation is far more complex. For example, there are many points in the color spectrum between color 2 and color 3 at which Anna or John could easily have drawn the line between things called 'blue' and things called 'green'. Accordingly, there are many different propositions Anna and John could easily have associated with Anna's utterance of (5), each corresponding to one way of drawing the line between things called 'green' and things called 'blue'. Given the huge number of such propositions, it would be extremely unlikely for Anna and John to associate the exact same one of those propositions with the utterance in question.

Observations of this kind hold for utterances of sentences involving terms whose appli-

cation depends on the properties an object has along one or more continuous dimensions. In order to see this, suppose for example that ‘F’ is a predicate of that kind and ‘a’ is a proper name. For any language user, there will be a huge number of extremely similar and equally natural (or otherwise eligible) properties—each corresponding to a slightly different cutoff point along one or more of the dimensions on which the application of ‘F’ depends—such that that language user could easily have believed that an utterance of ‘a is F’ has the truth-conditional content that *a* (the object denoted by ‘a’) has that property. If, for example, a language user believes that ‘F’ expresses the property corresponding to the cutoff point *x* along one of the dimensions relevant to the application of ‘F’, 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. Given the huge number of properties any language user could easily have believed ‘F’ to express, it is extremely unlikely for any two language users to believe that an utterance of ‘a is F’ has the truth-conditional content that *a* has the exact same one of those properties. Variance is supported by the fact that we can build arguments of this kind for every utterance involving a predicate whose application depends on one or more continuous dimensions (or even just dimensions with enormously many points).

The case for Variance shows that, for nearly every utterance, no two people associate the same proposition with that utterance. What the argument does *not* show is that the propositions those people associate with a given utterance will be so different that those people are generally incapable of having a partly factual dispute over that utterance’s truth—i.e. that those propositions share at least an actual truthmaker or at least an actual falsitymaker.¹⁶ Anna’s dispute with John illustrates this: in a possible world in which H is color 4, their dispute over the truth of Anna’s utterance of (5) is partly factual even though they don’t as-

¹⁶ I’m abstracting away from issues concerning Uncertainty for the time being. I address those issues in section 3.6.

sociate the same proposition with that utterance. More generally, it is entirely compatible with Variance that typical disputes are partly factual.

Now, the observation that typical disputes can be partly factual despite the truth of Variance takes us only half the way. Ideally, we would also like an argument to the effect that typical disputes really are partly factual. I will now provide such an argument. In particular, I will argue that the truth-conditional candidates for a given utterance generated in the way I just described are very likely to have at least one actual truthmaker (falsitymaker) in common. If that is the case, then as long as the participants in a dispute over an utterance's truth associate one of those truth-condition candidates with that utterance, their dispute is very likely to be partly factual.

Before proceeding to the arguments, I want to make two clarifications. The first is that, in arguing that typical disputes are partly factual, I don't mean to claim that *all* disputes are partly factual. Many everyday disputes resulting from misunderstandings or ambiguities do not even exhibit epistemic conflicts; in consequence, they are not partly factual either. This brings us to the second clarification, which is that I am assuming that, for the most part, typical disputes really do exhibit epistemic conflicts. In this sense, I will have very little to say in reply to the skeptic who thinks that, for the most part, the participants in ordinary disputes are just talking past each other.

To begin with, recall an important assumption from the case for Variance. That assumption is that the truth of the various truth-conditional content candidates an utterance has depends on the same object's properties along the same dimensions. The difference between the various truth-conditional content candidates an utterance may have concerns only the exact region in the space generated by those dimensions where that object's properties must fall in order for those candidates to be true. By assumption, that difference is very slight, so the region in the space generated by those dimensions in which the rele-

vant object must fall in order for two truth-conditional content candidates to have different truth-values is very small. In this way, it is unlikely that one of those two candidates would be true and the other false.¹⁷

On the basis of that same assumption, we can argue that if any two truth-conditional content candidates have the same truth-value, then they have at least one truthmaker in common if they are true, or at least one falsitymaker in common if they are false. In chapter 2 I said that if the truth of an atomic proposition Fx depends on x 's properties along various dimensions, then Fx 's possible truthmakers and falsitymakers are the minimal states of affairs that specify all of x 's maximally determinate properties along those dimensions. In particular, Fx 's possible truthmakers will be the minimal such states in which Fx is true, and its possible falsitymakers will be the minimal such states in which Fx is false. Thus, if the truth of two atomic propositions depends on the same object's properties along the same dimensions, those propositions will have the same truthmakers if one of them is true, and the same falsitymakers if one of them is false.

For example, consider the propositions *H is at least 11m tall* and *H is at least 11.1m tall*. The truth of both propositions depends exclusively on H's height. Accordingly, one possible truthmaker for both of them will be a minimal state of affairs in which H is exactly 11.5m tall; another will be a minimal state of affairs in which H is exactly 13m tall; etc. On the other hand, one possible falsitymaker for both *H is at least 11m tall* and *H is at least 11.1m tall* will be a minimal state of affairs in which H is 9m tall; another will be a minimal state in which H is 3m tall, and so on. In fact, the only possible states that are a truthmaker for one but not for the other will be states in which H is less than 11.1m tall, but at least 11m tall. Those possible states correspond to situations in which only one of the two propositions is true. Thus, if both *H is at least*

¹⁷ See chapter 4, section 6, for a precisification of this line of reasoning.

11m tall and H is at least 11.1m tall have the same truth-value, they will either have the same truthmakers or the same falsitymakers.

Now, though we can't argue for a similarly strong conclusion in the case of disjunctions and conjunctions, we can still make a case that typical disputes over the truth of conjunctions and disjunctions will be partly factual. Let's start with disjunctions. As long as P and R are very similar and depend on the same object's properties along the same dimensions, and Q and S are very similar and depend on the same object's properties along the same dimensions, $P \vee Q$ and $R \vee S$ will be likely to have at least one truthmaker in common if they are true, and will definitely have the same falsitymaker if they are false. For example, let P be H is at least 11m tall, Q be H is a color between 2 and 5, R be H is at least 11.1m tall, and S be H is a color between 2.1 and 5.1. Then the scenarios in which $P \vee Q$ and $R \vee S$ are both true but have no truthmakers in common will be those in which H is a height between 11 and 11.1m tall, and either a color between colors 2 and 2.1, or a color between colors 5 and 5.1. These scenarios constitute a very small portion of the possible scenarios in which the two disjunctions can be true, which makes it very unlikely that they obtain. That both disjunctions must have the same falsitymakers if they are false follows from the definition of falsitymakers for disjunction, the fact that the truth of P and R depends on the same object's properties along the same dimensions, and the fact that the truth of Q and S depends on the same object's dimensions.¹⁸

¹⁸ To see this, let s be an actual falsitymaker for $P \vee Q$. By the definition of falsitymakers for disjunctions, s must be the fusion of a falsitymaker for P, call it s_P^- , and a falsitymaker for Q, call it s_Q^- . Now, since P and R depend on the same object's properties along the same dimensions, s_P^- must be either a truthmaker or a falsitymaker for R. Since, by assumption, $R \vee S$ is false, s_P^- cannot be a truthmaker for R (otherwise, $R \vee S$ would be true), so it must be a falsitymaker. Similarly, since Q and S depend on the same object's properties along the same dimensions, s_Q^- must be either a truthmaker or a falsitymaker for S. But, by assumption, $R \vee S$ is false, so s_Q^- must be a falsitymaker for S. Finally, since the falsitymakers for $R \vee S$ are all the fusions of falsitymakers for R and S, s itself must be a falsitymaker for $R \vee S$. And since all the states involved in the present argument were chosen arbitrarily, it follows that if $P \vee Q$ and $R \vee S$ are both false, every falsitymaker for the former is a falsitymaker for the latter. An analogous argument shows that the same is true the other way around.

Analogous arguments show that if P and R are very similar and depend on the same object's dimensions, and Q and S are very similar and depend on the same object's dimensions, $P \wedge Q$ and $R \wedge S$ have the same truthmakers if they are both true, and are likely to have the same falsitymakers if they are both false.

This ends my arguments to the effect that typical disputes are likely to be partly factual. Those arguments rely on the assumption that the truth of all the truth-conditional content candidates for typical utterances depends on the same object's properties along the same dimensions. In the present dialectical context, that assumption seems warranted. However, an objector may claim that the same kind of reasons that support Variance and Variance* may also support *Dimensional Variance*, the thesis that, for nearly every utterance and any two language users, there is no set of dimensions such that those two language users believe that that utterance's truth-conditional content is a proposition whose truth depends exactly on the same object's properties along those dimensions. If Dimensional Variance is true, even partly factual disputes may be difficult to come by.

I don't think Dimensional Variance is true. In my development of the case for Variance in chapter 1, I offered several examples showing that, given a fixed set of dimensions, we can generate an enormous number of extremely similar and equally plausible truth-conditional content candidates for one and the same utterance. For example, when I presented the case for Variance, I said that there are many things one may reasonably call 'quesadillas', depending on the amount of cheese they had. Given the huge number of such different plausible denotations for the word 'quesadilla', it would be unlikely that any two language users took that word to have exactly the same denotation. And if that's the case, then (assuming basic principles of compositionality) it would be unlikely that any two language users took an utterance involving the word 'quesadilla' to have the exact same truth-conditional content.

I am skeptical that the present objector can offer similar examples which can motivate Dimensional Variance. For example, consider again the word ‘quesadilla’. Presumably, that word’s application depends not only on an object’s quantity of cheese, but also on its shape, size, how melted the cheese is, and perhaps on the amount of corn it has. It is difficult enough to think of further (natural enough) dimensions on which the application of the word quesadilla may depend, let alone enormously many equally natural and extremely similar sets of dimensions on which the application of the word ‘quesadilla’ may depend. And without a big enough number of such sets for a big enough number of utterances, the case for Dimensional Variance does not get off the ground.

To conclude, I acknowledge that Dimensional Variance really would be problematic for the view I articulated here. However, as I have noted, it is difficult to see how a compelling case for Dimensional Variance could be made. The next section considers another kind of problem for my claim that typical disputes are partly factual. That problem arises from the thesis I have called *Uncertainty*.

3.6 Troubles with Uncertainty

So far I have written as if people have determinate beliefs about the truth-conditional content of utterances they have disputes about. For example, throughout my discussion I assumed that Anna associates the proposition (6) with her utterance of (5) and that John associates the proposition (9) with the same utterance. However, as I pointed out in chapter 1, the kind of observations that support Variance also support the further thesis I called Uncertainty. Uncertainty is the thesis that ordinary language users rarely have determinate beliefs about the truth-conditional content of ordinary utterances. According to Uncertainty, ordinary language users will typically be undecided as to which of a large number

of propositions is a given utterance's truth-conditional content.

Uncertainty is problematic for my claim that typical disputes are partly factual. When I introduced the definition of partly factual disputes in section 3.4, I presupposed that the participants in partly factual disputes had determinate beliefs about the truth-conditional content of the utterances their disputes are about. Yet if Uncertainty is true, this presupposition rarely holds, and typical disputes are not partly factual. Accordingly, if Uncertainty is true, we can't explain the fact that knowledge of extralinguistic facts is at stake in typical disputes on the basis of the claim that typical disputes are partly factual.

We can address this problem by modifying the definition of a partly factual dispute so as to accommodate Uncertainty. I will start with an example. Suppose we name each color in the color spectrum between colors 1 and 7 with a real number. So, for example, color 2.5 will be some color between colors 2 and 3, and color 4.7 will be some color between colors 4 and 5. Anna, let's suppose, is undecided as to which of propositions (21)–(25) to associate with her utterance of (5): she believes that her utterance's truth-conditional content is one of those propositions, but she is just as certain that it is (21) as she is that it is (22), that it is (23), and so on.

(21) H is one of colors 2.15–4.95

(22) H is one of colors 2.3–4.8

(23) H is one of colors 2.2–4.9

(24) H is one of colors 2.05–5.1

(25) H is one of colors 2.1–4.7

John, on the other hand, is undecided as to which of propositions (26)–(28) to associate with the same utterance:

(26) H is one of colors 2.95–6.05

(27) H is one of colors 3.1–5.9

(28) H is one of colors 3.15–5.85

Like in the original example, suppose Anna's house is in fact color 4, and that John disputes the truth of Anna's utterance by saying 'that's not true' right after Anna utters (5). If Anna replied by uttering 'yes it is, my house is *this* color!' pointing to her house in a photograph that clearly shows H is color 4, it would still be unreasonable for John to reply something like

(29) Wait a moment! The fact that the house is that color does not make what you said true! That color is not green!

Though Anna and John don't associate any particular proposition with Anna's utterance, if Anna and John both know that H is color 4, they should both think that Anna's utterance is true. This is so because the fact that Anna's house is color 4 is responsible for the truth of (21)–(28). In this respect, Anna and John's present dispute closely resemble partly factual disputes.

Examples of this kind suggest we can modify the original definition of partly factual disputes so as to accommodate Uncertainty as follows:

Partly factual disputes*: A dispute between parties A and B over the truth of an utterance

U is partly factual* just in case:

- (i) If one of the propositions A takes to be live candidates for U's truth-conditional content is true, then each of those candidates has an actual truthmaker in common with a proposition B takes to be a live candidate for U's truth-conditional content, and vice versa—i.e. each proposition takes to be a live candidates

for U's truth-conditional content has an actual truthmaker in common with a proposition A takes to be a live candidate.

- (ii) If one of the propositions A takes to be live candidates for U's truth-conditional content is false, then each of those candidates has an actual falsitymaker in common with a proposition B takes to be a live candidate for U's truth-conditional content, and vice versa.^{19, 20}

Strictly speaking, the introduction of this new definition solves only one part of the problem. Above I said that if typical disputes are partly factual then they exhibit an epistemic conflict. Unfortunately, the definition of epistemic conflicts itself presupposes that Uncertainty is false, so we need a new way of defining what it is for knowledge of the facts the participants in a dispute take the utterance under dispute to be about to be at stake.

We can address this problem by modifying the notion of epistemic conflicts as well:

Epistemic conflicts*: A dispute between parties A and B over the truth of utterance U exhibits an epistemic conflict* just in case:

- (i) A doesn't know that any of the propositions she takes to be live-candidates for U's truth-conditional content is true, or B doesn't know that any of the propositions she takes to be live candidates for U's truth-conditional content is false, and

¹⁹ Note that if it is vague whether A or B take P to be a live candidate for U's truth-conditional content, it may be similarly vague whether A and B's dispute over the truth of U is partly factual*.

²⁰ For the sake of simplicity, I'm ignoring that language users may assign different weights to the different propositions they consider to be candidates for an utterance's truth-conditional content. For our purposes, we may consider a proposition to be a *live* candidate for a given language user just in case that language user assigns it a weight above a certain threshold. Note also that, given the present definition, for a dispute to be partly factual* all the propositions A takes to be live candidates for the truth-conditional content of U must have the same truth-value, as should all the propositions B takes to be live-candidates for the truth-conditional content of U. Thanks to Cian Dorr, Jim Pryor, and Stephen Schiffer for discussion of these issues in relation to the definition of partly factual disputes*.

- (ii) A doesn't know that any of the propositions she takes to be live-candidates for U's truth-conditional content is false, or B doesn't know that any of the propositions she takes to be live candidates for U's truth-conditional content is true.

If a dispute exhibits an epistemic conflict*, then at most one of the parties to the dispute knows the truth-value of the propositions she takes to be live candidates for being the truth-conditional content of the utterance the dispute is about.

Necessarily, if a dispute is partly factual* then it exhibits an epistemic conflict*. Thus, as long as typical disputes are partly factual*, we will be in a position to explain why typical disputes exhibit epistemic conflicts*. Furthermore, the kind of epistemic conflicts* exhibited by partly factual disputes* are grounded in facts all the participants in the dispute deem relevant to determining the truth of the utterance their dispute is about.

The case for the claim that typical disputes are partly factual* remains the same. If the truth of the different truth-conditional content candidates for a given utterance depends on the same object's properties along the same dimensions, it is very likely that they will all have the same truthmakers if one of them is true, and the same falsitymakers if one of them is false.

To summarize, I have acknowledged that the truth of Uncertainty poses problems for the explanations I developed in sections 3.4–3.5. However, I have also shown that we can modify the original definitions of partly factual disputes and epistemic conflicts so as to accommodate Uncertainty. Once we adopt the revised definitions, we are in a position to recover the original explanation from sections 3.4–3.5. The next section discusses further roles that the notion of a partly factual dispute may play. For the sake of simplicity, I will go back to writing as if Uncertainty was false. The reader can rest assured that the difference between partly factual and partly factual* disputes does not matter for the purposes of that

discussion.²¹

3.7 Further work for partly factual disputes

Throughout this chapter I have argued that the notion of a partly factual dispute can play an important explanatory role. I have argued that if typical disputes are partly factual, that explains why they exhibit epistemic conflicts and why those conflicts are grounded in facts all the participants in the dispute deem relevant to its resolution. In this way, the fact that typical disputes are partly factual explains why knowledge of the facts the participants in the dispute take the utterance to be about is at stake. In this section I will discuss two further theoretical roles that partly factual disputes can play. Those roles concern a debate about contextualism and disagreement, and a debate over what it takes for a dispute to be purely verbal.

3.7.1 Contextualism and the problem of lost disagreement

Contextualism is the view that many words besides the standard indexicals (e.g. ‘I’, ‘you’, ‘here’, ‘now’) are context sensitive: the truth-conditional content of utterances involving those words depends on the context in which the utterance is made. A common objection to contextualism is that it cannot account for various forms of disagreement.

²¹ As in previous chapters, if it turns out that language users are usually uncertain as to which truth-conditional content candidates they are uncertain between, we can iterate the present strategy as many times as necessary. For example, suppose that S is uncertain as to which of sets $\Gamma_1, \Gamma_2, \dots, \Gamma_n$ is such that S is undecided as to which of its members is U’s truth-conditional content. Then the set of propositions that should play a role in determining whether a dispute over the truth of U is partly factual* should be the union of $\Gamma_1, \Gamma_2, \dots, \Gamma_n$. Alternatively, we could just take the set Γ such that S is uncertain as to which of the propositions in Γ is U’s truth-conditional content, and it is definitely the case that S is uncertain as to which of the propositions in Γ is U’s truth-conditional content, and it is definitely definitely the case that... and so on, ad infinitum. I hope to explore these options in more detail in future work. Thanks to Stephen Schiffer for pressing me on these issues.

Take for example certain kinds of contextualism about taste predicates, according to which in uttering ‘apples are delicious’ the speaker says that apples are delicious relative to her standard of taste, while in replying ‘apples are not delicious’, her interlocutor says that apples are tasty relative to her own standard of taste. According to MacFarlane (2007),

If in saying ‘apples are delicious’ I am saying that they taste good *to me*, while in saying ‘apples are not delicious’ you are denying that they taste good *to you*, then we are no more disagreeing with each other than we would be if I were to say ‘My name is John’ and you were to say ‘My name is not John’. Intuitively, though, it does seem that we are disagreeing. (p. 18)

The problem, according to MacFarlane (2007, 2014) and others (e.g. Stephenson 2007, Lasersohn 2005, Kolbel 2004, Egan 2010, Huvenes 2012) is that contextualism about predicates of personal taste is in conflict with our intuition that disputes like the one in the example are genuine disagreements.²²

Or take more general remarks like those by Cappelen and Lepore (2005). According to them, if contextualism is true, “it would be a miracle if speakers in different contexts were ever able to agree, disagree, or more generally, share contents” (p. 124). According to theorists like Cappelen and Lepore, if contextualism is true then people who find themselves in different contexts would rarely associate the same propositions with a given utterance, in which case cross-contextual disagreements, agreements, and communication more generally, will be very rare.²³

²² Some people raise similar objections against certain forms of contextualism about moral terms. See Khoo and Knobe (2016), Khoo (forthcoming) for discussion.

²³ To clarify, the objection is that if contextualism is true, then genuine disagreements about the subject matter of the utterance under dispute will be very rare. For example, genuine disagreements about taste will be very rare. But, needless to say, objectors to contextualism will be happy to accept that contextualism is compatible with genuine disagreements about language: e.g. genuine disagreements whether this or that utterance is true.

Objections of this kind are often called objections from “lost disagreement”. Objections from lost disagreement may be global if they lead to the conclusion that contextualism in general is problematic, and local if they lead to the conclusion that certain forms of contextualism about certain particular kinds of discourse are problematic. The discussion so far puts us in a position to reply to these objections on behalf of contextualists. Objections from lost disagreement often presuppose that in order for a dispute over the truth of an utterance to constitute a genuine disagreement about extra-linguistic matters, they must associate the same proposition with that utterance.²⁴ Or, at the very least, they presuppose that in order for a dispute to constitute a genuine disagreement about extra-linguistic matters, the parties to the dispute must associate propositions that cannot all be accurate with the utterance under dispute. Call presuppositions of this kind *full-factuality presuppositions*

The case for Variance shows that objections from lost disagreement, whether in their global or in their local varieties, do not pose a special problem for contextualism. In particular, if Variance is true, full-factuality presuppositions are too strong. In conjunction with Variance, incompatibility presuppositions entail that disputes in which one of the parties makes an utterance and the other rejects it rarely count as genuine disagreements. For if in making an assertoric utterance the speaker expresses her belief that the proposition she associates with the utterance is true, and in rejecting an utterance the audience expresses

²⁴ The requirement is sometimes weakened so that a dispute about an utterance’s truth can be genuine as long as the disputants associate incompatible propositions with the utterance their dispute is about. See e.g. Egan (2010), who states:

There is a genuine conflict between an assertion of S by A and an assertion of \neg S by B iff neither party can consistently accept the other’s assertion without withdrawing, ceasing to stand by, and ceasing to be prepared to repeat their own. That is, A cannot consistently both stand by her original assertion and remain willing to assert S, while simultaneously accepting B’s assertion of \neg S, and B cannot consistently stand by her original assertion and remain willing to assert \neg S, while simultaneously accepting A’s assertion of S. When two assertions are in genuine conflict, each party’s making, and continuing to endorse, its own assertion commits it to rejecting the other’s. (p. 255)

The difference does not matter for present purposes, though see Abreu Zavaleta (Manuscript) for a more detailed discussion.

her belief that the proposition she associates with the utterance is false, the participants in ordinary disputes will often express their beliefs in entirely compatible claims. Since, presumably, genuine disagreements are very common, and provided that Variance is true, arguments from lost disagreement seem to pose more of a problem for incompatibility presuppositions than for contextualism itself.²⁵

Needless to say, it is one thing to point out that the problem of lost disagreement is much more general than opponents of contextualism seem to admit, and another to explain what genuine disagreements about non-linguistic matters consist in. This is where partly factual disputes can play an important role. So far I have argued that if a dispute over an utterance's truth is partly factual, then at most one of the participants in the dispute knows the facts she takes the utterance to be about. To some extent, this captures the idea that typical disputes are not exclusively about linguistic matters, but using the apparatus of truthmaker semantics we can go further and define what partly factual disputes are about:

Dispute aboutness: If A associates P with an utterance U, B associates Q with the same utterance, and their dispute is partly factual, A and B's dispute is about subject matter M just in case: every possible state of affairs in M is part of a possible state of affairs in P's exact subject matter and of a possible state of affairs in Q's exact subject matter, and (ii) the truthmakers (falsitymakers) for P and Q in the possible world in which the dispute takes place are in M.

As a reminder, subject matters are sets of possible states of affairs (see chapter 2), and a proposition's exact subject matter is the set of all its possible truthmakers and falsitymakers. The present definition states that every subject matter that a dispute is about must be

²⁵ Criticisms of incompatibility presuppositions are not new. For example, Sundell (2011) argues that disagreements about taste are really disagreements about what standard of taste to adopt and claim that such disagreements deserve to be called genuine, and Khoo and Knobe (2016) adopt a similar line on disagreements involving moral terms. What is new is the observation that if incompatibility presuppositions were true, then there would be nearly no genuine disagreements about non-linguistic matters.

encompassed by the subject matters of the propositions the participants in the dispute respectively associate with the utterance the dispute is about.²⁶ Among the subject matters a partly factual dispute is about, the dispute's *exact* subject matter is the subject matter M which encompasses all the subject matters the dispute is about, and such that any other subject matter that encompasses them encompasses M.

For example, recall Anna and John's dispute over the truth of Anna's utterance of (5)—'It is green'. Anna associates the proposition (6)—H is one of colors 2-5—with that utterance, but John associates the proposition (9)—H is one of colors 3-6—with the same utterance. I said in section 3.4 that Anna and John's dispute is partly factual in a possible world in which H is color 4, since that fact that is color 4 is a truthmaker for both (6) and (9). However, their dispute is not about the truth-value of (6), since John may have no opinion about its truth-value. Nor can the dispute be about the truth-value of (9), since Anna may have no opinion about its truth-value. So what is Anna and John's dispute about? According to the definition above, since (9) and (6) are both about H's color, so is Anna and John's dispute.²⁷

Needless to say, a full defense against particular local versions of the objection from lost disagreement should involve offering an empirically adequate semantics for those expressions which explains why disputes involving those expressions are typically partly factual. However, the present observations suffice to respond to more general arguments from lost disagreement—i.e. arguments like Cappelen and Lepore's—against contextualism. Using the notion of a partly factual dispute, contextualists are in a position to explain both the

²⁶ Recall that a subject matter encompasses another just in case every state in the latter is part of a state in the former.

²⁷ Here is a toy model. Let's say there are only seven possible states of affairs: s_1, \dots, s_7 , each corresponding to the state in which H is color 1, color 2, and so on. (6)'s truthmakers are s_2, \dots, s_5 , and its falsitymakers are s_1, s_6 and s_7 . (9)'s truthmakers are s_3, \dots, s_6 are all truthmakers for (6), and its falsitymakers are s_1, s_2 , and s_7 . Given these assumptions, (9) and (6) have the same subject matter, namely, the subject matter represented as the set containing states s_1, \dots, s_7 . Accordingly, the subject matter of Anna and John's dispute is also the subject matter represented as the set containing states s_1, \dots, s_7 .

sense in which there can be genuine cross-contextual disagreements, and the reason why such disagreements exhibit epistemic conflicts.

3.7.2 Purely verbal disputes

There is something dismissive about calling a dispute purely verbal: if a dispute is purely verbal, the parties in the dispute may both be right about the facts they take the utterance under dispute to be about; in that sense, there seems to be little point in engaging in that dispute. Using the notion of a partly factual dispute, we can offer a characterization of purely verbal disputes that explains this feature of purely verbal disputes. In particular, we can characterize *purely verbal* disputes as follows:

Purely verbal disputes: A dispute over the truth of an utterance is purely verbal just in case it is not partly factual.

If a dispute is not partly factual, then the parties to the dispute might each know the facts they take the utterance to be about. For example, consider a possible world in which Anna's house is color 2 and Anna and John have a dispute over the truth of (5). Since the fact that Anna's house is color 2 is a truthmaker for the proposition that said house is one of colors 2–5 and a falsitymaker for the proposition that it is one of colors 3–6, it may be that Anna knows that her house is one of colors 2–5 while John knows that Anna's house is not one of colors 3–6. In this sense, Anna and John may each know the facts they take Anna's utterance to be about. Accordingly, there is a sense in which it is a waste of time for Anna and John to debate whether (5) is true.

The present account of purely verbal disputes has significant advantages over a popular family of competitors. According to those competitors, a dispute over the truth of U is merely verbal if one of the parties believes that U has the truth-conditional content P, the

other believes that U has the truth-conditional content Q , and P and $\neg Q$ are consistent.²⁸ Given that partly factual disputes satisfy this characterization, the present approach does not capture the intuition that purely verbal disputes are pointless. For, as I argued in section 3.4, partly factual disputes exhibit epistemic conflicts; as such, at most one of the participants in a partly factual dispute knows the facts she takes the utterance under dispute to be about. My approach is superior insofar as it captures this important feature of merely verbal disputes.

My preferred approach to purely verbal disputes also has significant advantages over less standard characterizations. For example, according to Chalmers (2011)

A dispute over S is (broadly) verbal when, for some expression T in S , the parties disagree about the meaning of T , and the dispute over S arises wholly in virtue of this disagreement regarding T . (p. 522)

Taken at face value, this characterization seems extensionally inadequate. For example, suppose Anna's house is color 2, Anna utters (5) referring to her house, and John disputes the truth of Anna's utterance just to spite her. In such a case, the dispute seems merely verbal even though it does not arise *wholly* in virtue of a disagreement regarding the meaning of the word 'green'; it arises *partly* in virtue of John's desire to upset Anna. Or take a case in which Anna's house is color 4, John knows that he and Anna disagree about the meaning of the word 'green', and disputes the truth of Anna's utterance on that basis without knowing that Anna's house is color 4. In such a case, Anna and John's disagreement would be partly factual even though it would have arisen solely in virtue of Anna and John's disagreement over the meaning of the word 'green'. Yet we have seen that, though there is an element of mere verbalness to some partly factual disputes, they exhibit epistemic conflicts

²⁸ As I said in footnote 3, people who endorse this general approach include Hirsch (1993, p.181), Sider (2006, p.76), and Vermeulen (2018).

which are grounded in facts all the participants in the dispute would deem relevant to the truth of the utterance the dispute is about. In that sense, partly factual disputes are not *purely* or *merely* verbal.

Thus, Chalmers' characterization doesn't capture the pointlessness we tend to associate with purely verbal disputes. To be fair, Chalmers acknowledges that there are some senses of 'in virtue of' that do not really capture what he meant to say. However, absent a more precise characterization of what it is for a disagreement to arise "wholly in virtue of a disagreement regarding an expression", we have enough motivation to look for a better alternative. I believe my proposal is such an alternative.²⁹

In a way, the differences between my preferred approach to purely verbal disputes and Chalmers' may be due to a difference in explanatory aims. As it is, Chalmers' characterization seems like seems to track the motivational grounds on the basis of which we may think a dispute is purely verbal. In contrast, my characterization tracks what we may call "alethic" grounds on the basis of which we may think that a dispute is purely verbal.

Finally, consider Jenkins' (2014a, 2014b) account of merely verbal disputes. According to Jenkins,

Parties A and B are having a merely verbal dispute iff they are engaged in a sincere prima facie dispute *D*, but do not disagree over the subject matter(s) of *D*, and merely present the appearance of doing so owing to their divergent uses of some relevant portion of language. (Jenkins 2014a, p.21)³⁰

Jenkins does not offer an account of what it is for the participants in a dispute to disagree over the dispute's subject matter. Absent such an account, Jenkins' characterization does

²⁹ Note that defenders of Chalmers' approach need not reject my preferred characterization. In particular, they may co-opt my preferred approach and claim that what it is for a dispute to arise "wholly in virtue of a disagreement over the meaning of an expression" is for the dispute not to be partly factual.

³⁰ Jenkins goes on to refine this view in light of context sensitivity, but the refinement does not matter for present purposes.

not yield concrete predictions as to which disputes are merely verbal and which are not.

My approach is superior to Jenkins in that it yields concrete predictions as to which disputes are purely verbal and which are not. Furthermore, there is a sense in which my own approach may complement Jenkins': as I explained in my discussion of the argument from lost disagreement, we can define what it is for a partly factual dispute to be about a given subject matter. Thus, if it turns out that only partly factual disputes can be about a subject matter its participants take the disputed utterance to be about, my account of purely verbal disputes would be equivalent to Jenkins'.

To summarize, I claimed that a dispute is purely verbal just in case it is not partly factual, and argued that this characterization explains the intuition that no knowledge of extra-linguistic facts is at stake in purely verbal disputes. I also argued that the present approach has advantages over standard characterizations of purely verbal disputes, as well as Chalmers' and Jenkins' characterizations. As I argued, the notion of a partly factual dispute need not only play a role in explaining the epistemic significance and resolvability of typical disputes; it can also be used to explain the sense that purely verbal disputes are pointless.

3.8 Conclusion

I argued that, given the truth of Variance, certain features of ordinary disputes call for an explanation. First, if Variance is true, we lack an explanation of the fact that typical disputes exhibit epistemic conflicts; i.e. we lack an explanation of why at most one of the participants in those disputes knows the facts she takes the utterance under dispute to be about. Second, I argued that a straightforward explanation of epistemic conflicts appealing to material conflicts does not account for the fact that the epistemic conflicts exhibited by

typical disputes are grounded in facts that all the parties to the dispute deem relevant to the truth of the utterance the dispute is about.

I introduced the notion of a partly factual dispute and showed that, necessarily, if a dispute is partly factual, then it exhibits an epistemic conflict. I also claimed that if typical disputes are partly factual, that would account for the fact that the epistemic conflicts exhibited by typical disputes are grounded on facts that all the participants to the dispute deem relevant to the truth of the utterance their dispute is about. Then I argued that typical disputes are partly factual. As I showed, this kind of explanation is not only compatible with Variance, but also with the stronger thesis I called Uncertainty.

Towards the end of the chapter I discussed two further uses for the notion of a partly factual dispute. The first of those uses concerns so-called arguments from lost disagreement against contextualism. I argued that contextualists can reply to the argument from lost disagreement by arguing that certain cross-contextual disputes are partly factual. The second use concerns the characterization of purely verbal disputes. As I briefly argued, the definition of purely verbal disputes as disputes that are not partly factual has significant advantages over standard accounts of purely verbal disputes as well as more recent approaches by Chalmers and Jenkins. In this way, the notion of a partly factual dispute not only puts us in a position to explain important features of ordinary disputes; it can also play important roles in adjudicating between contextualists and invariantists on the one hand, and in fleshing out the notion of a purely verbal dispute on the other.

One important element I have left out are disagreement reports. For instance, if Anna and John have a dispute over the truth of Anna's utterance of (5), it seems natural to report Anna and John as disagreeing over whether Anna's house is green. However, given the truth of Variance, it is unclear what the truth of such a report would involve. The next chapter deals with this problem indirectly by providing a semantic account of indirect

speech reports capable of accommodating Variance, but I will leave a specific discussion of disagreement reports for future work.

Chapter 4

Homophonic speech reports

4.1 Introduction

Call reports of the form *S said that ϕ through U*, where ‘S’ is to be replaced with a term denoting a speaker, ‘ ϕ ’ is to be replaced with a declarative sentence in English, and ‘U’ is to be replaced with a term denoting an assertoric utterance, *targeted speech reports*. This chapter is about a seemingly plausible thesis concerning targeted reports and its interaction with Variance:

Strictness: Necessarily, for any sentence ϕ , and denoting terms S and U, \lceil S said that ϕ through U \rceil is true in English as used in context c only if the referent of U in English as used in c is an assertoric utterance whose truth-conditional content is the proposition expressed by ϕ in English relative to c.^{1, 2}

¹ We will see later on, in section 4.7.3, that Variance puts pressure on the very idea that there can be a proposition expressed by a sentence in a context. I will abstract away from those issues for the time being, but later I will discuss what to do if sentences in natural language don’t have unique semantic contents.

² There are other ways in which targeted reports may be false other than if the proposition embedded in the report is not the truth-conditional content of the utterance the report is about. For example, it may be that S didn’t in fact utter U, or that U made no utterance. For present purposes, we can ignore these ways for a targeted report to be false. Thanks to Jim Pryor for discussion.

For example, according to Strictness, it is necessarily the case that ‘Anna said that grass is green through her utterance of ‘grass is green’ is true in English as used in the context I currently occupy only if the truth-conditional content of Anna’s utterance of ‘grass is green’ is the proposition expressed by ‘grass is green’ in English relative to that context.

Together with instances of the disquotational schemas *necessarily, ‘ ϕ ’ is true in English as used in the context I currently occupy iff ϕ , necessarily, ‘ ϕ ’ expresses the proposition that ϕ in English relative to the context I currently occupy*, and *necessarily, ‘T’ as used in English in the context I currently occupy refers to T*, Strictness entails the corresponding instances of the schema:

Strictness-schema: Necessarily, S said that ϕ through U only if U’s truth-conditional content is the proposition that ϕ .³

For example, according to one instance of the Strictness-schema, it is necessarily the case that Anna said that grass is green through her utterance of ‘grass is green’ only if that utterance’s truth-conditional content is the proposition that grass is green. In this way, Strictness and the Strictness-schema relate the ordinary notion of saying-that with the technical notions of truth-conditional content and the proposition expressed by a given sentence.

Taken at face value, Strictness is too strong. For example, if Anna utters ‘grass is green and snow is white’, she said that grass is green through her utterance. However, the proposition that grass is green is not the truth-conditional content of Anna’s utterance; rather, it is only entailed by it. For the sake of simplicity, I will postpone discussion of these issues until section 4.7.2. For present purposes, we can focus on the instances of Strictness and the Strictness-schema in which the sentence embedded in the report’s that-clause is the

³ Substituting a particular sentence ϕ^* , and particular terms S^* and U^* into Strictness, we get a necessary conditional whose left hand side is necessarily equivalent to S^* said that ϕ^* through U^* via the first disquotational schema. Similarly, we can use the second and third disquotational schemas to get that the right hand side of Strictness is necessarily equivalent to U^* ’s truth-conditional content is the proposition that ϕ^* .

exact same sentence the speaker uttered. So, for example, we can focus on reports like ‘Anna said that grass is green through U’, where U is itself an utterance of ‘grass is green’. We may call targeted reports of this kind *homophonic*.

There are multiple ways of motivating the restricted version of Strictness. For instance, one may think that the report ‘S said that ϕ through U’ is true just in case S *said* the proposition expressed by ϕ in English in the context in which the report is made through U, and that part of what it takes for a speaker to *say* a proposition through a given utterance—as opposed to merely implicating it—is for that proposition to be that utterance’s truth-conditional content.⁴ Or one may think that if U is an utterance of ϕ , then U’s truth-conditional content just is the proposition semantically expressed by ϕ in English in the context in which U was made, so if one can be accurately reported by means of ϕ , then U’s truth-conditional content must be the proposition expressed by ϕ in the context in which the report is made. In general, Strictness will be compelling for those who think that, in the strictest and most fundamental sense, what one says through an utterance is that utterance’s truth-conditional content, and homophonic reports are the tool we have in natural language to capture what a speaker said in the strictest and most fundamental sense.

This chapter argues against Strictness. In particular, I will argue that, via the relevant instances of the Strictness-schema, Strictness makes it extremely difficult to know what somebody said through a given utterance. This is a worrisome conclusion, especially when it comes to homophonic reports. For example, according to this conclusion, if Strictness is true, then at most one person knows that Anna said that grass is green through her utterance

⁴ Defenders of this kind of view may add further constraints to the definition of what it is to say a proposition through an utterance. For example, following Grice (1989a), Schiffer (1972), Bach and Harnish (1979) they may add that in order for a speaker to say proposition P through a given utterance, she must (i) intend her audience to believe P, (ii) intend her audience to recognize intention (i), and (iii) intend her audience to fulfill intentions (i) and (ii) on the basis of its being common knowledge between speaker and audience that the truth-conditional content of the speaker’s utterance is P. The addition of such constraints is not relevant to the present discussion.

of ‘grass is green’. And if we don’t know that Anna said that grass is green through her utterance of ‘grass is green’, one starts to wonder whether we know what Anna said at all. As I will argue, we can avoid this worrisome consequence by abandoning Strictness and adopting a more flexible account of targeted reports.

Before we proceed, it is important to clarify the scope of the discussion. As I have stated Strictness, it is only a thesis about reports involving the verb ‘to say’. However, my discussion here applies equally to versions of Strictness involving other reporting verbs. For example, where ‘ ϕ ’ is to be replaced with a declarative sentence in English, and ‘U’ and ‘S’ are to be replaced with referring terms, such reports may be of the forms: ‘S *asserted* that ϕ through U’, ‘S *literally and strictly speaking said* that ϕ through U’, ‘S *expressed the proposition* that ϕ through U’, etc. For each of those kinds of reports, we can generate a suitable version of Strictness, and the discussion in this chapter applies to all of them.

The chapter is structured as follows. I start by explaining the problems Variance raises for Strictness (section 2). Then I offer an account of homophonic reports that rejects Strictness (sections 3–5), and show that it solves the problems arising from Variance (section 6). I discuss further issues towards the end of the chapter, including a generalization of my account to non-homophonic reports (section 7), and conclude (section 8).

4.2 Troubles with Variance

I will argue that if Strictness is true, then at most one language user knows what somebody said through a given utterance. In order to introduce the argument, consider the following example. Anna has just gotten a new toy, which she decided to call ‘Charlie’. When describing Charlie to her friend John, Anna utters

- (1) Charlie is green.

Now consider the report

(2) Anna said that Charlie is green through her utterance of ‘Charlie is green’

As it happens, there is no proposition Anna and John both believe to be the truth-conditional content of Anna’s utterance of (1).

Now, recall the definition of truth-conditional content from chapter 1:

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 relative to c in L if and only if P is true.⁵

It follows from this definition that, for any given utterance, at most one of any two propositions that are not necessarily equivalent is that utterance’s truth-conditional content.

Given this observation and the fact that there is no proposition Anna and John both believe to be the truth-conditional content of Anna’s utterance of (1), it follows that there is no proposition they both know to be the truth-conditional content of Anna’s utterance. And, given Strictness, it would seem to follow (via the corresponding instance of the Strictness schema) that at most one of them knows that Anna said that Charlie is green through her utterance of (1).⁶ Now, recall that, according to Variance, 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. Through a generalization of the reasoning above, it would seem to follow from Variance and Strictness (via instances of the Strictness schema) that, for nearly every utterance U, at most one language user knows that S said that ϕ through U—where S is the speaker who uttered U and ‘ ϕ ’ is to be replaced with a declarative sentence.

⁵ Sentences here should be thought of as disambiguated sentences.

⁶ Note that the argument here relies on an objectionable principle of epistemic closure (i.e. if A knows P, and P entails Q, then A knows Q), but we will see in a moment that we can develop a version of the argument that does not rely on that principle.

I say “it would seem to follow” because the present argument relies on an objectionable principle of epistemic closure: that if A knows P, and P entails Q, then A knows Q. However, we can formulate a similar argument that relies only on an unobjectionable version of closure as follows.

Let’s start by introducing the notion of *weak belief*. As I will understand weak belief here, A weakly believes proposition P just in case P is true in every possible world compatible with every proposition A believes. Given the present definition, someone can weakly believe propositions she doesn’t believe in the ordinary sense of belief; for example, if someone believes P, and Q follows from P, she need not believe Q in the ordinary sense, but she will *weakly* believe that Q.

The following two principles follow from the definition of weak belief:

Belief entails weak belief: If A believes P, then A weakly believes P.

Weak-belief closure: If A weakly believes P, and P entails Q, then A weakly believes Q.

It follows from the relevant instance of the Strictness-schema and weak-belief closure that if there is no proposition Anna and John *weakly believe* to be the truth-conditional content of Anna’s utterance of (1), then at most one of them weakly believes that Anna said that Charlie is green through her utterance of ‘Charlie is green’. And, given that belief entails weak belief, it follows from this that at most one of Anna and John believes (in the ordinary sense) that Anna said that Charlie is green through her utterance of ‘Charlie is green’. Finally, it follows from this and the fact that knowledge entails belief (i.e. if A knows P, then A believes P) that at most one of Anna and John knows that Anna said that Charlie is green through her utterance of ‘Charlie is green’.⁷

⁷ Some people may think that the notion of weak belief does not correspond to any realistic psychological notion. That may or may not be true, but it doesn’t matter for our purposes. For our purposes, what matters is that, given the definition of weak belief, belief entails weak belief, and weak belief is closed under logical consequence. Thanks to Stephen Schiffer for helpful discussion of these issues.

More generally, it follows from Strictness and the three principles used above—i.e. weak-belief closure, the principle that belief entails weak-belief, and the principle that knowledge entails belief—that, for all declarative sentences we replace ‘ ϕ ’ with, and for any speaker S, utterance U, and language users A and B: if there is no proposition A and B both weakly believe to be U’s truth-conditional content, then at most one of A and B knows that S said that ϕ through U. Now, let *Variance** be the thesis that nearly every utterance is such that there is no proposition which more than one language user with consistent beliefs *weakly believes* to be that utterance’s truth-conditional content. Together with the conditional I just stated, *Variance** entails that, for all declarative sentences we replace ‘ ϕ ’ with, for any speaker S and nearly every utterance U, at most one language user with consistent beliefs knows that S said that ϕ through U.

*Variance** is logically independent from Variance, but we can straightforwardly adapt the case for the latter so as to defend the former. The case for *Variance** is as follows. For nearly every utterance and any given language user, there is an enormous number of equally natural or otherwise eligible non-truthconditionally-equivalent propositions any given language user could easily have weakly believed to be that utterance’s truth-conditional content. Given the huge number of such propositions, and that any language user can consistently weakly believe at most one of them to be the truth-conditional content of the utterance in question, it would be extremely unlikely for any two language users with consistent beliefs to weakly believe the exact same one of those propositions to be that utterance’s truth-conditional content. Which is to say that it would be extremely unlikely for there to be a proposition which more than one language user with consistent beliefs weakly believes to be that utterance’s truth-conditional content.

Take for instance Anna’s utterance of ‘Charlie is green’. Presumably, its truth depends on where in the color spectrum Charlie’s color falls. But there is an enormous number of

extremely similar and equally natural regions of the color spectrum such Anna could easily have weakly believed ‘green’ to refer to. Accordingly, there is an enormous number of extremely similar and equally natural propositions Anna could easily have weakly believed to be her utterance’s truth-conditional content. For example, the propositions that Charlie is the color corresponding to region r_1 , that Charlie is the color corresponding to region r_2 , and so on, where r_1, r_2, \dots are extremely similar and equally natural regions of the color spectrum to be denoted by ‘green’. The same is true of John: there are enormously many propositions he could easily have weakly believed to be the truth-conditional content of Anna’s utterance, each corresponding to the proposition that Charlie’s color falls in a slightly different region of the color spectrum. Given the enormous number of such propositions, if Anna and John have consistent beliefs, it would be extremely unlikely for Anna and John to have weakly believed the exact same one to be the truth-conditional content of Anna’s utterance.

Observations of this kind hold for utterances of sentences involving terms whose application depends on the properties an object has along one or more sufficiently fine-grained dimensions. In order to see this, suppose for example that ‘F’ is to be replaced with a predicate of that kind and ‘a’ is to be replaced with a proper name. For any language user, there will be a huge number of extremely similar and equally natural (or otherwise eligible) properties—each corresponding to a slightly different cutoff point along one or more of the dimensions on which the application of the predicate we replace ‘F’ with depends—such that that language user could easily have weakly believed that an utterance of ‘ $\lceil a \text{ is } F \rceil$ ’ has the truth-conditional content that a (the object denoted by the term we replace ‘a’ with) has that property. If, for example, a language user weakly believes that the predicate we replace ‘F’ with expresses the property corresponding to the cutoff point x along one of the dimensions relevant to that predicate’s application, she could just as easily have weakly

believed that that predicate 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. Given the huge number of properties any language user could easily have weakly believed that predicate to express, it is extremely unlikely for any two language users to weakly believe that an utterance of $\lceil a \text{ is } F \rceil$ has the truth-conditional content that a has the exact same one of those properties.

Defenders of Strictness may call attention to the fact that Variance* is restricted to speakers with consistent beliefs. As they may point out, it is very likely that ordinary speaker's beliefs are inconsistent, in which case the truth of Variance* is compatible with Strictness and the claim that ordinary language users know what people say through their utterances. In this way, defenders of Strictness may claim that Variance* is not problematic at all.

I think this attitude misjudges the nature of the problem. While it may well be that ordinary people's beliefs are not usually consistent, what the present considerations show is that, if Strictness is true, ordinary people's beliefs *must* be inconsistent in order for them to even be in a position to know what speakers say through their utterances. And it would be odd, to say the least, if knowing such propositions was only possible if we had inconsistent beliefs. It would be odd if, for example, only somebody with inconsistent beliefs could know that Anna said that Charlie is green through her utterance of 'Charlie is green'. In this way, pointing to the inconsistency of ordinary people's beliefs seems to get little traction against the case for Variance*.

Other ways of resisting the case for Variance* may involve appeals to metaphysical or psychological naturalness, to common ground, or to social externalism. However, the considerations I offered against those strategies in chapter 1 apply here too. Thus, unless we're willing to accept that we don't know seemingly easy to know facts about what people

say through their assertoric utterances—e.g. that Anna said that Charlie is green through her utterance of ‘Charlie is green’—we must reject Strictness. The question is what a correct account of targeted reports should look like given that Strictness is false. The rest of this chapter answers that question.

4.3 Diagnosing the problem

If Strictness is true, then knowing (say) that Anna said that Charlie is green through her utterance of ‘Charlie is green’ requires weakly believing that the truth-conditional content of Anna’s utterance is the proposition that Charlie is green. But we have already seen that it would be unlikely for any two people to consistently weakly believe such a thing. What we need is an account according to which knowledge that (e.g.) Anna said that Charlie is green through her utterance of (1) is compatible with many different beliefs (or weak beliefs) about the truth-conditional content of Anna’s utterance. For example, according to such an account, John can be in a position to know that Anna said that Charlie is green through her utterance of ‘Charlie is green’ regardless of whether he believes that that utterance’s truth-conditional content is the proposition that Charlie’s color falls in region r_1 of the color spectrum, the proposition that Charlie’s color falls in region r_2 , etc.

As an example of a view that satisfies this desideratum, consider:

Similarity: Necessarily, for any sentence ϕ , and denoting terms S and U, ‘S said that ϕ through U’ is true in English as used in context c just in case the referent of U in English as used in c is an assertoric utterance whose truth-conditional content is *similar* to the proposition expressed by ϕ in English relative to c .

In section 4 we will see that this view is unsatisfying in certain respects, but examining it will help us pin down what structural features an account of homophonic reports must have

in order to avoid the problems from Variance and Variance*.

The problem with Strictness was that, according to it, knowing that Anna said that Charlie is green through her utterance of (1) would require weakly believing that the proposition that Charlie is green is the truth-conditional content of Anna's utterance. Given Variance*, it would be very unlikely for any two language users to have that weak-belief. In contrast, if Similarity is true, then knowing that Anna said that Charlie is green through her utterance of (1) only requires weakly believing that the proposition that Charlie is green is similar to the truth-conditional content of Anna's utterance. And since the proposition that Charlie is green may be similar to many different propositions, in principle there will be many different propositions one could weakly believe to be the truth-conditional content of Anna's utterance while still knowing that Anna said that Charlie is green through that utterance.

For example, suppose that $green_1$ and $green_2$ are slightly different overlapping regions of the color space, so that the proposition that Charlie is $green_1$ and the proposition that Charlie is $green_2$ are very similar. Suppose also that the proposition that Charlie is green in fact is the proposition

(3) `Charlie is green1`,⁸

and that John weakly believes that the truth-conditional content of Anna's utterance of (1) is

(4) `Charlie is green2`.

According to the relevant instance of the Strictness-schema, Anna said that Charlie is green through her utterance of (1) only if (3) is the truth-conditional content of Anna's utterance. Thus, if Strictness is true, a weak-belief that Anna said that Charlie is green through her utterance of (1) is incompatible with John's weak-belief that that utterance's truth-conditional

⁸ Sentences written in `this` font are names for propositions.

content is (4).

In contrast, according to the relevant instance of Similarity, Anna said that Charlie is green through her utterance of (1) only if (3) is similar to the truth-conditional content of Anna's utterance. Thus, if Similarity is true, weakly believing that Anna said that Charlie is green through her utterance of (1) is entirely compatible with John's belief that that utterance's truth-conditional content is (4). For, at least in principle, John may weakly believe that (4) and (3) are similar propositions.

Can John also *know* that Anna said that Charlie is green through her utterance of (1)? I think the answer is yes, but it will be easier to see why by examining an example first. Suppose somebody throws a dart onto a dartboard, and the dart lands right in the center of the bullseye. John looks at the dartboard from afar and believes that the dart has landed just a tiny bit away from the exact center of the bullseye, but still well within the bullseye. I take it that, despite John's false belief about the exact landing of the dart, he may still know that the dart landed in the bullseye. More generally, one can know that a dart landed in the bullseye as long as (i) the dart in fact landed in the bullseye, and (ii) the point in the dartboard one believes the dart to have landed in (if any) is well inside the bullseye.

What I want to propose is that, if Similarity is true, knowing that Anna said that Charlie is green through her utterance of (1) is very much like knowing that a dart landed in the bullseye. In the same way in which one can know that the dart landed in the bullseye without knowing (or even having any weak belief about) where exactly it landed, one may know that Anna said that Charlie is green without knowing (or any having any weak belief about) what proposition is the truth-conditional content of her utterance of 'Charlie is green'. In particular, if the relevant instance of Similarity is true, John may know (or be in a position to know) that Anna said that Charlie is green through her utterance of (1) as long as (i) the proposition that Charlie is green is in fact similar to the truth-conditional content

of Anna's utterance, and (ii) the proposition John takes to be the truth-conditional content of Anna's utterance (if any) is similar to the proposition that Charlie is green and to the actual truth-conditional content of Anna's utterance.⁹

Now, while my discussion in this section has focused on the thesis I called "Similarity", there are many related views which would allow us to block the problems from Variance and Variance*. In general, where X is a relation holding between the proposition expressed by ϕ in English in the context in which a report is made and the plausible candidates for the utterance that U refers to, any version of the following thesis will allow us to block those problems:

Looseness: Necessarily, for any sentence ϕ , and denoting terms S and U , 'S said that ϕ through U ' is true in English as used in context c just in case the referent of U in English as used in c is an assertoric utterance whose truth-conditional content *stands in relation X* to the proposition expressed by ϕ in English relative to c .

Given Looseness and the corresponding schema—necessarily, S said that ϕ through U only if U 's truth-conditional content stands in relation X to the proposition that ϕ —for all declarative sentences we replace ' ϕ ' with, if the proposition that ϕ and all the propositions one could easily have believed to be U 's truth-conditional content (including U 's actual truth-conditional content) stand in relation X to one another, ordinary language users will be in a position to know that S said that ϕ through U . The question now is what relation we should take X to be. The next section imposes some constraints on that relation, and the section after that presents an actual implementation of Looseness.

⁹ Note that, if Similarity is true, then knowing that Anna said that Charlie is green through her utterance of (1) does not require having any particular weak belief about what U 's truth-conditional content is. It only requires having the weak belief that that truth-conditional content is similar to the proposition that Charlie is green. Thus, if Similarity is true, knowing that Anna said that Charlie is green through her utterance of (1) is compatible with the thesis I called 'Uncertainty' in chapter 1.

4.4 Constraints on Looseness

In making a speech report, we attempt to characterize what the speaker said. We will see later, in section 4.7.3, that there are reasons to think that ordinary utterances don't have truth-conditional contents, but setting that aside for a moment, we can suppose that part of what a speaker says through an utterance is its truth-conditional content. And if part of what a speaker says through an utterance is that utterance's truth-conditional content, then a report of that utterance should be true just in case it characterizes that truth-conditional content reasonably accurately. The question is what it takes for a report to reasonably accurately characterize such truth-conditional content. We have seen that such characterizations need not be so accurate that the report attributes to the speaker the saying of the exact truth-conditional content of the speaker's utterance, but surely there must be some constraints on what an accurate characterization amounts to. In terms of the thesis I called Looseness, those constraints will come by way of constraints on the relation X that Looseness appeals to.

This section proposes three such constraints. In order to do so, I will find it helpful to focus on the proposal I called Similarity. By examining what is satisfying or unsatisfying about Similarity and the similarity relation it appeals to, we can learn more about what relation X should be.

One problem with similarity is that, presumably, homophonic reports are truth-preserving, but similarity between propositions is not. For example, if Anna said that Charlie is green through her utterance of 'Charlie is green', we should expect that if Anna's utterance of 'Charlie is green' is true, then so should be the proposition that Charlie is green. If two propositions are as similar in one possible world as they are in another, mere similarity between the proposition that grass is green and the truth-conditional content of Anna's

utterance does not guarantee this.

For example, suppose that the proposition that Charlie is green is in fact (3) and the truth-conditional content of Anna's utterance is (4). Suppose also that (4) is true and (3) is false. Given our assumption that (3) and (4) are very similar, Similarity predicts that Anna said that Charlie is green through her utterance of 'Charlie is green' in the present scenario, even though Anna's utterance is true and the report attributes to Anna the saying of something false.

A related problem arises if what the report attributes to the speaker is true, even though the speaker's utterance is false. If Anna said that Charlie is green through her utterance of 'Charlie is green', then if Anna's utterance is false, so should be what Anna said according to the report. On the basis of these considerations, we can impose the following constraint on the relation X we are after: if two propositions stand in relation X , then they are materially equivalent.

Another problem concerns the exact way in which we should understand similarity. For example, in a certain sense of 'similar', these two propositions are fairly similar:

(5) Charlie is green or exactly 1.6m tall

(6) Charlie is green or weighs exactly 1.6kg

(5) and (6) are similar in that they (presumably) have large areas of overlap in logical space, in that they are both disjunctions, in that one of their disjuncts attributes to Charlie a very precise quantity, etc. Despite these similarities, it would be odd for it to be true that Anna said that Charlie is green through her utterance of 'Charlie is green' in a scenario in which the proposition that Charlie is green just is proposition (5) and the truth-conditional content of Anna's utterance is (6). Thus, relation X can't be similarity in just about any respect.

What is the relevant difference between (5) and (6) such that they don't stand in relation

X despite their various similarities? I want to propose that an important difference concerns their subject matter. (5) tells us something about Charlie’s color *and* height, namely, that Charlie is not both a color other than green and a height other than 1.6m tall. On the other hand, (6) tells us something about Charlie’s color *and* weight: that it’s not both the case that Charlie is a color other than green and that Charlie has some weight other than 1.6kg. In this way, though (5) and (6) both tell us something that is partly about Charlie’s color, they are not entirely about the same subject matter: one is partly about Charlie’s weight, and the other is partly about Charlie’s height. In light of this observation, I want to propose the following constraint on relation X : two propositions stand in relation X only if they have “matching” subject matters. In the next section I will present an implementation of this idea, but for now let’s move on to the third constraint.

In the last few paragraphs I have been pointing out problems for Similarity, but Similarity seems to be right in that relation X must presumably require significant similarity between its relata. Otherwise, it could be true that Anna said that Charlie is green through her utterance of ‘grass is green’ in a scenario in which the proposition that Charlie is green is in fact the proposition (7), the truth-conditional content of Anna’s utterance is (8), and Charlie is color 4 (see fig. 4.1).

(7) Charlie is one of colors 1-4

(8) Charlie is one of colors 4-7



Figure 4.1: Some colors

I take it that if these conditions obtain, it will not be true that Anna said that Charlie is green through her utterance of ‘Charlie is green’. Though (7) and (8) are both true and have

matching subject matters, there is a sense in which the requirements (7)'s truth imposes on the world are still too different from the requirements (8)'s truth imposes. This leads to the third and last constraint I want to impose on X : two propositions stand in relation X only if they are truth-conditionally similar. As with the relation of having matching subject matters, I will flesh out the notion of truth-conditional similarity in the next section. For the time being, it is worth summarizing the discussion so far.

In the previous section I offered an argument to the effect that, if Strictness is true, we don't know what people say through their utterances. Given that argument, I have claimed we ought to reject Strictness and replace it with some version of Looseness. As I have shown, if Looseness is true, knowledge of what the speaker said through a given utterance is compatible with several different beliefs about the truth-conditional content of the utterance the report is about. Finally, I have imposed three constraints on the relation X that occurs in Looseness: first, if two propositions stand in relation X , they are materially equivalent; second, if two propositions stand in relation X , they have the same subject matter; third, if two propositions stand in relation X , they are truth-conditionally similar.

Before proceeding to my preferred implementation of the constraints I just introduced, it is worth noting an interesting consequence of adopting Looseness. Throughout this chapter I have focused on reports involving the reporting verb 'to say', but as I said in the introduction, there are other reporting verbs which may also generate candidates for strictness. Of particular interest here is the reporting verb 'to literally say the proposition that'. Because of the arguments from the previous section, if the corresponding version of Strictness is true of reports of the form *S literally said the proposition that ϕ through U* , then we rarely know that they are true. However, taking the corresponding version of Looseness to be true of such reports seems much more theoretically costly than taking it to be true of ordinary natural language reports. For it is a common assumption in semantics that the proposi-

tion a speaker literally says through a given utterance and that utterance’s truth-conditional content are one and the same. Thus, adopting Looseness for reports about the proposition a speaker literally says will require us to reject a dear theoretical commitment. In other words, we face a dilemma: on the one hand, if reports about what a speaker literally says are strict, we almost never know what proposition a speaker literally says through a given utterance; on the other, if Looseness is true of those reports, we must reject the equation between semantic content and truth-conditional content. I will leave the discussion of this consequence for future work; the rest of this chapter focuses on my preferred implementation of Looseness.

4.5 An implementation of Looseness

The constraints on relation X I discussed in the previous section are: that if P and Q stand in relation X , they are materially equivalent; that if P and Q stand in relation X , they have matching subject matters; and that if P and Q stand in relation X , they are truth-conditionally similar. The first of these constraints requires no explanation. The rest of this section develops particular implementations of the second and third constraints, and uses them to define the relation I call “descriptive similarity”. Descriptive similarity is the relation I will appeal to in my preferred implementation of Looseness.

4.5.1 Matching subject matters

In the previous section I said that the relation X that Looseness appeals to shouldn’t be such that propositions (5) (Charlie is green or exactly 1.6m tall) and (6) (Charlie is green or weighs exactly 1.6kg) stand in it. This is so because (5) gives us information partly about Charlie’s height, whereas (6) gives us no such information; on the

other hand, (6) gives us information partly about Charlie's weight, whereas (5) gives us no such information. In this way, (5) gives us information that goes beyond (6)'s subject matter, and vice versa. Based on this example, I said that relation X should be such that if two propositions stand in it, they should match in subject matter.

We can refine what it is for two propositions to match in subject matter using the framework introduced in chapter 2. There I said that we can represent subject matters as sets of possible states of affairs, and that a proposition's exact subject matter is the set of the possible states that could make it true were they to obtain (i.e. its possible truthmakers) and the possible states that could make it false were they to obtain (i.e. its possible falsitymakers). Using this notion of subject matter as the basis, we can define what it is for two propositions to have matching subject matters as the following equivalence relation:

Subject-matter matching: P and Q have matching subject matters just in case:

- (i) Every part of a state of affairs in P's exact subject matter has a part in common (i.e. overlaps) with a state of affairs in Q's exact subject matter; and
- (ii) Every part of a state of affairs in Q's exact subject matter has a part in common with a state of affairs in P's exact subject matter.¹⁰

I will explain how the resulting notion applies to the present case in a moment, but before that it is worth explaining the use of subject-matter matching instead of equivalence in exact subject matters.

¹⁰ That subject-matter matching is reflexive follows from the reflexivity of the parthood relation. That it is symmetric follows straightforwardly from the definition. That it is transitive can be seen as follows. Symbolize the subject-matter matching relation as ' \equiv '. Suppose M, N, O are subject matters, and that $M \equiv N$, and $N \equiv O$. Now, take an arbitrary part s of some $m \in M$. By assumption, s has a part in common with some $n \in N$. Call that part ' $s \bullet n$ '. $s \bullet n$ is a part of n , and since $N \equiv O$, it follows that there is some $o \in O$ with which $s \bullet n$ has a part in common. Since that common part will be part of $s \bullet n$, and the latter is part of s , it follows that s has a part in common with o . Since s and m were chosen arbitrarily, it follows that every part of a state in M has a part in common with a state in O . An analogous argument shows that every part of a state in O has a part in common with some state in M .

The reason for adopting subject-matter matching has to do with our definition of the possible truthmakers and falsitymakers for Boolean compounds. For example, in chapter 2 I said that the possible truthmakers for a disjunction are: the possible truthmakers for its disjuncts, and the fusions of compatible truthmakers for its disjuncts. On the other hand, its possible falsitymakers are the fusions of compatible falsitymakers for the disjunctions' disjuncts.¹¹

With this in mind, take the disjunctions $P \vee Q$ and $P \vee \neg Q$. For the sake of simplicity, suppose that the only possible truthmaker for P is p^T , the only possible truthmaker for Q is q^T , the only possible falsitymaker for P is p^F , and the only possible falsitymaker for Q is q^F . Then the possible truthmakers for $P \vee Q$ are p^T, q^T , and $p^T + q^T$, and its only possible falsitymaker is $p^F + q^F$. In turn, the possible truthmakers for $P \vee \neg Q$ are p^T, q^F , and $p^T + q^F$, and its only possible falsitymaker is $p^F + q^T$. Now, since a proposition's exact subject matter is the set of its possible truthmakers and falsitymakers, $P \vee Q$ and $P \vee \neg Q$ do not have the same exact subject matters. In particular, the state q^T is in the first of those two proposition's subject matter, but not of the second, and the same is true of other states. In this way, identity of subject matters is too constraining to capture the relation between subject matters that we are after. On the other hand, even though the two propositions do not have the same exact subject matters, they *do* have matching subject matters. This is so because all of $p^T, q^T, p^T + q^T$, and $p^F + q^F$ have a part in common with one of $p^T, q^F, p^T + q^F$, and $p^F + q^T$, and vice versa.

To bring the discussion back to the examples from the previous section, suppose for simplicity that the proposition `Charlie is green` has only one possible truthmaker, g^T , and one possible falsitymaker: g^F . Suppose also that the proposition `Charlie is`

¹¹ It is worth calling attention to the fact that I talk about fusions of *compatible* states, whereas Fine (MSc) does no such thing. One reason to follow Fine in this respect is that, without the compatibility requirement, certain results are more straightforward. See below, footnote 13, for an example.

exactly 1.6m tall has one possible truthmaker, t^T and one possible falsitymaker t^F . Finally, suppose that the proposition Charlie weighs exactly 1.6kg has one possible truthmaker w^T , and one possible falsitymaker, w^F .

Given these assumptions, the possible truthmakers for Charlie is green or exactly 1.6m tall are: g^T, t^T , and $g^T + t^T$; and its only possible falsitymaker is $g^F + t^F$. Thus, that proposition's subject matter corresponds to the set

$$(9) \{g^T, t^T, g^T + t^T, g^F + t^F\}.$$

On the other hand, suppose that the possible truthmakers for Charlie is green or weighs exactly 1.6kg are: g^T, w^T , and $g^T + w^T$; and its only possible falsitymaker is $g^F + w^F$. Thus, this proposition's exact subject matter corresponds to the set

$$(10) \{g^T, w^T, g^T + w^T, g^F + w^F\}$$

Since the state t^T does not overlap with any state in (10), (5) and (6)'s subject matters don't match. Accordingly, if two propositions stand in relation X only if they have matching subject matters, (5) and (6) do not stand in relation X. Thus, Looseness will predict that in a scenario in which the proposition that Charlie is green just is the proposition (5) and the truth-conditional content of Anna's utterance of (1) is (6), it is not true that Anna said that Charlie is green through her utterance of (1).

More generally, let's say that P and Q have non-overlapping subject matters just in case no state in P's subject matter has a part in common with a state in Q's subject matter. If P, Q, and R have non-overlapping subject matters (i.e. P's subject matter doesn't overlap with Q's or R's, and Q's subject matter doesn't overlap with R's), then $P \vee Q$ and $P \vee R$ will fail to have matching subject matters, and the same is true of $P \wedge Q$ and $P \wedge R$.¹²

In contrast, compare (5) with

¹²To see this, suppose P, Q, and R have non-overlapping subject matters. Given this assumption and the definition of the truthmakers for a disjunction (see ch. 2), any truthmaker for Q will be a truthmaker for $P \vee Q$

(11) Charlie is either green or not 1.6m tall

Given our assumptions above about the truthmakers and falsitymakers for *Charlie is green* and *Charlie is 1.6m tall*, (11)'s possible truthmakers will be the states g^T, t^F , and $g^T + t^F$, and its only possible falsitymaker will be the state $g^F + t^T$. Thus, (11)'s exact subject matter will be the set

(12) $\{g^T, t^F, g^T + t^F, g^F + t^T\}$

This is not the same subject matter as (5)'s exact subject matter. However, (5) and (11) have matching subject matters: every state in (5)'s exact subject matter has a part in common with a state in (11)'s subject matter, and vice versa. More generally, if P and R have matching subject matters, Q and S have matching subject matters, every state in P's exact subject matter is compatible with every state in Q's exact subject matter, and every state in R's exact subject matter is compatible with every state in S's exact subject matter, then $P \vee Q$ and $R \vee S$ will have matching subject matters as well.¹³

that does not overlap with any truthmaker or falsitymaker for $P \vee R$. Similarly, given the present assumption and the definition of the falsitymakers for a conjunction, any falsitymaker for Q will be a possible falsitymaker for $P \wedge Q$ that does not overlap with any truthmaker or falsitymaker for $P \wedge R$. Since a proposition's exact subject matter is just the set of its possible truthmakers and possible falsitymakers, it follows that $P \vee Q$ and $P \vee R$ don't have matching subject matters, nor do $P \wedge Q$ and $P \wedge R$. In fact, for that matter, neither $P \vee Q$ and $P \wedge R$ nor $P \vee R$ and $P \wedge Q$ have matching subject matters.

¹³ To see this, suppose that P and R have matching subject matters, Q and S have matching subject matters, every state in P's exact subject matter is compatible with every state in Q's exact subject matter, and every state in R's exact subject matter is compatible with every state in S's exact subject matter. Now take an arbitrary state s in $P \vee Q$'s exact subject matter. We know by the definition of exact subject matter and the definition of truth and falsitymaking for disjunction that s is one of (a) a possible truthmaker for P, (b) a possible truthmaker for Q, (c) a fusion of possible truthmakers for P and Q, and (d) a fusion of possible falsitymakers for P and Q. We now consider each of these cases.

- (a) s is a possible truthmaker for P. Since P and R have matching subject matters, every part of s has a part in common with some state s' in R's exact subject matter. Now, s' is either a possible truthmaker or a possible falsitymaker for R (since a proposition's exact subject matter is just the set of all its possible truthmakers and falsitymakers). If s' is a possible truthmaker for R, then it is in $R \vee S$'s exact subject matter, in which case s has a part in common with a state in $R \vee S$'s exact subject matter. If s' is a possible falsitymaker for R, then, since every state in R's exact subject matter is compatible with every state in S's subject matter, s' is compatible with a falsitymaker for S, in which case the fusion of s' and that falsitymaker for S will be a falsitymaker for $R \vee S$. If that is so, then every part of s has a

As I pointed out in the previous section, subject matter matching is not the only relevant dimension of similarity. To see why, consider the propositions (7)—`Charlie is one of colors 1-4`—and (8)—`Charlie is one of colors 4-7`. In chapter 2 I said that if the truth of a proposition of the form Fx depends on x 's properties along various

part in common with a state in $R \vee S$'s exact subject matter, namely, the falsitymaker(s) for $R \vee S$ that s' is part of.

- (b) s is a possible truthmaker for Q . The proof is analogous to the one above. Since Q and S have matching subject matters, every part of s has a part in common with some state s' in S 's exact subject matter. Now, s' is either a possible truthmaker or a possible falsitymaker for S (since a proposition's exact subject matter is just the set of all its possible truthmakers and falsitymakers). If s' is a possible truthmaker for S , then it is in $R \vee S$'s exact subject matter, in which case s has a part in common with a state in $R \vee S$'s exact subject matter. If s' is a possible falsitymaker for S , then, since every state in S 's exact subject matter is compatible with every state in R 's subject matter, s' is compatible with a falsitymaker for R , in which case the fusion of s' and that falsitymaker for S will be a falsitymaker for $R \vee S$. If that is so, then every part of s has a part in common with a state in $R \vee S$'s exact subject matter, namely, the falsitymaker(s) for $R \vee S$ that s' is part of.
- (c) s is the fusion of a truthmaker for P and a truthmaker for Q . Let's start by recalling our definition of fusion from chapter 2: a state x is the fusion of states x' and x'' just in case x' and x'' are both parts of x , and every part of x has a part in common with x' or x'' . It follows from this definition that every part of s has a part in common with a truthmaker for P or with a truthmaker for Q . Since we have seen in cases (a) and (b) that every part of a truthmaker for P and every part of a truthmaker for Q has a part in common with a state in $R \vee S$'s subject matter, it follows that every part of s has a part in common with a state in $R \vee S$'s subject matter.
- (d) s is the fusion of a falsitymaker for P and a falsitymaker for Q . It follows from the definition of fusion that every part s' of s has a part in common with a falsitymaker for P , call it s_P , or a falsitymaker for Q , call it s_Q . Let's start with the first disjunct. By assumption, every part of s_P has a part in common with a state in R 's exact subject matter. The latter is either a possible truthmaker or a possible falsitymaker for R and, as we saw in case (a), either way it is in $R \vee S$'s exact subject matter. Now let's consider the second disjunct. By assumption, every part of s_Q has a part in common with a state in S 's exact subject matter. The latter is either a truthmaker or a falsitymaker for S and, as we saw in case (b), either way it is in $R \vee S$'s exact subject matter. Thus, s' has a part in common with a state in $R \vee S$'s subject matter.

Since all the states we discussed, and s itself, were chosen arbitrarily, it follows that every state in $P \vee Q$'s exact subject matter has a part in common with a state in $R \vee S$'s subject matter. Analogous arguments show that the converse is also true. Thus, $P \vee Q$ and $P \vee R$ have matching subject matters. Further application of the same kind of reasoning shows that if P and R have matching subject matters, Q and S have matching subject matters, every state in P 's exact subject matter is compatible with every state in Q 's exact subject matter, and every state in R 's exact subject matter is compatible with every state in S 's exact subject matter, then $P \wedge Q$ and $R \wedge S$ also have matching subject matters.

Now, I am assuming that only compatible states can be fused, but dropping this assumption would in fact simplify the present discussion. For that would allow us to show, without further caveat, that if P and R have matching subject matters, and Q and S have matching subject matters, then $P \wedge Q$ and $R \wedge S$ also have matching subject matters, as do $P \vee Q$ and $R \vee S$. This, I think, is a good reason to abandon the compatibility requirement. Thanks to Cian Dorr for discussion.

continuous dimensions, then Fx 's possible truthmakers and falsitymakers will be the minimal states of affairs that specify all of x 's maximally determinate properties along all of those dimensions.¹⁴ Accordingly, since a proposition's exact subject matter is just the set of all its possible truthmakers and falsitymakers, Fx 's exact subject matter will be the set of all the possible states of affairs that specify all of x 's maximally determinate properties along such dimensions.

Thus, assuming that the truth of (7) and (8) depends exclusively on Charlie's color, (7) and (8) will have the same exact subject matters. For instance, assuming that colors 1–7 are all the maximally determinate colors there are, the minimal states of affairs in which Charlie is color 1, color 2, and so on, will all be possible truthmakers or possible falsitymakers for (7) and (8) (though, of course, some of those states will be possible truthmakers for one and possible falsitymakers for the other, and vice versa). Hence, for all we have said about the implementation of the constraints on relation X from the previous section, (7) and (8) may stand in relation X . If that is the case, the present implementation of Looseness does not yet exclude that Anna said that grass is green through her utterance of 'grass is green' in a scenario in which (7) and (8) are materially equivalent, the proposition that Charlie is green just is the proposition (7), and (8) is the truth-conditional content of Anna's utterance of (1). The implementation of the third constraint addresses this problem.

4.5.2 Truth-conditional similarity

In a truthmaker framework, we could think that two propositions are truth-conditionally equivalent just in case they have the same possible truthmakers. As a starter, one may think that two propositions should be truth-conditionally *similar* just in case there are very few

¹⁴ Recall, s is a minimal state in which P is true just in case P is true in x , and there is no part of s in which P is true.

possible truthmakers for one that are not possible truthmakers for the other, and vice versa.

I think this way of understanding truth-conditional similarity is roughly on the right track, but it needs refinement. The first problem is that, for many pairs of logically independent but intuitively similar propositions, there will be infinitely many possible truthmakers for one of the propositions in the pair that are not possible truthmakers for the other. For example, take the proposition that Anna is at least 1.6m tall, and the proposition that Anna is at least 1.61m tall. Presumably, these two propositions should count as truth-conditionally similar, yet there are continuum many possible truthmakers for the former that are not truthmakers for the latter; namely, the possible states in which Anna is at least 1.6m tall but below 1.61m tall. So what, then, does it mean that there are very few possible truthmakers for one of those propositions that are not possible truthmakers for the other?

The second problem is that two propositions may count as truth-conditionally similar in certain contexts but not in others. For example, if I want to make a dining table that will accommodate a certain number of people out of a certain plank of wood x , we may think that the propositions x is between 2 and 3m long and x is between 2.05 and 3.05m long are similar enough. But we may not think the same in the context of a delicate experiment that requires using a plank at most 3m long and not a millimeter more. Thus, if we are to capture truth-conditional similarity in terms of the number of possible truthmakers two propositions don't have in common, we should think that what counts as very few possible truthmakers in some contexts may count as too many in others.

In order to solve the first of these problems, we need a measure of the sets of states of affairs that are possible truthmakers for various propositions. In order to address the second problem, we can assume that there are contextually determined thresholds for what counts as too many states of affairs that are not possible truthmakers for the two propositions we are comparing. Given such a measure and threshold, we could define truth-conditional

similarity as follows (where P and Q are propositions, and μ is the measure):

Truth-conditional similarity: P and Q are truth-conditionally similar relative to threshold τ just in case: $\mu(\{s : s \text{ is a possible truthmaker for P but not for Q, or } s \text{ is a possible truthmaker for Q but not for P}\}) \leq \tau$

That is, two propositions are truth-conditionally similar relative to threshold τ just in case the measure of the set of states that are possible truthmakers for only one of them is smaller than or equal to τ .

Now, there are many different ways of measuring the sets of possible truthmakers for the propositions we are concerned with—for example, there are many different probability measures all of which differ from each other only slightly in the values they assign to a given proposition and its negation. Since the same two propositions may be truth-conditionally similar relative to some of them but not relative to others, the question is how to choose a measure among all the possible alternatives which, when input to our definition of truth-conditional similarity, yields intuitively compelling results. In what follows I will assume that there is no unique such measure; rather, different measures will yield the right results when comparing certain pairs of propositions but not when comparing others. What I will do is outline a general way of defining reasonable measures for pairs of propositions whose properties depend on the same object's properties along the same dimensions. In this way, perhaps it is better to think about the present view more as a proof of concept than as a definitive implementation of the kind of measures required to flesh out the notion of truth-conditional similarity.

To begin with, recall that the truth of proposition Fx depends exactly on x 's properties along dimensions d_1, d_2, \dots, d_n —e.g. size, weight, degree of similarity to paradigmatic members of a certain class, etc.—just in case there is a set of points in the space generated by those dimensions such that, necessarily, Fx is true if and only if x 's properties fall in one

of those points (see chapter 2). For example, *Charlie is between 1 and 2m tall* depends exactly on Charlie's height since, necessarily, that proposition is true if and only if Charlie is some height between 1 and 2m.¹⁵

As I said in chapter 2, if the truth of a proposition depends exactly on x 's properties along d_1, d_2, \dots, d_n , that proposition's possible truthmakers will be the minimal states determining x 's maximally determinate properties along d_1, d_2, \dots, d_n in which it is true. Thus, if the truth of a proposition depends exactly on x 's properties along d_1, d_2, \dots, d_n , we can think of that proposition as a region (i.e. a set of points) in the space generated by those dimensions (i.e. $d_1 \times d_2 \times \dots \times d_n$). My proposal is to assign the set of possible truthmakers for a given proposition the measure of the region that proposition occupies in the space generated by d_1, d_2, \dots, d_n . For example, we can think of the proposition *Charlie is between 1 and 2m tall* as the interval $[1, 2]$ in the real line, where each point r corresponds to the minimal state of affairs in which Charlie's height is r meters tall. In turn, if we take the set of minimal possible states that determinate Charlie's exact height, and assume that every set of such states has the measure of the corresponding set of points in the space $d_1 \times d_2 \times \dots \times d_n$, we should think that the set of *Charlie is between 1 and 2m tall*'s possible truthmakers has measure $2 - 1$, i.e. 1.

For the sake of simplicity, my discussion will be focused on comparing simple propositions I will call "rectangular". However, that discussion extends to all propositions that can be thought of as unions of countably many rectangles in a given space. In general, where P depends on o 's properties along dimensions d_1, d_2, \dots, d_n , we can think of the measure of

¹⁵ Strictly speaking, we should also add the condition that any other set of dimensions of which the condition in the main text is true is a superset of d_1, d_2, \dots, d_n . For example, we would like to say that *Charlie is between 1 and 2m tall* depends exactly on Charlie's height since, necessarily, that proposition is true if and only if Charlie is some height between 1 and 2m. But it is also the case that, necessarily, *Charlie is between 1 and 2m tall* is true if and only if Charlie is some height between 1 and 2m *and* some weight between 0 and infinitely many grams. Yet we wouldn't want to say that that proposition's truth depends on Charlie's weight. By adding the condition in this footnote, we exclude that possibility.

the set of possible truthmakers for P as the Lebesgue measure of the set of points such that P is true iff o 's properties falls in one of those points. With this in mind, let's move on.

Among the propositions whose truth depends exactly on x 's properties along d_1, d_2, \dots, d_n , there are some especially simple ones which we can think of as n -dimensional "rectangles" in the space $d_1 \times d_2 \times \dots \times d_n$. Given a set of possible states each of which corresponds to a point in the space $d_1 \times d_2 \times \dots \times d_n$, we can think of the proposition Fx as an n -dimensional rectangle in $d_1 \times d_2 \times \dots \times d_n$ just in case the set of its possible truthmakers corresponds to some set of n -tuples $I_1 \times I_2 \times \dots \times I_n$, where each I_i is a closed interval on d_i . I say that we can think of that proposition as a rectangle because each n -tuple in $I_1 \times I_2 \times \dots \times I_n$ is a point in the space $d_1 \times d_2 \times \dots \times d_n$, and those points form a rectangle whose sides are parallel to the axes d_1, d_2, \dots, d_n .¹⁶ The proposition *Charlie is between 1 and 2m tall* illustrates this: the set of its possible truthmakers corresponds to the closed interval $[1, 2]$ in a scale of heights measured by meters, and we may think of that interval as a 1-dimensional rectangle.

If we think of these simple propositions as n -dimensional rectangles, we can define their measure straightforwardly as follows. Suppose we can think of Fx as the n -dimensional rectangle $I_1 \times I_2 \times \dots \times I_n$. This rectangle's volume is $\ell(I_1) \cdot \ell(I_2) \cdot \dots \cdot \ell(I_n)$ —where $\ell(I)$ is I 's length, given by the absolute value of the difference of its endpoints—and this volume is the measure of the set of Fx 's possible truthmakers. That is, if we take the set of minimal states each of which corresponds to one of the points in the space generated by dimensions d_1, d_2, \dots, d_n , we can assign to each subset of such states the measure of the corresponding region of the space generated by d_1, d_2, \dots, d_n . So, for example, if the volume of the n -rectangle corresponding to Fx in the space generated by d_1, d_2, \dots, d_n is $v \in \mathbb{R}$, so will be the measure of the set of Fx 's possible truthmakers. Take for example *Charlie is*

¹⁶ To clarify, $I_1 \times I_2 \times \dots \times I_n$ denotes a set of points in the space generated by dimensions d_1, d_2, \dots, d_n .

green. If this proposition corresponds to the 1-dimensional rectangle $[2, 3]$ in a scale of heights measured by meters, then the measure of the set of its possible truthmakers will be 1 (i.e. $3 - 2$).

We can compare simple “rectangular” propositions whose truth depends on the same object’s properties along the same dimensions as follows. Let \mathbf{P} and \mathbf{Q} be “rectangular” propositions whose truth depends exactly on x ’s objects along dimensions d_1, d_2, d_n , \mathbf{P} be the set of \mathbf{P} ’s possible truthmakers, and \mathbf{Q} be the set of \mathbf{Q} ’s possible truthmakers. Then, where $I_{P_1} \times I_{P_2} \times \dots \times I_{P_n}$ is the n -dimensional rectangle corresponding to \mathbf{P} in the space generated by d_1, d_2, \dots, d_n , and $I_{Q_1} \times I_{Q_2} \times \dots \times I_{Q_n}$ is the n -dimensional rectangle corresponding to \mathbf{Q} in the same space:

$$(13) \mu(\mathbf{P}) = \ell(I_{P_1}) \cdot \ell(I_{P_2}) \cdot \dots \cdot \ell(I_{P_n})$$

$$(14) \mu(\mathbf{Q}) = \ell(I_{Q_1}) \cdot \ell(I_{Q_2}) \cdot \dots \cdot \ell(I_{Q_n})$$

$$(15) \mu(\mathbf{P} \cap \mathbf{Q}) = \ell(I_{P_1} \cap I_{Q_1}) \cdot \ell(I_{P_2} \cap I_{Q_2}) \cdot \dots \cdot \ell(I_{P_n} \cap I_{Q_n})$$

$$(16) \mu(\mathbf{P} \cup \mathbf{Q}) = \mu(\mathbf{P}) + \mu(\mathbf{Q}) - \mu(\mathbf{P} \cap \mathbf{Q})$$

$$(17) \mu((\mathbf{P} \cap \bar{\mathbf{Q}}) \cup (\bar{\mathbf{P}} \cap \mathbf{Q})) = \mu(\mathbf{P}) + \mu(\mathbf{Q}) - (2 \cdot \mu(\mathbf{P} \cap \mathbf{Q}))$$

Here, $\ell(I_{P_1} \cap I_{Q_1}) \cdot \ell(I_{P_2} \cap I_{Q_2}) \cdot \dots \cdot \ell(I_{P_n} \cap I_{Q_n})$ is the volume of the rectangle $I_{P_1} \cap I_{Q_1} \times I_{P_2} \cap I_{Q_2} \times \dots \times I_{P_n} \cap I_{Q_n}$, which corresponds to the region (itself a rectangle) in which the rectangles associated with \mathbf{P} and \mathbf{Q} overlap.¹⁷ (15) tells us that the measure of the set of possible truthmakers for both \mathbf{P} and \mathbf{Q} is just the measure of that rectangle. (16) tells us that the measure of the possible truthmakers for either \mathbf{P} or \mathbf{Q} is the sum of the measures in which each is true, minus the measure of the set of possible truthmakers for

¹⁷ Here, for each $i \leq n$, $I_{P_i} \cap I_{Q_i}$ corresponds to that rectangle’s side along dimension i . A key assumption in the present discussion is that the sides of the rectangles we are dealing with are parallel to the axes.

both. Finally, (17) tells us that the measure of the set of possible truthmakers for only one of \mathbf{P} and \mathbf{Q} is just the measure of the set of possible truthmakers for either \mathbf{P} or \mathbf{Q} — $\mu(\mathbf{P}) + \mu(\mathbf{Q}) - \mu(\mathbf{P} \cap \mathbf{Q})$ —minus the measure of the set of possible truthmakers for both—i.e. $\mu(\mathbf{P} \cap \mathbf{Q})$.

Given this way of measuring sets of possible states of affairs, we can think of similarity thresholds as follows. If we think that, for each property along certain dimensions d_0, d_1, \dots, d_n , we have a certain interval of tolerance relative to a certain context, we can think of the similarity threshold in that context as the product of the lengths of each of those intervals.

A couple of examples will help us clarify these ideas. Consider again (7) and (8)—respectively, the propositions `Charlie is one of colors 1-4` and `Charlie is one of colors 4-7`. For the sake of simplicity, above I assumed that colors $1, 2, \dots, 7$ were all the maximally determinate colors there are. But we can now simplify a little bit less and assume that all the colors there are are those in the color spectrum between colors 1 and 7, and call each of those colors by the real number corresponding to their position in the line segment between colors 1 and 7. Given these assumptions, we can think of (7) as the 1-dimensional rectangle (i.e. the line segment) $[1, 4]$ and of (8) as the 1-dimensional rectangle $[4, 7]$. And if we take the set of minimal possible states that determine Charlie's exact color between 1 and 7, and assign to each subset of those states the measure of the corresponding line segment in the line between colors 1 and 7, the measure of the set of possible truthmakers for (7) will be just the length of the interval $[1, 4]$; i.e. 3. The same goes, *mutatis mutandis*, for (8).

More relevant to our purposes is the measure of the set of states that are truthmakers for only one of (7) and (8). Given their definitions, (7) and (8) are both made true only by the minimal state of affairs in which Charlie is color 4. Since the set of worlds in which Charlie

is color 4 corresponds to a single point in the line between colors 1 and 7, and the distance from any point to itself is 0, the measure of the set of possible truthmakers for both (7) and (8) is 0 as well. With this in mind, we can calculate the measure of the set of states that make only one of (7) and (8) true. First, the measure of the set of states that are possible truthmakers for (7) but not for (8) is 3, and so is the measure of the set of states that are possible truthmakers for (8) but not for (7). Thus, the measure of the possible truthmakers for only one of (7) and (8) is 6 (since $3 + 3 - (2 \cdot 0) = 6$).

Now, whether (7) and (8) count as similar will be determined by the similarity threshold we use. For example, suppose we care about Charlie's color give or take a distance of 1 in the color spectrum. Then we will be willing to adjust the similarity threshold so that if the set of possible truthmakers for only one of (7) and (8) has at most measure 2, those two propositions will be similar for present purposes. Since that measure is in fact greater than 2, using that threshold will predict that the two propositions are not similar. In contrast, if we only care about Charlie's color give or take a distance of 3 in the color spectrum, we will adjust the similarity threshold differently. In particular, we will adjust it so that if the set of possible truthmakers for only one of (7) and (8) has at most measure 6, and in that case (7) and (8) will count as similar.

Here is another example. Generally, quesadillas are folded tortillas filled with cheese, but (let's suppose) not just any amount of cheese in a folded tortilla or just any amount of corn that the tortilla is made of suffice for a folded tortilla with cheese to count as a quesadilla. Now, consider the following two propositions (where a is an object):

- (18) a is a folded tortilla with between 20 and 40g of cheese,
and
with between 30 and 60g of corn.

(19) a is a folded tortilla with between 19 and 39g of cheese,
and
with between 25 and 55g of corn.

Assuming for simplicity that the only properties of a on which the truth of (18) and (19) depends are a 's quantity of cheese and corn, we could assign a measure to the set of those propositions respective possible truthmakers as follows.

Take all the minimal states which correspond to a point in the space generated by the dimensions *a's exact quantity of cheese* and *a's exact quantity of corn*. For example, one such state will be the one in which a has exactly 20g of cheese and 40g of corn and nothing else happens; another will be the possible state in which a has exactly 20.1355g of cheese and 45.6g of corn and nothing else happens; etc. Given that set of states, we can assign each of its subsets the measure of the corresponding region in the space generated by the dimensions *a's exact quantity of cheese* and *a's exact quantity of corn*. In particular, we can think of the measure of the set of possible truthmakers for (18) and of the set of possible truthmakers for (19) as the measure of the rectangles that (18) and (19) (respectively) correspond to in the space generated by *a's exact quantity of cheese* and *a's exact quantity of corn*.

That is, the measure of the set of possible truthmakers for (18) will be the area of the rectangle $R_{(18)} = [20, 40] \times [30, 60]$ (where the first interval corresponds to the quantity of cheese a must have for (18) to be true, and the second corresponds to the quantity of corn a must have for (18) to be true); i.e. 600. Similarly, the measure of the set of possible truthmakers for (19) will be the area of the rectangle $R_{(19)} = [19, 39] \times [25, 55]$; i.e. 600 (see figure 4.2).

Given measures for the sets of possible truthmakers for (18) and for (19), we can determine the measure of the set of states that are possible truthmakers for only one of them as

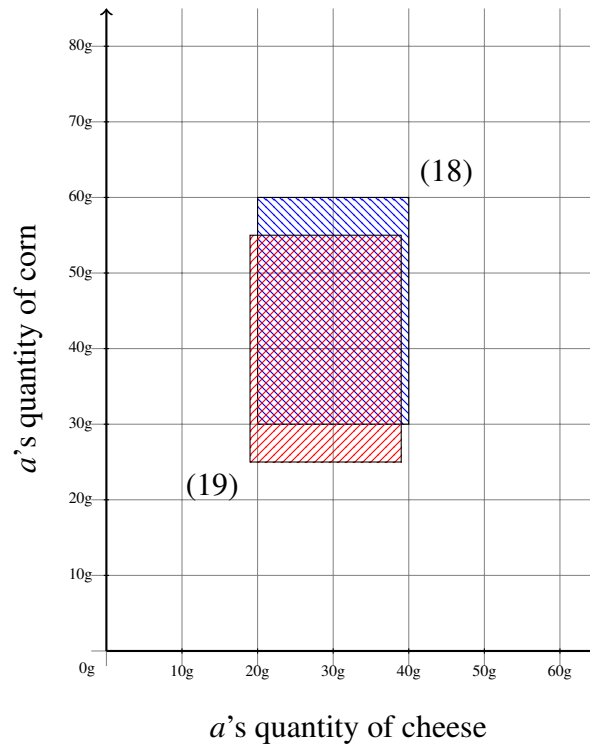


Figure 4.2: (18) and (19) in the space generated by a 's quantity of cheese and corn.

follows. The set of possible truthmakers for both (18) and (19) corresponds to the rectangle $R_{(18)\wedge(19)} = ([20, 40] \cap [19, 39]) \times ([30, 60] \cap [25, 55]) = [20, 39] \times [30, 55]$; accordingly, the measure of the set of possible truthmakers for both (18) and (19) is that rectangle's area: 475. Accordingly, the measure of the set of possible states that are truthmakers for only one of (18) and (19) is 250: i.e. it is the measure of the total area (18) and (19) occupy ($600+600-475 = 725$), minus the measure of the set of states that are truthmakers for both (475), which results in 250.

Given this measure of the set of states that are possible truthmakers for only one of (18) and (19), whether those two propositions count as similar will depend on the similarity threshold we use. For example, if we only care about a 's quantity of cheese and corn give or take 10g, then two propositions will be similar just in case the measure of the set of states that are truthmakers for only one of them is less than 400. In that case, (18) and (19)

will count as similar. On the other hand, if we care about a 's quantity of cheese and corn give or take 3g, then two propositions will be similar just in case the measure of the set of states that are possible truthmakers for only one of them is less than 36. In that case, (18) and (19) will not count as similar.

As the examples illustrate, the present account of truth-conditional similarity together with my preferred way of measuring the sets of possible truthmakers for propositions whose truth depends on the same object's properties along the same dimensions seems to yield the right results. There are, however, two salient issues. The first is that, even once we settle the relevant similarity threshold, similarity comes in degrees, but the present definition of truth-conditional similarity does not. The second, more pressing concern, is that the present notion of truth-conditional similarity yields some results that do not accord with our intuitions of similarity. I will discuss each of these issues in turn.¹⁸

The first of these issues is not too difficult to address. Given a minimum threshold for truth-conditional similarity τ , it is not too difficult to define a function that takes us from the measure of the set of states that make only one of P and Q true to a degree of similarity between P and Q. In particular, the degree of similarity between P and Q relative to threshold τ , $\delta_\tau(P, Q)$, can be defined as follows:

$$(20) \quad \delta_\tau(P, Q) = 1 - \frac{\mu((P \cap \neg Q) \cup (\neg P \cap Q))}{\mu((P \cap \neg Q) \cup (\neg P \cap Q)) + \tau}$$

Given this definition, the closer the degree of similarity between P and Q is to 1, the more similar P and Q are (see figure 4.3 for illustration).¹⁹

¹⁸ A less pressing issue is the following. In my discussion of the example with (18) and (19) I used grams to measure a 's quantities of cheese and corn, but I could just as well have used ounces, kilograms, pounds, or any other weight unit. Though the use of those different units would have resulted in different measures for (18) and (19), and would have required an adjustment of the similarity threshold, it would not have affected the ultimate judgments as to which propositions are similar.

¹⁹ (20) may not be the only natural way of assigning degrees of similarity to pairs of propositions given a similarity threshold, but it seems reasonably natural. In any case the notion of degrees of similarity will not play a substantive role in what follows. Thanks to Jim Pryor for discussion.




Figure 4.3: Mapping from values for $\mu((\mathbf{P} \cap \neg \mathbf{Q}) \cup (\neg \mathbf{P} \cap \mathbf{Q}))$ to degrees of similarity, with $\tau = 20$.

The second, more pressing problem, is this. Take for instance the following extremely precise propositions:

(21) x is exactly 1m tall

(22) x is exactly 400m tall

The set of possible truthmakers for each of these propositions corresponds to a single point in the space generated by x 's height. Accordingly, there are only two possible states of affairs that make only one of (21) and (22) true. But given the present way of measuring sets of propositions, any set with countably many members will receive measure 0. Accordingly, (21) and (22) will be counted as similar regardless of what similarity threshold we use.²⁰

A potential solution is to adopt a different kind of measure when comparing propositions whose sets of possible truthmakers have only countably many members. For example, in those cases we could use a counting measure instead of the Lebesgue measure we have been using so far.

²⁰ Thanks to Cian Dorr for pressing me on these issues.

One potential objection to this approach comes from considering the propositions:

(23) x is between 2 and 4m tall

(24) x is between 2 and 4m tall, or 400m tall

(25) x is between 2 and 4m tall, or 4.1m tall

Given that all of (23)–(25) have uncountably many possible truthmakers, using a counting measure to compare them is a non-starter. Thus, we should compare these propositions using the kind of measures I have been discussing throughout this section. However, if we compare those propositions using those kinds of measures, we predict that (23)–(25) are maximally similar, and some people may think that this is an odd prediction.

I don't find that prediction to be so odd: the measure-theoretic notion of similarity compares propositions on the basis of the size of their overlap in possible truthmakers. If we keep this way of comparing propositions firmly in mind, the prediction is to be expected. Nevertheless, I do acknowledge that the present prediction may be motivation enough to consider other alternatives, so I develop an alternative in appendix B. For present purposes, however, we can set the present discussion aside. For those purposes, we are considering only sentences ordinary language users would be prone to utter, and propositions that would be reasonably natural candidates for those utterances truth-conditional contents. But, presumably, ordinary speakers will rarely make utterances for which propositions like (24) or (25) would be natural candidates.

This ends the presentation of my preferred implementation of truth-conditional similarity. To summarize, I introduced a way of measuring the sets of possible truthmakers for propositions whose truth depends on the same object's properties along the same dimensions. Using that measure, we can make sense of the idea that two propositions are truth-conditionally similar just in case there are very few possible truthmakers for one that

are not also possible truthmakers for the other. Finally, I considered a potential objection to the resulting notion of truth-conditional similarity, but claimed it can be set aside for present purposes.

Even setting aside the objections I just discussed, there is still work to do. For example, I have focused on the comparison of propositions whose truth depends on the same object's properties along the same dimensions, but we may wonder whether propositions whose truth depends on different dimensions can be meaningfully compared too. Given the present dialectical situation, we won't need to worry about such cases. For the purposes of solving the problems from Variance and Variance*, all we need is the ability to compare propositions whose truth depends on the same object's properties along the same dimensions. I will argue for this point in the next section. Before continuing, I will state my preferred implementation of Looseness using the elements introduced in this section.

4.5.3 Descriptive similarity

Let me summarize the discussion so far. In the previous section I introduced the schema I called Looseness. That schema defined necessary conditions for the truth of an homophonic report in terms of a relation X between the semantic content of the sentence embedded in the report's complement clause and the truth-conditional content of the utterance the report was about. In this section I developed particular ways of implementing such constraints. Putting those constraints together, we can define the following relation between propositions:

Descriptive similarity: P and Q are descriptively similar relative to similarity threshold τ just in case:

- (i) P and Q are materially equivalent,
- (ii) P and Q have matching subject matters, and

(iii) P and Q are truth-conditionally similar relative to τ .

Very roughly, the notion of descriptive similarity tries to capture the way in which two different propositions may be similarly good descriptions of the same facts.

Using the relation of descriptive similarity, we can state the following implementation of Looseness:

Descriptive looseness: Necessarily, for any sentence ϕ , denoting terms S and U, and context c with similarity threshold τ , \lceil S said that ϕ through U \rceil is true in English as used in context c just in case the referent of U in English as used in c is an assertoric utterance whose truth-conditional content is *descriptively similar to* the proposition expressed by ϕ in English in c *relative to* τ .

According to Descriptive looseness, the truth of a targeted report depends on a contextually determined similarity threshold. The next section discusses this view's predictions in relation to the problems arising from Variance and Variance*.

4.6 Knowledge of homophonic reports

In section 3 I introduced the schema I called “Looseness”, which made use of a relation X between propositions. There, I said that for all sentences we replace ‘ ϕ ’ with and utterances U, as long as the proposition that ϕ and all the propositions a normal language user could easily have believed to be U’s truth-conditional content (including U’s actual truth-conditional content) stand in relation X to one another, ordinary language users will be in a position to know that S said that ϕ through U.

Something analogous is true of Descriptive Looseness. As long as the proposition expressed by ϕ in the context in which the report is made and all the propositions a normal

language user could easily have believed to be the truth conditional content of the utterance denoted are descriptively similar relative to the threshold τ that is relevant to context c , ordinary language users will be in a position to know that $\lceil S$ said that ϕ through $U \rceil$ is true as used in c , and they will also be able to know that the proposition that report expresses in that context is true. This section shows that this condition is satisfied at least in cases in which the report and the utterance are made in the same context. That is, in general, the proposition an homophonic report characterizes the speaker as having said through a given utterance relative to the context in which that utterance is made and all the propositions normal language users could easily have believed to be the truth-conditional content of the reported utterance are descriptively similar to one another relative to the relatively lax similarity thresholds we would expect to be relevant in ordinary contexts.²¹

To begin with, recall the way in which I argued for Variance and Variance*. I said that, for nearly every utterance, there are enormously many propositions any language user could easily have believed to be that utterance's truth-conditional content; given the number of such propositions, it would be extremely unlikely for any two language users to believe the same proposition to be that utterance's truth-conditional content.

Part of the support for Variance and Variance* came from a method for generating different propositions any language user could easily have believed to be a given utterance's truth-conditional content. In particular, I said that if the truth of an utterance depends on an object's properties along certain continuous (or sufficiently fine-grained) dimensions, we could generate a huge number of equally plausible candidates for that utterance's truth-conditional content. Each of those candidates corresponds to slightly different cutoff points along the dimensions on which the utterance's truth depends. For example, given that the

²¹ It is worth noting that the less similar c is from the context U occupies in whatever respects are relevant to determining the proposition expressed by ϕ , the less likely it is that the proposition expressed by ϕ in c will be descriptively similar to U 's truth-conditional content relative to reasonable similarity thresholds.

truth of Anna's utterance of 'Charlie is green' depends exclusively on Charlie's color, and that there are enormously many extremely similar and equally natural ways of distinguishing between regions of the color spectrum that qualify as green and regions that don't, there will be an equally large number of propositions one could easily have taken to be the truth-conditional content of Anna's utterance. Each of those propositions will correspond to one of the ways of dividing the color spectrum into colors that qualify as green and colors that don't.

What I will argue now is that, typically, for all sentences ϕ and utterances U of ϕ , the proposition expressed by ϕ in English relative to the context in which U is made and all the truth-conditional content candidates for U generated in the way I just outlined are descriptively similar to one another—given the kind of measures I introduced in section 5.2 and relatively lax similarity thresholds. The first step is to point out that, given standard semantic assumptions, if U is an utterance of ϕ , then U 's truth-conditional content is the proposition expressed by ϕ in the context in which U is made.²² Given this, all we need to show is that, typically, all of U 's truth-conditional content candidates are descriptively similar to one another. In particular, we need to show that, typically, those candidates are materially equivalent, have matching subject matters, and are truth-conditionally similar relative to reasonable similarity thresholds.

Material equivalence. The various truth-conditional content candidates an utterance may have are not necessarily equivalent. However, by assumption, their truth depends on the same object's properties along the same dimensions. The difference between the various truth-conditional content candidates an utterance may have is just the region in the space generated by those dimensions where that object's properties must fall in order for those candidates to be true. By assumption, that difference is very slight, so the region in the

²² See e.g. von Stechow and Heim (2011).

space generated by those dimensions in which the relevant object must fall in order for two truth-conditional content candidates to have different truth-values is very small. In this way, it is unlikely that one of those two candidates would be true and the other false.

Using the way of measuring propositions I introduced in the previous section, we could put the point as follows. Take for example, (18) and (19), repeated here:

(18) *a* is a folded tortilla with between 20 and 40g of cheese,
and with between 30 and 60g of corn

(19) *a* is a folded tortilla with between 19 and 39g of cheese,
and with between 25 and 55g of corn

Now, suppose these two propositions are both truth-conditional content candidates for an utterance of ‘*a* is a quesadilla’. Above I said that, given plausible assumptions, we could assign the set of states that make both propositions true measure 475, and the set of states that make only one of them true measure 250. Thus, even if we confined ourselves to the set of states that make at least one of (18) and (19) true, there are nearly two times more possible states that make both propositions true than states that make only one of them true.

The measure of the set of states that make both (18) and (19) false is the measure of the whole space generated by the dimensions on which the truth of (18) and (19) depends—in this case, ∞ —minus the measure of the set of worlds in which at least one of them is true (i.e. 725), which results in measure ∞ . Clearly, this measure added to the measure of the set of worlds in which (18) and (19) are both true (i.e. 475) is considerably larger than the measure of the set of worlds in which only one of them is (i.e. 250). In this sense, it is much more likely that a state relevant to the truth of (18) and (19) makes both of these propositions true or both false, than that it makes only one of them true. In turn, this means that it is mmuch more likely for (18) and (19) to have the same truth-value than not.

Now, one problem with the present argument is that it defines the measure of the set of possible states that make both (18) and (19) true or both false relative to the set of all possible states that are fully relevant to those propositions' truth-value. However, one may think that some of those states obtain only in worlds that are too far from the actual world to deserve consideration. For example, in a possible world in which *a* has 300 tons of cheese, (18) and (19) will both be false, but that possibility is so remote that we shouldn't even consider it. The likelihood that (18) and (19) have the same truth-value should be defined relative only to close possibilities.

I agree with this suggestion, but I don't think it makes a difference to the argument. For example, suppose we restrict ourselves to worlds in which *a* has at most 100g of cheese and at most 100g of corn. Then the measure of the set of possible states that make both (18) and (19) false will be the measure of the totality of the space generated by *a*'s quantity of cheese and corn up to 100g, minus the measure of the set of worlds in which at least one of them is true (i.e. 725). That is, $(100 \cdot 100) - 725 = 9275$. This, added to the measure of the set of states that make both (18) and (19) true yields measure 9750, which is still considerably greater than the measure of the set of states that make only one of (18) and (19) true (i.e. 250).

More generally, because the differences between the truth-conditional content candidates for the same utterances are so small, it will be more likely than not that they all have the same truth-value. In addition, as Schiffer (Manuscript) puts it, careful speakers will usually make sure that if they attribute a property through a given object by means of an assertoric utterance, that object falls well within the limits of that property. For example, someone who thinks that an utterance of 'Anna is tall' in the context she finds herself in has the truth-conditional content that Anna is at least 1.6m tall will be hesitant to make such an utterance if she knows that Anna is exactly 1.6m tall. Given this observation, cases in

which two different truth-conditional content candidates have different truth-values will be rare (though, of course, such cases may obtain).

Matching subject matters. Earlier in chapter 2 and then in section 5.1 of the present chapter, I said that if the truth of a proposition Fx depends on x 's properties along various dimensions, then Fx 's possible truthmakers and falsitymakers are the minimal states of affairs that specify all of x 's maximally determinate properties along those dimensions. Thus, since a proposition's exact subject matter is just the set of all its possible truthmakers and falsitymakers, Fx 's exact subject matter will be the set of all possible states of affairs that specify all of x 's maximally determinate properties along those dimensions.

By assumption, the truth of the different truth-conditional content candidates for an utterance of a simple predication depends on the same object's properties along the same dimensions. Thus, given what I just said in the previous paragraph, all those truth-conditional content candidates have the same exact subject matter. And if they have the same subject matter, it follows that they have matching subject matters. Moreover, as I claimed in section 2.1 (p. 164), if P and Q have matching subject matters, R and S have matching subject matters, every state in P 's exact subject matter is compatible with every state in R 's exact subject matter, and every state in Q 's exact subject matter is compatible with every state in S 's exact subject matter, then $P \vee R$ and $Q \vee S$ will have matching subject matters, as will $P \wedge R$ and $Q \wedge S$.

Truth-conditional similarity. In the previous section I presented a method for comparing propositions whose truth depends on the same object's properties along the same dimensions. The measures defined using that method are the ones that figure in the definition of descriptive similarity. Now, by assumption, the truth of an utterance's truth-conditional content candidates depends on the same dimensions, so those content candidates can be

compared using the method I introduced. Also by assumption, all the truth-conditional content candidates for a given utterance differ from each other only slightly. That is, if the truth of those candidates depends on an object's properties along dimensions d_1, d_2, \dots, d_n , the regions of the space generated by those dimensions in which those candidates are true will differ from each other very slightly. In this way, for any two such candidates, the region of the space generated by those dimensions in which one candidate is true and the other false will be very small. Now, presumably, in ordinary contexts we will adopt relative lax similarity thresholds, in which case the various truth-conditional content candidates one could easily take to be a given utterance's truth-conditional content will count as similar by those similarity thresholds.

This ends my argument to the effect that the candidate truth-conditional contents for typical utterances are all descriptively similar. If those arguments are successful, the truth of Descriptive Looseness puts us in a position to explain why we typically know true homophonic reports despite the truth of Variance and Variance*. Before moving on to some further issues pertaining to Descriptive Looseness, I want to consider an objection.

My arguments in this section all rely on the assumption that the truth of all the truth-conditional content candidates for typical utterances depends on the same object's properties along the same dimensions. In the present dialectical context, that assumption seems warranted. However, an objector may claim that the same kind of reasons that support Variance and Variance* may also support *Dimensional Variance*, the thesis that, for nearly every utterance and any two language users, there is no set of dimensions such that those two language users believe that that utterance's truth-conditional content is a proposition whose truth depends exactly on the same object's properties along those dimensions. If Dimensional Variance is true, even Descriptive Looseness may be too strong to account for

our knowledge of homophonic reports.

I don't think Dimensional Variance is true. In my development of the case for Variance in chapter 1, I offered several examples showing that, given a fixed set of dimensions, we can generate an enormous number of extremely similar and equally plausible truth-conditional content candidates for one and the same utterance. For example, when I presented the case for Variance, I said that there are many things one may reasonably call 'quesadillas', depending on the amount of cheese they had. Given the huge number of such different plausible denotations for the word 'quesadilla', it would be unlikely that any two language users took that word to have exactly the same denotation. And if that's the case, then (assuming basic principles of compositionality) it would be unlikely that any two language users took an utterance involving the word 'quesadilla' to have the exact same truth-conditional content.

I am skeptical that the present objector can offer similar examples which can motivate Dimensional Variance. For example, consider again the word 'quesadilla'. Presumably, that word's application depends not only on an object's quantity of cheese, but also on its shape, size, how melted the cheese is, and perhaps on the amount of corn it has. It is difficult enough to think of further (natural enough) dimensions on which the application of the word *quesadilla* may depend, let alone enormously many equally natural and extremely similar sets of dimensions on which the application of the word 'quesadilla' may depend. And without a big enough number of such sets for a big enough number of utterances, the case for Dimensional Variance does not get off the ground.

To conclude, I acknowledge that Dimensional Variance really would be problematic for the view I articulated here. However, as I have noted, it is difficult to see how a compelling case for Dimensional Variance could be made. Before concluding, I want to briefly discuss a few further issues. First, I will discuss the question of how to expand the view so as to

explain our knowledge of homophonic reports of homophonic reports. Second, I will consider a potential generalization of Descriptive Looseness to all kinds of indirect reports (as opposed to merely homophonic ones). Finally, throughout my discussion I have assumed that ordinary sentences in English really do express unique propositions and that utterances really do have truth-conditional contents, but the case for Variance casts doubt on these assumptions; the question is how to adapt the view I have presented here if it turns out that typical utterances do not have truth-conditional contents or that sentences do not express propositions.

4.7 Further issues

4.7.1 Homophonic reports of homophonic reports

Suppose John utters (2), repeated below, and consider the report (26)

(2) Anna said that Charlie is green through her utterance of ‘Charlie is green’.

(26) John said that Anna said that Charlie is green through her utterance of ‘Charlie is green’ through his utterance of ‘Anna said that Charlie is green through her utterance of ‘Charlie is green’.

According to Descriptive Looseness, (26) (and the proposition it expresses) is true as used in a context in which the similarity threshold is τ just in case the proposition expressed by ‘Anna said that Charlie is green through her utterance of ‘Charlie is green’ in the context in which the report is made and the truth-conditional content of John’s utterance of (2) are descriptively similar relative to τ . That is, just in case those two propositions: (a) are materially equivalent; (b) have matching subject matters; and (c) are similar relative to τ .

The problem is that so far I have only defined subject matters for simple predications and Boolean compounds thereof, and I have only given a way of constructing reasonable measures for simple predications whose truth depends on an object's properties along certain dimensions. The question thus arises how to extend the present notions of subject matter and measure so as to account for reports like (26). That is the question this subsection is concerned with. I'll start by describing the subject matters of homophonic reports, and then I will describe a way of meaningfully comparing the propositions expressed by homophonic reports.

Subject matters. I have been stressing that a proposition's exact subject matter is just the set of all its possible truthmakers and falsitymakers. So the question is what the possible truthmakers and falsitymakers are for the proposition (2) expresses. In order to answer this question, let c be a context with similarity threshold τ . My proposal is that if a given proposition P is descriptively similar to the proposition expressed by 'Charlie is green' in context c relative to τ , then every minimal state in which P is the truth-conditional content of Anna's utterance of 'Charlie is green' is a possible truthmaker for (2) as it is used in c . Similarly, if P is *not* descriptively similar to the proposition expressed by 'Charlie is green' in c relative to τ , then every minimal state in which P is the truth-conditional content of Anna's utterance of 'Charlie is green' is a possible falsitymaker for (2) as used in c .

Note that, since (2)'s possible truthmakers or falsitymakers are only defined relative to a context, so is its exact subject matter. However, for every context and proposition, either that proposition is descriptively similar to the proposition expressed by 'Charlie is green' in that context relative to that context's similarity threshold, or it is not. Thus, relative to every context, (2)'s exact subject matter is the set of minimal states s such that, for some proposition P , s is a state in which P is the truth-conditional content of Anna's utterance of

‘Charlie is green’. Hence, we may just say without further relativization that (2)’s subject matter is the set of minimal states s such that, for some proposition P , s is a state in which P is the truth-conditional content of Anna’s utterance of ‘Charlie is green’.

More generally, we can define the possible truthmakers and falsitymakers for homophonic reports as follows. Where ϕ is a declarative sentence, and S and U are denoting terms,

(27) s is a possible truthmaker for the proposition expressed by ‘ S said that ϕ through U ’ as used in context c just in case, for some proposition P descriptively similar to the proposition expressed by ϕ in c relative to c ’s similarity threshold, s is a minimal state in which P is U ’s truth-conditional content.

(28) s is a possible falsitymaker for the proposition expressed by ‘ S said that ϕ through U ’ as used in context c just in case, for some proposition P that is *not* descriptively similar to the proposition expressed by ϕ in c relative to c ’s similarity threshold, s is a minimal state in which P is U ’s truth-conditional content.

Given what I said in the paragraph immediately above, we can define the exact subject matter of an homophonic report as follows. The exact subject matter of the proposition semantically expressed by ‘ S said that ϕ through U ’ as used in context c is the set of minimal states in which, for some proposition P , P is U ’s truth-conditional content.

It follows from this definition that any two propositions about what somebody said through the same utterance have the same exact subject matter. Thus, any two such propositions will have matching subject matters.

Truth-conditional similarity. In section 5 I offered a method for comparing propositions whose truth depends on the same object’s properties along the same dimensions. In particular, I introduced a way of assigning degrees of similarity to pairs of propositions on

the basis of the measure of the set of states that make only one of them true. That method takes advantage of the fact that, if the truth of a proposition depends on an object's properties along various dimensions, we can generate a space using those dimensions and then translate the measure from those regions to a measure of various sets of possible states.

Unfortunately, I know of no entirely natural way of extending that method to the present case. The reason is simply that I know of no natural way of thinking of two different propositions about what somebody said through a given utterance as depending on the same object's dimensions. As a provisional solution to this problem, I want to propose that two propositions about what somebody said through a given utterance are truth-conditionally similar relative to threshold τ just in case the propositions they attribute to the speaker a saying of are themselves truth-conditionally similar relative to τ .

Consider for instance (29) and (30):

(29) Anna said that Charlie is one of colors 2-5 through her utterance of 'Charlie is green'

(30) Anna said that Charlie is one of colors 3-6 through her utterance of 'Charlie is green'.

According to the present proposal, (29) and (30) will be similar relative to threshold τ just in case the proposition that Charlie is one of colors 2-5 is similar to the proposition that Charlie is one of colors 3-6 relative to τ . For example, if the measure of the set of states that make only one of those propositions true is 2, and the similarity threshold is 3 (i.e. a difference of 3 colors in the present scale would not make a difference to our purposes), then Charlie is one of colors 2-5 and Charlie is one of colors 3-6 will be similar relative to the present threshold, and so are (29) and (30).²³

²³ In the same way, the degree of similarity between these two propositions relative to a given similarity

Admittedly, this is not entirely satisfying. For, in principle, (29) and (30) may turn out to have relatively few truthmakers in common despite the similarity between the propositions they characterize Anna as having said. However, the present approach has the advantage of allowing Descriptive Looseness to make concrete predictions about the truth of homophonic reports of homophonic reports whenever it makes a prediction about the truth of the reported homophonic report in the first place.

4.7.2 A generalization of Descriptive Looseness

Throughout this chapter I have restricted my discussion to homophonic reports, but homophonic reports are only one kind of indirect report, and one may wonder whether my account can be generalized to other indirect reports. For those who wonder, the answer is that it can. To begin with, let's start with a tentative proposal:

- (31) $\llbracket S \text{ said that } \phi \rrbracket^{c,w} = T$ if and only if, in w , S made an utterance whose truth-conditional content is descriptively similar to $\llbracket \phi \rrbracket^c$ relative to c 's similarity threshold.

There are two salient ways in which this proposal is too restrictive. The first is that, in general, we can truly report what somebody said by ascribing to her the saying of something entailed by the truth-conditional content of an utterance she made. For example, if Anna utters

- (32) Bob and Carla passed the exam,

the following reports of her utterance are true (though, admittedly, a bit misleading)

- (33) Anna said that Bob passed the exam,

threshold equals the degree of similarity between the propositions `Charlie is one of colors 2-5` and `Charlie is one of colors 3-6` relative to that similarity threshold.

(34) Anna said that Carla passed the exam.

But the proposition that Bob passed the exam is not descriptively similar to the proposition that Bob and Carla passed the exam, since (among other things) they don't have matching subject matters.

The second problem is that indirect reports may be true even if they don't report what the speaker said through a single utterance. For example, if Anna utters (35) and (36), we can truly report what she said through (37):

(35) Roses are red,

(36) Violets are blue,

(37) Anna said that roses are red and violets are blue.

However, there is no utterance of Anna's whose truth-conditional content is descriptively similar to the proposition that roses are red and violets are blue—since the proposition that red and violets are blue's exact subject matter does not match that of the proposition that roses are red or of the proposition that violets are blue.

In light of these problems, we can amend (31) as follows (where $\bigwedge \Gamma$ is the conjunction of all the propositions in the set Γ):

(38) $\llbracket S \text{ said that } \phi \rrbracket^{c,w} = T$ if and only if, for some set of propositions Γ each of which is the truth-conditional content of an utterance S made in w , $\llbracket \phi \rrbracket^c$ is descriptively similar to a proposition entailed by $\bigwedge \Gamma$ relative to c 's similarity threshold.²⁴

A problem with this new proposal is that it counts too many intuitively false reports as true. For example, it predicts that Anna said that either grass is green or it is not if she

²⁴ Entailment here is to be understood in the classical sense. However, we could also use what we may call "relevant entailment": P relevantly entails Q just in case every possible truthmaker for P has a truthmaker for Q as a part, and Q 's subject matter is encompassed by P 's exact subject matter. Unfortunately, appealing to relevant entailment will not help solve the kind of problems I discuss below, in footnote 25.

utters ‘pandas are cute’. However, it can be straightforwardly combined with a proposal I offer elsewhere (Abreu Zavaleta Forthcoming) in order to make the right predictions in this case. As it stands, the present view also fails to account for phenomena related to differences in cognitive significance, de re reports, and what I have called information-sensitivity elsewhere.²⁵

For present purposes, we can set those problems aside. What I want to discuss instead is a potential problem for (38) which arises specifically from the use of descriptive similarity. Suppose Anna utters (39); presumably, (40) is false, but (38) may seem to count it as true:

(39) The table is between 3 and 4m long

(40) # Anna said that the table is between 3.001m and 3.999m long

Suppose that in fact the table is 3.5m long, that we assign measures to (39) and (40) relative to the table’s length in meters, and that for our purposes we can tolerate a discrepancy in the table’s length of at most one tenth of a meter. Given these assumptions, the proposition that the table is between 3 and 4m long and the proposition that the table is between 3.001 and 3.999m tall will be descriptively similar. Thus, (38) predicts that (40) is true relative to the present possible world, measure, and similarity threshold.

I don’t think this sort of problem requires a modification of (38). For, according to (38), the descriptive similarity between (39) and (40) depends in part on a contextually determined similarity threshold, and it seems natural to think that, by using precise vocabulary such as ‘between 3.001 and 3.999m long’ (whether in an unembedded context or in the context of a speech report), one would be changing the contextually relevant similarity

²⁵ Even if we understand entailment as relevant entailment, the resulting view will make some inaccurate predictions. For example, that view inaccurately predicts that if Anna utters ‘cardinals are crimson’, she can be accurately reproted as having said that cardinals are crimson or cerulean. The proposal I develop in “Weak speech reports” avoids this and other problems.

threshold so as to make it much more constraining. For example, by using such vocabulary, one may affect the context so that even a difference of one micrometer would make a difference to our present purposes, hence modifying the similarity threshold so that the proposition that the table is between 3 and 4m long and the proposition that the table is between 3.001 and 3.999m tall are not descriptively similar relative to that threshold.

To be clear, my proposal is that the contextually-relevant similarity threshold can be affected both by an utterance of ‘the table is between 3.001m and 3.999m long’ and by an utterance of (40). The assumption that we can modify the similarity threshold indirectly, through an utterance of an indirect report, allows us to explain the following interesting asymmetry.²⁶

As we just saw, an utterance of (40) seems false as a report of Anna’s utterance of (39); however, a report of Anna’s utterance of (41) by means of (42) would seem perfectly acceptable in many contexts:

(41) The table is between 3.001m and 3.999m long

(42) Anna said that the table is between and 3 and 4m long.

The assumption that we can modify the contextually relevant similarity threshold indirectly, by means of an indirect report, explains this. If by uttering (42) we can relax the contextually relevant similarity threshold, that would explain why (42) strikes us as a true report of (41) at least in certain cases.

Needless to say, this kind of mechanism for modifying the contextually determined similarity threshold may not always be effective. For example, in the context of a delicate experiment (42) would strike us as false even if its typical effect would be to loosen the relevant similarity threshold. Attempting to describe the exact mechanisms by which the

²⁶ Thanks to Jim Pryor for discussion.

contextually relevant similarity threshold is determined falls out of the scope of this chapter, but I am confident such mechanisms can be captured in a systematic way.

4.7.3 Troubles with truth-conditional nihilism

Variance, Variance*, and the thesis I called Uncertainty in chapter 1, are all theses about ordinary speakers' beliefs about the truth-conditional content of ordinary utterances. However, we can make a case that the truth of those theses supports the further thesis I'll call *truth-conditional nihilism*, or *nihilism*, for short. This is the thesis that virtually no assertoric utterances have truth-conditional contents.

Nihilism is problematic for the views I have presented in this chapter. For example, according to Descriptive Looseness, the truth of an homophonic report requires that the utterance it is a report of has a truth-conditional content. Thus, if Nihilism is true, Descriptive Looseness entails that nearly every homophonic report is false, and that view's generalization from the previous section entails that nearly every speech report is false. Another problem is that the views I have considered here presuppose that sentences in English taken at a given context really do express propositions. Yet, if the proposition expressed by a sentence at a context is just the truth-conditional content of an utterance of that sentence in that context, then that presupposition fails.²⁷

I will leave a discussion of arguments for Nihilism for future work. Instead of examining such arguments, I will take for granted that Nihilism is true and explain how Descriptive Looseness can be adapted to take Nihilism into account.²⁸

²⁷ Thanks to Jim Pryor and Stephen Schiffer for pressing me on some of the issues in this subsection.

²⁸ An argument for Nihilism roughly in the spirit of the case for Variance is due to Schiffer (Manuscript). That argument goes roughly as follows. For nearly every utterance, there are enormously many extremely similar and equally eligible truth-conditional contents that utterance could easily have had. Given that all those candidates are equally eligible, and that at most one of them is the utterance's actual truth-conditional content, it would be hard to explain in virtue of what the one candidate that is the utterance's truth-conditional content is that utterance's truth-conditional content—why that candidate, instead of another? Absent such an

As a first step, we can let the propositions various language users take to be an utterance's truth-conditional content play the role of truth-conditional contents. In particular, we can reformulate Descriptive Looseness as follows:

Descriptive Looseness[†]: Necessarily, for any sentence ϕ , denoting terms S and U , and context c with similarity threshold τ in which the speaker is S' , $\lceil S$ said that ϕ through $U \rceil$ is true in English as used in context c just in case the proposition S' believes to be expressed by ϕ in c is descriptively similar, relative to τ , to the proposition the referent of S in English in c believes to be the truth-conditional content of the referent of U in English in c .

Now, Descriptive Looseness[†] is entirely compatible with Nihilism, but it is incompatible with Uncertainty—the thesis that, for nearly every utterance and any language user, there is no proposition that language user takes to be that utterance's truth-conditional content. In previous chapters I have accommodated Uncertainty by assuming that, though ordinary language users may not have determinate beliefs about which proposition is an utterance's truth-conditional content, we can think of them as being uncertain over which of many candidates are the utterance's truth-conditional content.

Adapting this strategy to the present discussion, we can reformulate Descriptive Looseness as follows:

Descriptive Looseness^{††}: Necessarily, for any sentence ϕ , denoting terms S and U , and context c with similarity threshold τ in which the speaker is S' , $\lceil S$ said that ϕ through $U \rceil$ is true in English as used in context c just in case every proposition S' believes to be a live candidate for the proposition expressed by ϕ in c is descriptively similar,

explanation, none of those candidates is the utterance's actual truth-conditional content. I hope to discuss Schiffer's argument in future work.

relative to τ , to a proposition the referent of S in English in c believes to be the truth-conditional content of the referent of U in English in c , and vice versa.

In other words, if Uncertainty is true, what we need to compare are the propositions the reporter takes to be live candidates for the proposition expressed the sentence she uses to characterize what someone else said, and the propositions the latter takes to be live candidates for her utterance.²⁹

4.8 Conclusion

In this chapter I considered an account of homophonic reports I called Strictness. I argued that, given the truth of Variance and the related thesis Variance*, Strictness makes homophonic reports very difficult to know. But, presumably, one of the easiest ways of knowing what the speaker said through a given utterance is through an homophonic report. So Strictness must be false.

I offered an alternative account, Descriptive Looseness, according to which homophonic reports are true just in case they attribute to the speaker the saying of a proposition descriptively similar to the truth-conditional content of the utterance they report. I showed that, if Descriptive Looseness is true, Variance and Variance* represent no obstacle to our knowledge of homophonic reports.

²⁹ As in previous chapters, if it turns out that language users are usually uncertain as to which truth-conditional content candidates they are uncertain between, we can iterate the present strategy as many times as necessary. For example, suppose that S is uncertain as to which of sets $\Gamma_1, \Gamma_2, \dots, \Gamma_n$ is such that S is undecided as to which of its members is U's truth-conditional content. Then the set of propositions that should play a role in determining the truth-value of an homophonic report about U should be the union of $\Gamma_1, \Gamma_2, \dots, \Gamma_n$. Alternatively, we could just take the set Γ such that S is uncertain as to which of the propositions in Γ is U's truth-conditional content, and it is definitely the case that S is uncertain as to which of the propositions in Γ is U's truth-conditional content, and it is definitely definitely the case that..., and so on, ad infinitum. I hope to explore these options in more detail in future work. Thanks to Stephen Schiffer for pressing me on these issues.

Part of the interest of the discussion concerns Dorr and Hawthorne's (2014) puzzle of semantic plasticity. According to Dorr and Hawthorne, the semantic content of ordinary utterances is very modally plastic: even small microphysical changes can lead to changes in what proposition an utterance expresses. They argue that this result has puzzling consequences for the semantics of speech reports. In particular, they argue that if the plasticity thesis is true, then we are rarely in a position to know that reports of the form 'If it had been that ϕ , S would have said that ψ ' are true.

The difficulties arising from Variance which I discussed here are not very different from Dorr and Hawthorne's puzzle about speech reports. As such, the strategy I adopted to resolve the difficulties arising from Variance can also be of help in solving their puzzle. In particular, since adopting Descriptive Looseness can make sense of the idea that a speech report may be true even if it does not attribute to the speaker the saying of the exact same proposition she semantically expressed, Descriptive Looseness is in a position to explain our knowledge of the counterfactual speech reports that Dorr and Hawthorne are interested in. I hope to develop an implementation directly addressing Dorr and Hawthorne's worries about counterfactual reports in future work.

Chapter 5

Variance, competence, and pluralism

Throughout this dissertation I have argued for Variance and discussed its consequences for various linguistic phenomena: the transmission of relevant information, disputes over an utterance's truth, and indirect speech reports. I have argued that available accounts of those phenomena fail given the truth of Variance, and offered novel accounts that explain some of those phenomena's key features while accommodating Variance.

This chapter offers no new arguments or explanations. Instead, it offers more speculative remarks on the consequences of the discussion so far. In particular, I will focus on its consequences for a general theory of the nature of communication and linguistic competence. I conclude with a brief enumeration of future areas of inquiry.

The nature of communication. We have seen that, despite the truth of Variance, we can exchange relevant information with each other through ordinary assertoric utterances, have substantive disputes over those utterances' truth, and even know what other people say by means of them. Towards the end of chapter 1 I observed, together with the coordination of action by linguistic means, all these phenomena are independent from one another.

I believe the independence of those seemingly related phenomena suggests we should

adopt a pluralistic picture of communicative success. According to this picture, there is no single natural phenomenon that we talk about when we talk about communication. Rather, there are several independent and equally natural communication-like features a conversation may have. From this pluralistic perspective, the assumption that ordinary language users have the same beliefs about the truth-conditional content 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-conditional content of their utterances, conversations between them will tend to have all or most of those features, among others.

From this perspective, Vraiance illuminates the way to a more complex conception of communicative success, according to which communication as standard accounts conceive it can be thought of as a limit or ideal of communicative success. Ordinary conversations rarely, if ever, reach that ideal, but they may still succeed in having some or all the communication-like features I have described, among others. 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 they are communication-like in those ways.

Linguistic competence. According to a popular view in semantics, to know the meaning of a sentence is to know its truth-conditional content. For example, Heim and Kratzer start their famous textbook by stating “To know the meaning of a sentence is to know its truth-conditions” (1998, p. 1), and Portner (2005) follows them when he states “The knowledge of meaning involves (at least) knowledge of the conditions under which a sentence is true, and those under which it is false” (p.13). These brief statements do not take into account the

possibility of context sensitivity, but it is reasonable to think that, once context sensitivity is taken into account, linguists like Heim, Kratzer, and Portner, would claim that knowing the meaning of a sentence requires knowing the truth-conditional content of utterances of that sentence.

The case for Variance directly challenges this account. If Variance is true, ordinary language users rarely know the truth-conditional content of ordinary utterances. Thus, unless we are willing to postulate a very radical error theory, knowledge of an utterance's meaning cannot require knowledge of its truth-conditional content. So what, then, does it take to know the meaning of a sentence?

I believe this question is misguided. To ask what it takes to know the meaning of a sentence suggests that there is something to be known in the first place—perhaps a proposition of the form *S means such-and-such*. But, as the discussion in chapters 2–4 suggests, we can explain several linguistic phenomena without making any assumption about what the meaning of an utterance is, or even about what its truth-conditional content is. In particular, all the explanations I have offered so far appeal only to people's beliefs about the truth-conditional content of various utterance, without assuming that those utterances in fact had this or that truth-conditional content. Even in my discussion of indirect reports, in which the assumption that utterances had truth-conditional contents played a prominent role in the beginning of the discussion, towards the end I argued that this assumption was not necessary.

I believe the right question to ask is what it takes to be *competent* in the use of a sentence. Importantly, talk about competence here is meant not to presuppose that there is some particular proposition one must know, or even believe, to be competent in the use of a sentence. What I want to propose is that being competent in the use of a sentence is just being able to do various things with it. For example, being able to exchange relevant

information through utterances of that sentence with other language users; being able to have disputes with other language users over the truth of utterances of that sentence which exhibit epistemic conflicts; being able to know what other people said through utterances of that sentence; etc.

More precisely, my proposal is this:

A language user S is competent in the use of a sentence E as used by population X just in case, typically,

- (i) S can exchange relevant information with members of X through utterances of E;
- (ii) S can have partly factual disputes over the truth of utterances of E with members of X;
- (iii) S knows what members of X said through utterances of E; etc.

Needless to say, this list is incomplete, but it can at least give us an idea of what an adequate theory of competence should look like.

Interestingly, this approach constitutes a kind of social externalism about linguistic competence. For, if this proposal is true, then being competent in the use of a sentence does not depend exclusively on one's intrinsic properties. Given the accounts of information transmission, partly factual disputes, and speech reports I developed in chapters 2–3, competence depends both on the propositions S takes to be live candidates for an utterance's truth-conditional content and on the propositions members of X take to be live candidates for that utterance's truth-conditional content. This kind of social externalism, it is worth noting, is entirely independent of standard versions of social externalism about mental content.

Further issues. Aside from a further investigation of the various ways in which a conversation can be communication-like and of the notion of linguistic competence, there are several relevant questions this dissertation has left untouched. Those questions include:

- **Do ordinary utterances have truth-conditional contents?** If Variance is true, then for almost every utterance, at most one language user knows its truth-conditional content. Given how epistemically inaccessible truth-conditional contents are, it is worth asking whether an utterance's *actual* truth-conditional content could play any role in psychological or semantic theorizing. And if it doesn't, there is real pressure to reject the idea that ordinary utterances really have truth-conditional contents. The question is then whether we can continue to understand disciplines like semantics as the study of the way in which an utterance's truth-conditional content is determined on the basis of the semantic contents of its components in the context in which the utterance is made.
- **Can knowledge be transmitted by testimony?** Principles roughly like the following have played a prominent role in the literature on knowledge by testimony: *if a speaker knows P and asserts P to her audience, and that audience accepts P on the basis of the speaker's testimony, then the audience knows P*. Variance raises a distinctively linguistic problem for principles of this kind: if it is true, we are rarely in a position to know what propositions the speaker asserted in the first place, let alone come to accept those propositions on the basis of the speaker's assertion. And if this is so, it is an open question how (and even, *if*) knowledge can be transmitted by testimony, and what such transmission consists in.
- **Can the accounts in this dissertation be extended?** Throughout this dissertation I have focused on very simple cases of information transmission, partly factual dis-

putes, and speech reports. Most of those cases involve very simple predications of the form Fx and Boolean compounds thereof. Furthermore, in most of those cases, disagreement over an utterance's truth-conditional content arises from disagreement over the property an utterance attributes to a given object. However, it is natural to wonder, first, whether the approach can be extended so as to account for cases involving quantification and non-monadic properties; second, whether the approach can be extended to cases in which people disagree about what objects an utterance is about. These are pressing questions which I leave for future work.

- **Is Variance related to vagueness?** Throughout this dissertation I omitted any discussion of the relationship between Variance and vagueness. However, one may wonder if the truth of Variance depends on the fact that natural languages are vague. I believe the answer to this question is 'no': Variance is true due to the fact that different people may have slightly different evidence on the basis of which to form beliefs about the truth-conditional content of ordinary utterances, and that those people may be differently attuned to the evidence. This would be so even if natural language was fully precise. If anything, the relation between vagueness and Variance goes the other way around: the vagueness of terms in natural language arises due to the truth of Variance. The question is whether a view of that kind could explain vagueness in thought, and not only in language.

This list is not exhaustive. Without a doubt, there are many other ways in which Variance interacts with standard accounts of various linguistic and psychological phenomena. In this way, our general understanding of language and communication may benefit from further study of Variance and its consequences.

Appendix

A distance-based notion of truth-conditional similarity

In section 5.2 I considered a measure-theoretic notion of truth-conditional similarity. That notion of similarity is enough for the purposes of solving the problems arising from Variance, and can help us make sense of the idea that two propositions have large regions of overlap in logical space. However, it also yields the wrong results when it comes to comparisons of propositions that have measure zero. In this appendix I wish to offer an alternative notion of truth-conditional similarity that does not have this problem: the distance-based conception of truth-conditional similarity.

The underlying intuition behind this approach is that two propositions are truth-conditionally similar just in case every possible truthmaker for one is “close enough” to a possible truthmaker for the other, and vice versa. As we will see now, the notion of closeness at play here is to be taken quite literally, as a distance relation between possible states of affairs relative to a certain space.

To begin with, recall that the truth of a proposition Fx depends on x 's properties along dimensions d_1, d_2, \dots, d_n —e.g. size, weight, degree of similarity to paradigmatic mem-

bers of a certain class, etc.—just in case there is a set of points in the space generated by those dimensions such that, necessarily, Fx is true if and only if x 's properties fall in one of those points. As I also said in chapter 4 and earlier, in chapter 2, if the truth of a proposition depends exactly on x 's properties along d_1, d_2, \dots, d_n , that proposition's possible truth-makers will be the minimal states determining x 's maximally determinate properties along d_1, d_2, \dots, d_n in which Fx is true. Thus, if the truth of a proposition depends exactly on x 's properties along d_1, d_2, \dots, d_n , we can think of that proposition as a set of points in the space generated by those dimensions—i.e. $d_1 \times d_2 \times \dots \times d_n$.

Given such a space, we can calculate the distance between two points in the standard way. That is, if $a = (a_1, a_2, \dots, a_n)$ and $b = (b_1, b_2, \dots, b_n)$ are two points in $d_1 \times d_2 \times \dots \times d_n$, their distance $\delta(a, b)$ is defined as follows:

$$\delta(a, b) = \sqrt{(a_1 - b_1)^2 + (a_2 - b_2)^2 + \dots + (a_n - b_n)^2}$$

What I want to propose is to think of the distance between two minimal possible states of affairs that determine a possible way for x 's properties to be with respect to dimensions d_1, d_2, \dots, d_n as the distance between the corresponding points in $d_1 \times d_2 \times \dots \times d_n$.

Given that notion of distance, we can define a new notion of truth-conditional similarity as follows. Propositions P and Q are similar relative to similarity threshold τ just in case, for every possible truthmaker p^T for P, there is a truthmaker q^T for Q such that the distance between p^T and q^T is less than or equal to τ . So, for example, the propositions x is exactly 3m tall and x is exactly 3.1m tall will be similar as long as the similarity threshold is less than or equal than 0.1m.

As compelling as this view may seem, there is significant room for improvement. For example, above I said that the truth of the proposition x is a folded tortilla

filled with between 20 and 40g of cheese, and with between 30 and 60g of corn depends on x 's quantity of cheese and corn. But it may well be that, for certain purposes, we are willing to tolerate greater differences in x 's quantity of cheese than in x 's quantity of corn. Thus, we would be much better off using a relation of similarity that was sensitive to different similarity thresholds for different dimensions.

Thus, instead of relativizing truth-conditional similarity to a single similarity threshold, we should relativize it to a sequence of similarity thresholds. In particular, if we think of the possible truthmakers for a proposition whose truth depends on dimensions d_1, d_2, \dots, d_n as points in the space generated by those dimensions, we can define truth-conditional similarity as follows:

Distance-based similarity: P and Q are similar relative to similarity thresholds $\tau_1, \tau_2, \dots, \tau_n$

just in case: for every possible truthmaker for P $p^T = (p_1, p_2, \dots, p_n)$, there is a truthmaker for Q $q^T = (q_1, q_2, \dots, q_n)$ such that, for each pair of p_i and q_i , $|p_i - q_i| \leq \tau_i$, and vice versa.¹

For example, if our standard for similarity tolerates a discrepancy of at most 3g of cheese and at most 5g of corn, the present view predicts that the propositions

(18) x is a folded tortilla filled with between 20 and 40g of cheese,
and with between 40 and 60g of corn

(19) x is a folded tortilla filled with between 19 and 39g of cheese,
and with between 25 and 55g of corn

are similar relative to those similarity thresholds. On the other hand, if our standard for similarity tolerates a discrepancy of at most 1g of cheese, the view will predict that those two propositions are not similar.

¹ Here, $|x - y|$ denotes the absolute value of $x - y$.

The present approach's main advantage over the measure-theoretic approach is that it can meaningfully compare propositions regardless of the size of the sets of their possible truthmakers. Another significant advantage is that it allows for more flexibility in the determination of similarity thresholds.

Bibliography

- Abreu Zavaleta, Martín. Forthcoming. “Weak Speech Reports.” *Philosophical Studies* .
- Abreu Zavaleta, Martín. Manuscript. “Disagreement Lost.”
- Bach, K. and Harnish, R. 1979. *Linguistic Communication and Speech Acts*. MIT Press.
- Bach, Kent. 1994. “Conversational Implicature.” *Mind and Language* 9:124–162.
- Barwise, Jon and Perry, John. 1981. “Situations and Attitudes.” *Journal of Philosophy* 78:668–691.
- . 1983. *Situations and Attitudes*. MIT Press.
- 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 Hawthorne, John. 2009. *Relativism and Monadic Truth*. Oxford University Press.
- Cappelen, Herman and Lepore, Ernest. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Wiley-Blackwell.
- Chalmers, David J. 2011. “Verbal Disputes.” *Philosophical Review* 120:515–566.
- Choi, Sungho. 2008. “The Incompleteness of Dispositional Predicates.” *Synthese* 163:157–174.
- Christensen, David. 2007. “Epistemology of Disagreement: The Good News.” *Philosophical Review* 116:187–217.

Bibliography

- Ciardelli, Ivano, Groenendijk, Jeroen, and Roelofsen, Floris. 2013. "Inquisitive Semantics: A New Notion of Meaning." *Language and Linguistics Compass* 7:459–476. ISSN 1749-818X. doi: 10.1111/lnc3.12037.
- Cohen, Jonathan. 2013. "Indexicality and the Puzzle of the Answering Machine." *Journal of Philosophy* 110:5–32.
- Cresswell, Max J. 1977. "The semantics of degree." In Barbara Partee (ed.), *Montague Grammar*, 261–292. New York: Academic Press.
- DeRose, Keith. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52:913–929.
- 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.
- Egan, Andy. 2010. "Disputing About Taste." In Ted Warfield and Richard Feldman (eds.), *Disagreement*, 247–286. Oxford, UK: Oxford University Press.
- Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41:478–502.
- Evans, Gareth. 1982. *Varieties of Reference*. Oxford University Press.
- Fara, Delia Graff. 2001. "Descriptions as Predicates." *Philosophical Studies* 102:1–42.
- 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.
- . MSA. "A theory of truth-conditional content, part I." URL:https://www.academia.edu/9186494/Truthconditional_Content_-_Part_I, DOA: August 25, 2016.
- . MSb. "A theory of truth-conditional content, part II." URL:https://www.academia.edu/9206919/Truthconditional_Content_-_Part_II, DOA: August 25, 2016.
- . MSc. "Truthmaker Semantics." URL: https://www.academia.edu/10908756/survey_of_truthmaker_semantics, DOA: August 25, 2016.

Bibliography

- Fodor, Jerry A. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. MIT Press.
- Frances, Bryan and Matheson, Jonathan. 2018. “Disagreement.” In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2018 edition.
- 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.
- Groenendijk, Jeroen and Roelofsen, Floris. 2009. “Inquisitive semantics and pragmatics.” In *Meaning, Content, and Argument: Proceedings of the ILCLI International Workshop on Semantics, Pragmatics, and Rhetoric*. www.illc.uva.nl/inquisitive-semantics.
- Groenendijk, Jeroen and Stokhof, Martin. 1989. *Type-Shifting Rules and the Semantics of Interrogatives*, 21–68. Dordrecht: Springer Netherlands. ISBN 978-94-009-2723-0. doi: 10.1007/978-94-009-2723-0_2.
- Groenendijk, Jeroen and Stokhof, Martin. 1982. “Semantic Analysis of Wh-Complements.” *Linguistics and Philosophy* 5:175–233.
- Hamblin, Charles L. 1973. “Questions in Montague English.” *Foundations of Language* 10:41–53.
- 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. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, UMass Amherst.
- . 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.
- Hintikka, Jakko. 1969. “Semantics for propositional attitudes.” In *Models for Modalities*, 87–111. Springer.
- Hirsch, Eli. 1993. *Dividing Reality*. Oxford University Press.

Bibliography

- Hovda, Paul. 2009. "What Is Classical Mereology?" *Journal of Philosophical Logic* 38:55–82.
- Huemer, Michael. 2011. "Epistemological Egoism and Agent-Centered Norms." In Trent Dougherty (ed.), *Evidentialism and its Discontents*, 17. Oxford University Press.
- Huvenes, Torfinn Thomesen. 2012. "Varieties of Disagreement and Predicates of Taste." *Australasian Journal of Philosophy* 90:167–181.
- JBiz et al. 2010. "How fast do you have to be going to count it as running?" URL: http://www.letsrun.com/forum/flat_read.php?thread=3576961, DOA: July 27, 2017.
- Jenkins, C. S. I. 2014a. "Merely Verbal Disputes." *Erkenntnis* 79:11–30.
- . 2014b. "Serious Verbal Disputes: Ontology, Metaontology, and Analyticity." *Journal of Philosophy* 111:454–469.
- 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. "The Disagreement Challenge to Contextualism." In Jonathan Jenkins Ichikawa (ed.), *The Routledge Handbook of Epistemic Contextualism*.
- 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.
- Kolbel, Max. 2004. "Faultless Disagreement." *Proceedings of the Aristotelian Society* 104:53–73.
- Kratzer, Angelika. 1989. "An Investigation of the Lumps of Thought." *Linguistics and Philosophy* 12:607–653.
- . 2012. "Facts: Particulars or Information Units?" In *Modals and Conditionals*, 160–183. New York: Oxford University Press.

Bibliography

- . 2016. “Situations in Natural Language Semantics.” In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Spring 2016 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.
- Lasersohn, Peter. 2005. “Context Dependence, Disagreement, and Predicates of Personal Taste.” *Linguistics and Philosophy* 28:643–686.
- 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.
- . 1988. “Statements Partly About Observation.” *Philosophical Papers* 17:1–31.
- . 1991. *Parts of Classes*. Blackwell.
- 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.
- MacFarlane, John. 2005. “Making Sense of Relative Truth.” *Proceedings of the Aristotelian Society* 105:321–339.
- . 2007. “Relativism and Disagreement.” *Philosophical Studies* 132:17–31.
- . 2014. *Assessment Sensitivity: Relative Truth and its Applications*. OUP Oxford.
- 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.

Bibliography

- Michaelson, Eliot. 2014. "Shifty Characters." *Philosophical Studies* 167:519–540.
- Partee, Barbara Hall. 2004. "Binding Implicit Variables in Quantified Contexts." In *Compositionality in Formal Semantics: Selected Papers of Barbara Partee*. Blackwell.
- paul. 2004. "Ki." URL:<https://www.urbandictionary.com/define.php?term=ki>, DOA: April 13, 2018.
- 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.
- Predelli, Stefano. 1998. "I Am Not Here Now." *Analysis* 58:107–115.
- . 2013. *Meaning Without Truth*. Oxford University Press.
- Prior, E. W. 1985. *Dispositions*. Humanities Press.
- Putnam, Hilary. 1975. "The Meaning of 'Meaning'." *Minnesota Studies in the Philosophy of Science* 7:131–193.
- Roberts, Craige. 2012. "Information structure in discourse: Towards an integrated formal theory of pragmatics." *Semantics and Pragmatics* 5:1–69. doi: 10.3765/sp.5.6.
- Rooth, Mats. 1992. "A Theory of Focus Interpretation." *Natural Language Semantics* 1:75–116.
- 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.
- . Manuscript. "When Meaning Meets Vagueness."
- Schoubye, Anders Johan. 2013. "Ghosts, Murderers, and the Semantics of Descriptions." *Noûs* 47:496–533.
- Schwarz, Florian. 2012. "Situation pronouns in determiner phrases." *Natural Language Semantics* 20:431–475.
- Sidelle, Alan. 2007. "The Method of Verbal Dispute." *Philosophical Topics* 35:83–113.
- Sider, Theodore. 2006. "Quantifiers and Temporal Ontology." *Mind* 115:75–97.
- . 2011. *Writing the Book of the World*. Oxford University Press.

Bibliography

- Stalnaker, Robert. 1974. "Pragmatic Presuppositions." In Robert Stalnaker (ed.), *Context and Content*, 47–62. Oxford University Press.
- . 1978. "Assertion." *Syntax and Semantics (New York Academic Press)* 9:315–332.
- . 1999. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford University Press UK.
- . 2009. "On Hawthorne and Magidor on Assertion, Context, and Epistemic Accessibility." *Mind* 118:399–409.
- Stanley, Jason and Szabó, Zoltan Gendler. 2000. "On Quantifier Domain Restriction." *Mind and Language* 15:219–261.
- Stephenson, Tamina. 2007. "Judge Dependence, Epistemic Modals, and Predicates of Personal Taste." *Linguistics and Philosophy* 30:487–525.
- Stevens, Jon Scott. 2017. "Pragmatics of Focus." URL: <http://linguistics.oxfordre.com/view/10.1093/acrefore/9780199384655.001.0001/acrefore-9780199384655-e-207>, DOA: February 2, 2018.
- 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.
- Tarski, Alfred. 1956. *Logic, Semantics, Metamathematics*. Oxford, Clarendon Press.
- van Inwagen, Peter. 1996. "It Is Wrong, Everywhere, Always, for Anyone, to Believe Anything Upon Insufficient Evidence." In Jeff Jordan and Daniel Howard-Snyder (eds.), *Faith, Freedom and Rationality*, 137–154. Savage, Maryland: Rowman and Littlefield.
- Vermeulen, Inga. 2018. "Verbal Disputes and the Varieties of Verbalness." *Erkenntnis* 83:331–348.
- Vision, G. 1985. "I Am Here Now." *Analysis* 45:198–199.
- von Fintel, Kai and Heim, Irene. 2011. "Intensional semantics. Ms., Massachusetts Institute of Technology."
- von Fintel, Kai-Uwe. 1994. *Restrictions on quantifier domains*. Ph.D. thesis, University of Massachusetts Amherst.

Bibliography

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.

—. forthcoming. "Relevance Without Minimality." In Andy Egan and Dirk Kindermann (eds.), *Unstructured Content*. Oxford: Oxford University Press.