

© 2019

Cameron Domenico Kirk-Giannini

ALL RIGHTS RESERVED

COMMITMENT, COMMUNICATION, AND CONTENT:

TOWARD A THEORY OF ASSERTION

By

CAMERON DOMENICO KIRK-GIANNINI

A dissertation submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Philosophy

Written under the direction of

Ernie Lepore and Jeffrey C. King

And approved by

New Brunswick, New Jersey

May, 2019

ABSTRACT OF THE DISSERTATION

Commitment, Communication, and Content:

Toward a Theory of Assertion

by CAMERON DOMENICO KIRK-GIANNINI

Dissertation Directors:

Ernie Lepore and Jeffrey C. King

According to a familiar and appealingly simple picture of assertion, the propositional content asserted by a speaker is both that which she communicates to members of her audience and that to which she undertakes a distinctive sort of commitment. In what follows, I develop a criticism of this familiar picture and propose an alternative account according to which assertoric content is more intimately associated with commitment than communication.

I begin by motivating a common claim about rational communication: that it cannot proceed when interlocutors are uncertain which contents utterances contribute to discourse. It emerges during my discussion of this claim that, given certain natural assumptions about the conditions under which it is rational for interlocutors to assert, it is possible to construct an argument from the premise that the propositional content asserted by a speaker is the information she communicates to members of her audience to the conclusion that speakers always assert *diagonal propositions* of their utterances — that is, propositions which characterize the information interlocutors can learn from their utterances by assuming that the propositions they semantically determine are true.

I proceed to argue that this latter claim, to which I refer as *Diagonalism*, systematically conflicts with our intuitive judgments about the conditions under which

the contents of speakers' assertions would be true or false. It follows that the failure of Diagonalism requires us to abandon the claim that the propositional content asserted by a speaker is the information she communicates to members of her audience.

I then argue that, even if assertoric content does not play the role in communication it has often been thought to play, the familiar picture of assertion is at least half right: the propositional content asserted by a speaker is indeed that to which she undertakes a distinctive sort of commitment. I motivate this claim in two ways. First, I propose a diagnostic for identifying the propositional content to which a speaker becomes primarily committed in asserting and show that this diagnostic singles out assertoric content even in cases where what is asserted is not what is communicated. In the course of this argument, it emerges that the content of an assertion is usually its *horizontal proposition* — that is, the proposition it semantically determines in the context in which it is in fact uttered. Second, then, I argue that a linguistic community which adopts a practice of regarding its speakers as assertorically committed to the horizontal propositions of their utterances enjoys benefits not enjoyed by communities which adopt alternative committal practices, such as the practice of regarding speakers as assertorically committed to the diagonal propositions of their utterances.

Acknowledgements

To say that this dissertation would not have existed without the support of Ernie Lepore would in fact be to understate his significance in my intellectual development, since it was his mentorship and enthusiasm which first led me to conceive of myself as a philosopher of language shortly after I arrived at Rutgers. Without Ernie, then, neither this dissertation nor any even remotely like it would have existed. As co-director of my dissertation, he has worked tirelessly to help me distill the ideas underlying my arguments in as broadly accessible, framework-independent a way as possible. In the process, he has helped me to see a number of new connections between my work and traditional questions in the philosophy of language which I anticipate will shape the direction of my research for years to come.

During my time at Rutgers, I have come to rely on Jeff King's detailed, perspicuous, no-nonsense feedback as a crucial test of the viability of my ideas and arguments. If an argument is enthymematic, Jeff will notice. If a term is introduced as a predicate of objects of one type and then used as a predicate of an object of some other type, Jeff will notice. As co-director of my dissertation, Jeff has helped me tremendously to clarify my thoughts about assertion, communication, commitment, and the relationships between these topics and the metasemantic question of how contexts of utterance determine the semantic values of lexical items.

The other members of my committee have also played crucial roles in the writing process. Andy Egan has been unfailing in his support and encouragement, giving me detailed feedback on numerous drafts of every chapter as well as helping me to revise my manuscripts in light of comments from journal referees. I

am especially grateful to him for his intimate understanding of and irrepressible enthusiasm for the Stalnakerian possibility-carving framework, which have buoyed me through many a dark night of the soul concerning the tenability of coarse-grained approaches to content. Liz Camp has challenged me to think more carefully about the notion of assertoric commitment and the connection between semantics and pragmatics. She has also given me invaluable advice about how to approach philosophical writing in a way which puts ideas first without compromising formal rigor. Paul Pietroski has kept me honest by challenging me to defend the foundational assumptions of my work: assumptions about the natures of representation, linguistic and mental content, and grammar.

John Hawthorne has been the best external committee member anyone could have hoped for. I will always have fond memories of my semester as a visiting student at USC, and I am tremendously grateful to him for the many hours he spent during that time reading drafts of my work and giving me feedback on it. More than anyone else, John has been able to occupy multiple theoretical perspectives in order to anticipate not only how philosophers of various persuasions might object to my arguments but also what a committed possibility-carver might say in reply to them.

Ofra Magidor also deserves special thanks here. During my frequent summer visits to Oxford, she has always been ready to meet to discuss my work. Indeed, it was her feedback on an early draft of the second chapter of this dissertation which made me realize that I needed to write the first.

I am grateful to my fellow graduate students in the philosophy of language, both for their feedback on my work and for their role in creating and sustaining the quirky, hilarious, and deeply intellectually serious language community at

Rutgers. Simon Goldstein, Una Stojnić, Bob Beddor, and Peter van Elswyk were early role models for me. Sam Carter, Caley Howland, Kelsey Laity-D'Agostino, Lauren Richardson, and Andrew Rubner have been wonderful colleagues during the time we have overlapped in the program. I am also grateful to the graduate students in other areas who have given me feedback on my work as part of the Dissertation Seminar and Third-Year Seminar here — Austin Baker, Ben Bronner, Eddy Keming Chen, Nate Flores, Chris Hauser, and Oli Odoffin — as well as to the faculty members who taught those seminars — Jonathan Schaffer and Susanna Schellenberg.

Finally, I would like to thank all those who, though their friendship, have made my time at Rutgers one of the happiest in my life. I am especially grateful to Eli Shupe, Carolina Flores, Simon and Kathryn Goldstein, and Peter and Kat van Elswyk.

This is, of course, a very incomplete list of those to whom I am grateful for help with the papers in this dissertation. A more thorough (though almost certainly still incomplete) list of those whose feedback I have incorporated into the work contained here is given at the end of each chapter.

Note: Chapter 1 (as well as parts of the description thereof in the Introduction) first appeared in the journal *Linguistics and Philosophy* (2018, volume 41, pages 665–684). Chapter 3 is currently under review for publication at another journal.

Table of Contents

Abstract	ii
Acknowledgements	iv
Introduction	1
1 Uniformity Motivated	4
1.1 Introduction	4
1.2 From (Assertion Rule) to (Uniformity)	7
1.3 Inquiry and Possibility-Carving Methodology	11
1.4 The First Argument	15
1.5 The Second Argument	24
1.6 Uniformity and Introspection	30
2 Do Not Diagonalize	37
2.1 Introduction	37
2.2 (Horizontalism) and (Diagonalism)	42
2.3 The Case Against (Diagonalism)	47

2.4	Objections and Replies	54
2.5	Assertion and Communication	61
2.6	Assertion and Commitment	67
2.7	Conclusion	74
3	Why Horizontalism	75
3.1	Introduction	75
3.2	Committal Practices Introduced	80
3.3	Six Practices	88
3.4	Generality	91
3.5	Utility I: Non-Eliminativity and Unintentional Liability	93
3.6	Utility II: Mere Intentional Communication	95
3.7	Utility II: Secondary Commitment	101
3.8	Modal Coherence	105
3.9	Partiality	109
3.10	Conclusion	113
	Bibliography	114

Introduction

Assertion is a topic of central importance to a range of subdisciplines within linguistics and philosophy. In semantics and pragmatics, for example, it is often taken to demarcate the subject matters of the two fields: the assertoric content of an utterance is both the content which a suitable semantic theory must predict and the content on which the explanatory resources of pragmatic theory are subsequently brought to bear. Similarly, epistemologists have long been interested in the power of assertion to transmit propositional knowledge and justification between speaker and audience, as well as the role assertion plays in collective processes of inquiry.

Speakers assert in order to communicate information. It is tempting, therefore, to hold that the propositional content of an assertion *just is* the information it communicates to its audience. In cases involving uncertainty about the semantic values of context-sensitive lexical items, moreover, it is tempting to hold that the information an assertion communicates to its audience *just is* what audience members are in a position to recover from it by assuming that the proposition it semantically determines is true. According to this sort of picture, which I call *diagonalism*, a speaker's utterance of 'She speaks Farsi', produced in a context involving uncertainty about who is being referred to, counts as an assertion of the proposition *that the individual being referred to speaks Farsi*.

Diagonalism is an influential and intuitively appealing theory of assertoric content. Nevertheless, I argue in the three chapters contained here, it must be rejected.

I begin in Chapter 1 ("Uniformity Motivated") by considering the question

of whether rational communication can proceed when interlocutors are uncertain which contents utterances contribute to discourse. An influential negative answer to this question is embodied in the Stalnakerian *principle of uniformity*, which requires speakers to produce only utterances that express the same content in every possibility treated as live for the purposes of the conversation. The principle of uniformity enjoys considerable intuitive plausibility and, moreover, seems to follow from platitudes about assertion; nevertheless, it has recently proven controversial. In the chapter, I defend the principle by developing two arguments for it based on premises reflecting the central aims and assumptions of *possibility-carving frameworks* for modeling inquiry—that is, frameworks which describe the evolution of individuals’ attitudinal states in terms of set-theoretic operations defined over a domain of objects representing possibilities. It emerges during my discussion that, given certain natural assumptions about the conditions under which it is rational for interlocutors to assert, it is possible to construct an argument for diagonalism from the premise that the propositional content asserted by a speaker is the information she communicates to members of her audience. The connection between diagonalism and the idea that the propositional content of an assertion is the information it communicates to its audience is thus significantly tighter than might otherwise have been thought.

In Chapter 2 (“Do Not Diagonalize”), I argue that diagonalism systematically conflicts with our intuitive judgments about the circumstances in which the contents of speakers’ assertions would be true or false. The failure of diagonalism requires that we either abandon a familiar way of thinking about information and rational assertion or hold that the propositional content of an assertoric utterance is not the information it communicates. I suggest that we respond to this dilemma by embracing its second horn — assertoric content is better character-

ized in terms of the commitments speakers undertake in asserting than in terms of the information they communicate to their interlocutors.

In Chapter 3 (“Why Horizontalism”), I explore a novel strategy for adjudicating between diagonalism and *horizontalism* — the thesis that what a speaker asserts in literally and sincerely uttering an indicative sentence is some horizontal proposition of her utterance. The strategy involves considering the advantages and disadvantages which would accrue to a linguistic community as a result of adopting different committal practices — that is, practices of associating utterances with the propositions to which speakers undertake assertoric commitments in uttering them. I ultimately conclude that a horizontalist practice has important advantages over its competitors.

1 Uniformity Motivated

1.1 Introduction

Consider the following two claims about assertion:¹

(Assertion Rule): If accepted, an assertoric utterance changes the context by adding its content to the common ground.

(Uniformity): In cases of rational communication, an assertoric utterance expresses the same proposition in each possible world in the context set.

(Assertion Rule) renders formally tractable the platitude that assertion is a species of literal linguistic communication: a speaker asserts a content, and, if all goes according to plan, her audience comes to take that same content for granted for the purposes of the conversation. What is more, it is tempting to think that (Uniformity) follows from (Assertion Rule). For, temptation suggests, if (Uniformity) were violated, interlocutors would not know how to apply (Assertion Rule). Stalnaker (2009) has this sort of argument in mind when he writes:

“...an assertion is a move in conversational game with a rule that says how the context evolves. To allow assertions with content to which the players of the game do not have access is [to] allow a situation in which the players will be unable to apply the rule. It would be like a card game in which [you are] dealt a card, face down, and the rule requires that you draw another card if and only if the card you were dealt is red. If you are not allowed to look at the card you are dealt, then you [are] not in a position to apply the rule.” (Stalnaker 2009, 407)

¹See e.g. Stalnaker 1978, 2009, 2014.

If (Uniformity) follows from (Assertion Rule), and if (Assertion Rule) approximates the status of a platitude about assertion, then there is a strong case to be made for (Uniformity). But, tempting though it may be, the inference from (Assertion Rule) to (Uniformity) is controversial. Why think that violations of (Uniformity) *must* put interlocutors in a position from which they are unable to apply (Assertion Rule)? Hawthorne and Magidor (2009, 2010), for example, have argued that if individuals can sometimes come to believe a content by accepting a sentence that expresses it, then interlocutors can sometimes update the common ground with the content expressed by an assertion even in cases of (Uniformity) violation. To the extent that this is a reasonable account of what it takes to come to believe a content, there is reason to doubt that cases in which (Uniformity) is violated are ipso facto cases in which interlocutors cannot apply (Assertion Rule), and so reason to doubt (Uniformity) itself.

The friend of (Uniformity) must therefore confront the question of whether it is possible to construct an argument for her thesis from more compelling premises. An appealing way to begin addressing this question is to focus on the class of models of assertion in inquiry within which a convincing argument for (Uniformity) can be constructed. If some such models are explanatorily successful, their success constitutes a non-question-begging argument for (Uniformity). In what follows, I pursue this strategy by showing how certain general and explanatorily fruitful assumptions about how assertion is to be modeled can serve as the basis for two arguments for (Uniformity). In so doing, I demonstrate that one can argue persuasively for (Uniformity) without appealing to a controversial theory of belief.

Dialectically, my project combines concessiveness towards an initial skepticism about (Uniformity) with a robust defense of it as a constraint on rational commu-

nication. I will not dispute that, according to some ways of thinking about belief, there is no important connection between (Assertion Rule) and (Uniformity). Instead, I will show that if one wishes to develop a theory on which the common ground of a conversation plays a certain kind of explanatory role, one ought to accept (Uniformity) if one accepts (Assertion Rule). To this end, I will help myself to a number of premises which proponents of alternative frameworks might reject but which, unlike the claim about ability required by Stalnaker's argument, capture basic facts about the explanatory aims of the Stalnakerian framework.

The interest of this project of showing explicitly and formally why anyone with certain natural theoretical commitments ought to embrace the connection between (Assertion Rule) and (Uniformity) is acknowledged, at least implicitly, by both Stalnakerians and their critics. Hawthorne and Magidor (2009, 395–6), for example, present a brief, informal argument similar to part of the first argument I develop below. Stalnaker (2009) seems to be gesturing toward the possibility of constructing such an argument when he writes that Hawthorne and Magidor's model of coming to believe a content "does not take [his] model of discourse seriously" (2009, 408), and Almotahari and Glick (2010) make a similar point when they worry that that model "sever[s] the connection between communication theory and the theory of action" (2010, 1085). Yet no one, Stalnakerian or otherwise, has seriously attempted to carry the project to completion.

The remainder of this chapter is organized into five sections. In the first, I review the details of Stalnaker's framework for modeling inquiry and present a formal reconstruction of his argument for (Uniformity) and of Hawthorne and Magidor's criticism of that argument. In the second, I describe the motivations for my methodological assumptions and introduce an ancillary premise which will be invoked in later arguments. In the third and fourth, I present my two arguments

for (Uniformity). In the fifth and final section, I argue that the second of these suggests a way for the Stalnakerian to respond to Hawthorne and Magidor's (2009) argument that (Uniformity) is sometimes violated.

1.2 From (Assertion Rule) to (Uniformity)

Stalnaker's work on assertion in inquiry stands within an influential tradition in the study of content-bearing mental states, which models the contents of those states as sets of entities (*points*): each point represents a distinct possibility, and each set of points represents the content that is compatible with all and only those possibilities which correspond to its members.² Such *possibility-carving frameworks* are useful in that they permit facts about changes in individuals' attitudinal states to be described set-theoretically, leading to particularly elegant characterizations of processes like inquiry, hypothetical reasoning, and plan formation.

In the Stalnakerian framework, every conversation is associated with a set of *participants* or *interlocutors*. These interlocutors perform assertoric speech acts, which are governed by conversational rules. It will be helpful to separate out two aspects of each such speech act: the *sentence* assertorically uttered and the *content* (set of points) thereby asserted. The same sentence, uttered assertorically, can constitute the assertion of different contents in different contexts.

Central to the Stalnakerian framework is the notion of *presupposition*. Presupposition is a theoretically primitive attitude interlocutors bear to sets of points; an interlocutor presupposes a content just in case she takes it for granted for the

²I refer to *contents* rather than *propositions* throughout so as not to suggest that I take the possibilities represented by points to be metaphysically possible worlds. In the interest of maximal generality, this is an issue on which I wish to avoid taking a stand.

purposes of the conversation.³ The set of contents presupposed by every interlocutor in conversation c at point w is the *common ground* of c at w (henceforth CG_w^c); the set-theoretic intersection of these contents is the *context set* of c at w (henceforth CS_w^c). CS_w^c thus represents the set of possibilities treated as live for the purposes of the conversation. For the sake of simplicity, I will assume that all interlocutors in a given conversation presuppose the same contents, so that the common ground of the conversation can be identified with the presuppositions of any interlocutor.⁴

I will primarily be concerned in what follows with interactions between the contents of assertions and the context sets of the conversations in which they are performed. Let us call the set of points formed by intersecting the context set of conversation c at point w with the content of an assertoric utterance u at point w' the CS_w^c -*content* of u at w' (when the values of c and w are not important, we can also speak of *CS-content* more generally). Since the context set plays a distinctive role in Stalnaker's theory of inquiry, failures of (Uniformity) which involve CS-content are especially problematic. Indeed, in what follows I will restrict my attention to a version of the uniformity principle which concerns CS-content exclusively.⁵

Having introduced these Stalnakerian concepts, we can now give a formal

³Stalnaker (1978, 84) clarifies that this requires taking for granted that her interlocutors also take the content for granted for the purposes of the conversation.

⁴This exposition follows Stalnaker 1978. Stalnaker 2014 defines the notions of presupposition and common ground in terms of an iterated and collectivized notion of *acceptance for the purposes of the conversation* (2014, 25; 2014, 4). The resulting conception of the common ground is so stringent that it is not clear whether any nontrivial content meets its requirements, and, as it is not required for my arguments below, I will set it aside.

⁵Hawthorne and Magidor (2009, 384) call this version of the principle (*Weak Uniformity*) and note that there is textual evidence that Stalnaker prefers it to the unrestricted version of the principle.

characterization of the principle of uniformity which makes precise the sense in which it is concerned with the CS-content of an utterance:

(Uniformity): In cases of rational communication: for all conversations c , assertoric utterances u , contents p , and points w, w' : if w' is in CS_w^c and p is the CS_w^c -content of u in w' , then for all points w'' : if w'' is in CS_w^c , then p is the CS_w^c -content of u in w'' .

As we have seen, Stalnaker's argument from (Assertion Rule) to (Uniformity) relies on motivating an intuition to the effect that in cases where (Uniformity) is violated, interlocutors are not in a position to add the content of an assertion to the common ground. We can render this claim explicit as follows:

(Stuck): For all conversations c , assertoric utterances u , contents p , and points w, w' : if w' is in CS_w^c , and if p is the CS_w^c -content of u in w' , and if there is some point w'' such that w'' is in CS_w^c and the CS_w^c -content of u in w'' is not p , then some interlocutor in c is not in a position at w to apply (Assertion Rule) to CG_w^c in response to u .

Stalnaker's argument also requires a principle linking what counts as rational communication on the part of a speaker with the abilities of her audience. The intuitive idea is that a speaker ought not to assert if her audience cannot respond appropriately:

(Fair Play): In cases of rational communication: for all interlocutors s , conversations c , assertoric utterances u , and points w : s performs u at w only if every interlocutor in c is in a position at w to apply (Assertion Rule) to CG_w^c in response to u .

(Stuck) and (Fair Play) together secure the truth of (Uniformity).

Consider now the following scenario, a version of which is discussed by Hawthorne and Magidor and which provides a good illustration of the intuition behind Stalnaker's argument:

[FIRE]: Smith is watching Jones through a doorway. She can see the corridor in which Jones is standing, but not the room into which Jones is looking. It is common ground between Smith and Jones that the room contains either Bill or Ben and no one else. Bill or Ben (whoever it is) is performing a dangerous chemical experiment. Something goes horribly awry, and Jones turns to Smith and exclaims 'He is on fire!'.

Let us refer to the point representing the situation described in [FIRE] as w and the conversation between Smith and Jones as c . Since Smith does not know at w whether the referent of Jones's 'he' is Bill or Ben, there are distinct points w', w'' in CS_w^c immediately after Jones's utterance u such that in w' Jones refers to Bill and in w'' Jones refers to Ben. Suppose that in w' , u expresses a content which is true just in case Bill is on fire, and that in w'' , u expresses a content which is true just in case Ben is on fire. Because CS_w^c contains both points at which Bill but not Ben is on fire and points at which Ben but not Bill is on fire, the CS_w^c -content of u at w' is different from its CS_w^c -content at w'' . By (Stuck), then, some interlocutor (here, Smith) is not in a position at w to apply (Assertion Rule) to CG_w^c in response to u — which of the two contents of u is she supposed to add to the common ground? It follows, by (Fair Play), that Jones's utterance u does not count as a case of rational communication.⁶

⁶Stalnaker (1978) therefore proposes that speaker and audience perform a process of reinterpretation to recover a uniform content for the utterance.

Hawthorne and Magidor (2009, 2010) seek to undermine the inference from (Assertion Rule) to (Uniformity) by calling (Stuck) into question. As noted above, they argue that if an individual can come to be belief-related to a content by accepting a sentence that expresses it, then interlocutors can update the common ground with the content expressed by an assertion even in cases of (Uniformity) violation. In the case of [FIRE], what this suggestion amounts to is that Smith is able to apply (Assertion Rule) to the common ground in response to u by accepting the sentence ‘He is on fire’ and thereby coming to believe its content, whatever that happens to be.

1.3 Inquiry and Possibility-Carving Methodology

I have promised to derive (Uniformity) from premises reflecting basic facts about the explanatory aims of possibility-carving frameworks. In this section, I clarify the content of two of the possibility carver’s methodological commitments and introduce an auxiliary principle on which my arguments in later sections will rely.

The first commitment, which is central to the possibility-carving approach and will therefore be assumed without defense in what follows, is that the contents of both assertions and content-bearing mental states can be modeled as subsets of a suitably chosen set of points. The second is that there is a tight connection between facts about which possibilities interlocutors rule in and out for the purposes of the conversation and facts about what linguistic actions it is acceptable for them to perform.

This second commitment is both substantive and important to my arguments

in what follows, so it is worth discussing its motivation in some detail. The central insight which possibility-carving frameworks are supposed to capture is that when interlocutors engage in inquiry, their aim is to coordinate on an ever-more-precise conception of the world. Because the context set represents interlocutors' shared conception of the world, there must be a tight connection between what is ruled in or out for the purposes of a conversation and what constitutes reasonable linguistic behavior for its interlocutors. Thus Stalnaker:

“One cannot normally assert, command, promise, or even conjecture what is inconsistent with what is presupposed. Neither can one assert, command, promise or conjecture what is itself presupposed. There is no point in expressing a proposition unless it distinguishes among the possible worlds which are considered live options in the context.” (Stalnaker, 1970, 280)

On the possibility-carving picture, then, the theoretical role of the context set of a conversation is, in part, to constrain the speech acts which count as intelligible continuations of that conversation. Of the many ways in which the context set constrains acceptable linguistic behavior, I will focus on two — one in each of my arguments for (Uniformity). First, I will rely on the principle that, *ceteris paribus*, it is acceptable for an interlocutor to non-rhetorically ask a polar question \lceil Is it the case that ϕ ? \rceil ⁷ only if the content expressed by ϕ is neither entailed by nor inconsistent with the common ground. Second, I will rely on the principle that, *ceteris paribus*, it is acceptable for an interlocutor to assert a content just in case doing so would result in narrowing the context set without eliminating it entirely. These principles will be formulated more precisely in the next two sections.

My arguments will also depend on an ancillary premise which ensures that,

⁷Following Quine 1940, I use $\lceil \dots \mu \dots \rceil$ so that it is synonymous with the expression $\lceil \dots \frown \dots \frown \mu \frown \dots \rceil$, where \frown indicates concatenation.

despite the possibility of violations of (Uniformity), the negation operator behaves as it does in propositional logic — an assertoric utterance of a sentence always expresses the complement of the content expressed by an assertoric utterance of its negation:

(Negation): For any assertoric utterance u of sentence ϕ in conversation c at point w , if u expresses contents p_1, \dots, p_n at different points in CS_w^c , then there is a sentence $\lceil \neg\phi \rceil$ such that, for each point w' in CS_w^c and natural number m such that $1 \leq m \leq n$: if u expresses p_m at w' , then a corresponding assertoric utterance u' of $\lceil \neg\phi \rceil$ at w' expresses W/p_m .^{8,9}

Note that (Negation) is plausible only on the assumption that interlocutors are not uncertain about the meaning of ' \neg '. If they were — if, for example, they treated as live the possibility that ' \neg ' denoted the identity truth function (thus meaning *it is true that* rather than *not*) — then there would be points in the context set of the conversation at which ϕ and $\lceil \neg\phi \rceil$ expressed the same content. I will therefore make this assumption in what follows. Similar considerations

⁸Here ' W ' denotes the set of all points in the model and ' W/p'_m ' denotes the complement of content p_m in that set.

⁹Let us say that u' corresponds to u in the required way at a point w just in case it differs from u at w to the minimal extent necessary for it to constitute an assertoric utterance of $\lceil \neg\phi \rceil$ rather than ϕ . Thus u' and u are required to have the same speaker, addressee, accompanying demonstration, and so forth. More generally, two utterances of sentences ϕ, ψ correspond to each other in the required way at a point just in case they differ at that point to the minimal extent necessary for them to constitute sincere utterances of the relevant sentences. The qualification that u' must correspond to u ensures that (Negation) does not have the unintuitive consequence that (e.g.) my utterance of 'I am hungry' and your utterance of 'I am not hungry' express incompatible contents at every point. I will assume that if two utterances u and u' correspond to each other at at least one point, then they correspond to each other at any point where at least one exists, so that we may say that they correspond to each other *simpliciter*. Note also that here and in what follows my talk of utterances expressing contents should not be taken to entail that they are performed. Thus (Negation) does not require that any utterance of a sentence be accompanied by an utterance of its negation; rather, it requires that any time a sentence could be uttered so as to express certain contents at various points in the context set, its negation could also be uttered so as to express certain systematically related contents.

require that I assume that interlocutors are not uncertain about the meaning of the interrogative construction ‘Is it the case that ... ?’. The necessity of these assumptions should not be taken to undermine the generality of my arguments, however, for two reasons. First, because ‘¬’ and the English syntax for polar questions are arbitrary means of expressing negation and interrogative force, respectively, my arguments show, by parity of reasoning, that (Uniformity) must be satisfied in any conversation where there is coordination on any syntactic means of expressing negation and interrogative force which allows for the same kinds of syntactic embeddings as English. Second, if an argument for (Uniformity) can be constructed using a group of speakers who presuppose the grammar of some particular language, then there is reason to think that (Uniformity) must be satisfied even when there is uncertainty about the grammar of the language being spoken. (Uniformity) is a general principle of rational communication, and it would be odd for the applicability of a general principle of rational communication to depend on what interlocutors can take for granted about the grammar of their language. In the context of the first of my arguments below, for example, the assumption that interlocutors have coordinated on conventions for expressing negation and polar questions is important only because it allows them to express their uncertainty about certain contents. There is no reason to think that violations of (Uniformity) would not result in the same uncertainty if interlocutors lacked the means to express it. We thus have good reason to endorse the conditional claim that if there is a compelling argument for (Uniformity) when we restrict our attention to conversations in which interlocutors are certain of the meanings of the negation- and question-forming constructions of their language, then there is a compelling argument for (Uniformity) *tout court*.¹⁰

¹⁰My reliance on (Negation) for the purposes of my arguments is thus compatible with the proposal, advanced in section 6 below, that the assertoric content of any utterance is to be identified

1.4 The First Argument

In this section, I present the first of my arguments for (Uniformity). The principal premise on which I will rely explicates the intuitive relationship between what is taken for granted in a conversation and which questions are felicitous continuations of that conversation:

(Correspondence): For any conversation c , sentence ϕ , and point w , if it is acceptable for some interlocutor in c to sincerely and non-rhetorically ask the polar question \lceil Is it the case that ϕ ? \rceil at w , then there is a point in CS_w^c which is an element of the content of a corresponding assertoric utterance of ϕ at w , and there is a point in CS_w^c which is an element of the content of a corresponding assertoric utterance of $\lceil\neg\phi\rceil$ at w .¹¹

(Correspondence) enforces a connection between what counts as acceptable linguistic behavior for the participants in an inquiry and the ways in which that

with its semantic superdiagonal content.

¹¹One might object to (Correspondence) on the grounds that when interlocutors presuppose one or more falsehoods, w is not an element of CS_w^c , and therefore the acceptability of asking the polar question \lceil Is it the case that ϕ ? \rceil cannot be characterized in terms of the content a corresponding utterance of ϕ expresses at w . One moved by this sort of objection might instead endorse (i) the principle that there must be some point w' in CS_w^c such that there is a point in CS_w^c which is an element of the content of a corresponding assertoric utterance of ϕ at w' and there is a point in CS_w^c which is an element of the content of a corresponding assertoric utterance of $\lceil\neg\phi\rceil$ at w' (call this principle (\exists -Correspondence)), or (ii) the principle that every point w' in CS_w^c must be such that there is a point in CS_w^c which is an element of the content of a corresponding assertoric utterance of ϕ at w' and there is a point in CS_w^c which is an element of the content of a corresponding assertoric utterance of $\lceil\neg\phi\rceil$ at w' (call this principle (\forall -Correspondence)), or, drawing on the terminology of section 5, (iii) the principle that there must be a point in CS_w^c which is an element of the assertoric superdiagonal content of a corresponding utterance of ϕ , and there must be a point in CS_w^c which is an element of the assertoric superdiagonal content of a corresponding utterance of $\lceil\neg\phi\rceil$ (call this principle (F -Correspondence)). Though I have plumped for (Correspondence) in virtue of its simplicity, I do not need to decide between it and its three alternatives for the purposes of my argument — in the presence of (Sentential Plenitude) below, and given that I am concerned exclusively with polar questions formed using sentences which express uniform contents, any of the four premises just canvassed suffices to establish my conclusion.

inquiry can be modeled, precluding particular modeling choices which do not respect this connection. For example, as Allan Gibbard (ms) has shown, the kind of permissive view of coming to believe a content which Hawthorne and Magidor suggest leads quickly to the conclusion that inquiry is a trivial enterprise.¹² For it is a feature of the linguistic competence of a normal English speaker that she knows that the sentence ‘The world is like *this*’ expresses a truth as uttered in any context. If accepting a sentence as uttered on a particular occasion is sufficient for ruling out the points incompatible with its content, then one way for conversational participants to narrow the context set to the unique point representing the way things actually are is to utter this sentence, intending with ‘this’ to rigidly designate all that is, and then accept it; the content expressed by ‘The world is like *this*’, uttered with the aforementioned intention, excludes every point except the point representing the world of utterance.

Intuitively, no model which allows speakers to coordinate on a single point so easily could provide insight into inquiry as it actually occurs. (Correspondence) solves this sort of problem, since it ensures that if it is acceptable for interlocutors to sincerely ask a variety of polar questions even after accepting the sentence ‘The world is like *this*’ (as it evidently is), then accepting that sentence must not involve their ruling out all points but one from the context set of the conversation.

My first argument for (Uniformity) requires two additional premises. The first of these guarantees that there are sentences in our language which can be used

¹²Gibbard’s argument concerns individual propositional attitudes rather than the context set and attempts to demonstrate that, if knowing a proposition under any guise is sufficient for knowing it *simpliciter*, speakers can know any true proposition *a priori*. The similarity between Gibbard’s argument and the argument given here, however, should be apparent. Schlenker (2011) discusses a related sort of argument, due to von Stechow and Zimmerman (2005), which demonstrates that, given certain permissive assumptions about when sentences of the form ‘ S believes that ϕ ’ are true, speakers can correctly be credited with believing every true proposition.

uniformly to express certain contents:

(Sentential Plenitude): For any conversation c , content p , point w , and assertoric utterance u at w , for each point w' in CS_w^c , if u expresses p at w' , then there is a sentence ψ such that an assertoric utterance u' of ψ expresses p at w and every point in CS_w^c .^{13,14}

The idea behind (Sentential Plenitude) is that, if there is an assertoric utterance which violates (Uniformity), there must also be assertoric utterances which do not violate (Uniformity) and which express the various contents that could, given the context set, be expressed by the original assertion. This will usually be the case when (Uniformity) is violated because of the use of a context-sensitive expression, for there will usually be context-insensitive expressions which have each of the possible semantic values of the original. For example, in [FIRE], Jones's utterance of 'He is on fire' putatively violates (Uniformity) because the context set doesn't settle whether 'he' refers to Ben or to Bill. But corresponding to each possible content of Jones's utterance is a sentence which, uttered assertorically, would satisfy (Uniformity) and express that content: the two relevant sentences are 'Bill is on fire' and 'Ben is on fire'. It is plausible that this phenomenon generalizes to other varieties of context sensitivity. If gradable adjectives, relational predicates with implicit arguments, quantifiers, and modals come along with contextually supplied reference classes, thresholds, restrictor properties, and so forth, then, in-

¹³Note that, so long as it is sincere and literal, u' need not correspond to u . The idea is that an utterance of ψ expresses p regardless of speaker, addressee, accompanying demonstration, etc.

¹⁴Those suspicious of (Correspondence) because of the role played in it by the point of utterance w will likely also be suspicious of the requirement in (Sentential Plenitude) that each ψ express the corresponding p at w as well as throughout CS_w^c . Fortunately, adopting any of the three alternative correspondence principles suggested above obviates the need to hold that each ψ express the corresponding p at w .

sofar as these contextually supplied parameters can be explicitly signaled ('tall for an elephant', 'local to the Upper West Side', 'permissible according to the law of New Jersey', and so forth), (Sentential Plenitude) will be satisfied. Even in cases of (Uniformity) violation resulting from ambiguity, expressions can be freely introduced into the language to unambiguously express the various possible disambiguated contents.

Moreover, even if natural languages tend not to be rich enough to satisfy (Sentential Plenitude) in all cases, if there is a possible language which does satisfy it, and if an argument for (Uniformity) can be constructed using a group of speakers who presuppose the grammar of that language, then there is reason to think that speakers of less expressive languages must satisfy (Uniformity), as well. Just as it would be odd for the applicability of a general principle of rational communication to depend on what interlocutors can take for granted about the grammar of their language, it would be odd for the applicability of a general principle of rational communication to depend on their language's expressive power. We thus have reason to endorse the conditional claim that if there is a compelling argument for (Uniformity) when we restrict our attention to conversations in which interlocutors speak a language which satisfies (Sentential Plenitude), then there is a compelling argument for (Uniformity) *tout court*.

The second additional premise guarantees that, in cases of (Uniformity) violation, it is never presupposed that the context set would have to be different than it in fact is for the speaker to express one of the potential contents of her utterance:

(Independence): For any conversation c , point w , and assertoric utterance u , if u violates (Uniformity) at w , then for each content p expressed by u at a point in CS_w^c there is a point w' in CS_w^c such that (i) $CS_{w'}^c = CS_w^c$ and (ii) u

expresses p at w' .

Unlike (Sentential Plenitude), (Independence) is a substantive constraint on the structure of the context set. Nevertheless, (Independence) is quite plausible, holding in intuitive cases of (Uniformity) violation. In [FIRE], for example, it would be odd for either Smith or Jones to presuppose that if it were Bill in the room, the common ground would have to be different than it actually is. The common ground encodes the fact that who is in the room is left open for the purposes of the conversation; it would be odd for this to be so, but for it to be taken for granted that if Ben is in the room, then who is in the room is not left open for the purposes of the conversation.¹⁵

We are now in a position to construct an argument for (Uniformity) which does not rely on (Stuck). If (Uniformity) is violated, then there will be at least two distinct contents expressed by a given utterance at different points in the context set. For any two such contents, we can distinguish two types of cases: cases in which the two contents are not ordered by strength, and cases in which they are. Before proceeding to the formal proof, it will be helpful to consider a putative example of each type.

Hawthorne and Magidor's [FIRE] provides a convenient illustration of the first kind of case. Recall that Hawthorne and Magidor suggest that by accept-

¹⁵It is worth making two small points in connection with (Independence). First, (Independence) is trivially true if we hold that, for every conversation c and point w , every point in CS_w^c has the same context set as w , so that my first argument proceeds even more smoothly if I am granted this common (though not uncontroversial) assumption. Second, (Independence) features in my argument below only in the case where (Uniformity) is violated and the various CS-contents expressed by the violative utterance at points in the context set are ordered by strength. Even if (Independence) were deemed unacceptable for some reason, the first argument would still show that all cases of (Uniformity)-violation are cases in which the relevant CS-contents are ordered by strength. This is not a conclusion many critics of (Uniformity) will be happy to accept. Thanks to an anonymous referee for emphasizing these points.

ing Jones's assertion, Smith can come to believe whatever content was in fact expressed: if Bill is in the room, Smith comes to believe that Bill is on fire, while if Ben is in the room, she comes to believe that Ben is on fire. In either case, they argue, there is no reason to believe that Smith is not in a position to apply (Assertion Rule) to the common ground of the conversation: (Stuck) is false.

But Hawthorne and Magidor's account of [FIRE] does not obey (Correspondence). For suppose it is Ben in the room. Then, according to Hawthorne and Magidor's account, when Smith accepts Jones's assertion, she thereby rules out all points in the context set at which Ben is not on fire. But it seems reasonably clear that, even after she accepts Jones's assertion, it is acceptable for Smith to sincerely and non-rhetorically ask the question 'Is it the case that Ben is on fire?'. So, by (Correspondence), there is a point in the context set which is an element of the assertoric content of 'Ben is not on fire' as uttered by Smith — that is, a point where Ben is not on fire. It follows that Smith both does and does not rule out all points in the context at which Ben is not on fire: Contradiction.¹⁶

¹⁶Almotahari and Glick (2010) offer an informal argument similar to this one except in that it involves non-linguistic action:

"Suppose the auditor in the Bill and Ben example loves Ben but hates Bill, vowing never to save him from a fire. If accepting the sentence 'He is on fire' amounts to accepting the proposition that Ben is on fire (since Ben is in fact the man in the room), then one would expect the auditor to come to the rescue. But since the auditor does not know whether Bill or Ben is on fire, she does not know whether she should come to the rescue. However, Hawthorne and Magidor's account predicts that she would come to the rescue, since accepting 'He is on fire' is enough to accept that Ben is on fire." (2010, 1085)

This is a good argument, as far as it goes. However, the plausibility of a connection between the common ground and rational non-linguistic action is much diminished in cases where what is presupposed is not believed, such as when interlocutors are engaged in hypothetical reasoning or plan formation. Even when what is presupposed is believed, the connection between the common ground and rational non-linguistic action is mediated in complex ways by individuals' utility functions. For both of these reasons, the enthymematic non-linguistic analog of (Correspondence) employed by Almotahari and Glick's argument is not suitable for the purposes of providing a fully general argument for (Uniformity).

As an illustration of the second kind of case, consider the following scenario:¹⁷

[BIRTHDAY]: Arabella and Barbarella are wrapping presents on Thursday night for their daughter Cinderella's birthday on Friday. They plan to present Cinderella with half of her presents on Friday morning and the other half on Friday evening. Arabella leaves for some time to take a work-related phone call. When she returns, Barbarella asserts, 'All the presents are wrapped; we can go to sleep.' It is common ground between the two that Barbarella either intends to convey (only) that all the morning presents are wrapped or to convey that all the morning and evening presents are wrapped, but it is not common ground between the two which of these two scenarios obtains.

Let us call the conversation between Arabella and Barbarella c and the point which represents the situation described in [BIRTHDAY] w . Since Arabella does not know at w whether Barbarella has only the morning presents or both the morning presents and the evening presents in mind, there are distinct points w', w'' in CS_w^c immediately after Barbarella's utterance u such that in w' Barbarella has only the morning presents in mind and in w'' Barbarella has both the morning and the evening presents in mind. Suppose that in w' , u expresses a content which is true just in case all the presents for Friday morning have been wrapped, and that in w'' , u expresses a content which is true just in case both the presents for Friday morning and the presents for Friday evening have been wrapped. Then the two contents Barbarella expresses at w' and w'' are ordered by strength: the

¹⁷Here I assume that quantifier domain restriction is a semantic phenomenon, with the domain of quantification determined by the speaker's intentions. Those who accept accounts of quantifier domain restriction incompatible with these assumptions will have to devise a different illustration of the second kind of case.

proposition that all the morning and evening presents are wrapped entails the proposition that all the morning presents are wrapped, but not vice versa. Moreover, even if Arabella accepts Barbarella's assertion, it is acceptable for her to sincerely and non-rhetorically ask the question 'Is it the case that all the presents for Friday evening have been wrapped?'. It follows, by (Correspondence), that there is a point in CS_w^c which is an element of the assertoric content of the sentence 'Not all the presents for Friday evening have been wrapped' as uttered by Arabella — that is, a point at which not all the presents for Friday evening have been wrapped. But, since by (Assertion Rule) all such points are ruled out if Barbarella in fact intends to convey that both the morning and evening presents have been wrapped, it follows that Barbarella must not have intended to convey this. So, using only the information that her utterance violates (Uniformity), we can conclude (somewhat surprisingly) that in w Barbarella does not intend to convey that the evening presents have been wrapped. But the same reasoning can be applied to *any* point in CS_w^c which has the same context set as w . By (Independence), at least one such point must be one where Barbarella intends to convey that both the morning and the evening presents have been wrapped. So we have a contradiction: at this point, Barbarella both does and does not intend to convey that all the evening presents are wrapped.

Now for the formal statement of the proof:

Suppose we have a violation of (Uniformity): a case of rational communication in which there is an assertoric utterance u at a point w in a conversation c and distinct points w', w'' in CS_w^c such that the CS_w^c -content of u at w' differs from the CS_w^c -content of u at w'' . Then there is some set ρ of points in the context set at w such that whether the points ρ contains are ruled out of the context set by applying (Assertion Rule) to u in w depends on which

content u expresses. Since this is, by stipulation, a case of rational communication, there is no problem with u being accepted. Suppose it is. Then one of two cases must obtain.

Case 1: Neither of the two CS_w^c -contents of u is a proper subset of the other. By (Sentential Plenitude), there are sentences ϕ and ψ which uniformly express the contents of u at w' and w'' respectively. So, by (Assertion Rule) and (Negation), accepting u involves ruling out either all the points in the CS_w^c -content of a corresponding utterance of $\lceil \neg\phi \rceil$ or ruling out all the points in the CS_w^c -content of a corresponding utterance of $\lceil \neg\psi \rceil$. But even after accepting u in w , it is acceptable for interlocutors to sincerely and non-rhetorically ask \lceil Is it the case that ϕ ? \rceil and \lceil Is it the case that ψ ? \rceil . By (Correspondence) as applied to the first of these sentences, at least one of the the points in CS_w^c is in the assertoric content of a corresponding utterance of $\lceil \neg\phi \rceil$ at w ; similarly for $\lceil \neg\psi \rceil$ and the second sentence. But then neither all the points in the CS_w^c -content of a corresponding utterance of $\lceil \neg\phi \rceil$ nor all the points in the CS_w^c -content of a corresponding utterance of $\lceil \neg\psi \rceil$ have been ruled out of the context set at w . Contradiction.¹⁸

Case 2: One of the two CS_w^c -contents of u at w' and w'' is a proper subset of the other. Then ρ is the set obtained by set-theoretically subtracting the smaller (logically stronger) CS_w^c -content from the larger (logically weaker). By (Assertion Rule), either all the points in ρ are ruled out of the context set when u is accepted in w , or none are. By (Sentential Plenitude), there is a sen-

¹⁸Since (Sentential Plenitude) guarantees that an utterance of ϕ or ψ expresses the corresponding content at any point in $\{w\} \cup CS_w^c$ regardless of speaker, addressee, accompanying demonstration, and so forth, (Negation) guarantees that an utterance of $\lceil \neg\phi \rceil$ or $\lceil \neg\psi \rceil$ expresses the complement of the corresponding content in W at any point in $\{w\} \cup CS_w^c$ regardless of speaker, addressee, accompanying demonstration, and so forth. It follows that it is not important which utterances of ϕ and ψ (or \lceil Is it the case that ϕ ? \rceil , or \lceil Is it the case that ψ ? \rceil) we take our utterances of $\lceil \neg\phi \rceil$ and $\lceil \neg\psi \rceil$ to correspond to. Similar remarks apply to case 2 below.

tence ϕ that uniformly expresses the content expressed by u at whichever of w' and w'' corresponds to the stronger of the two CS_w^c -contents. By (Negation), a corresponding utterance of $\lceil \neg\phi \rceil$ uniformly expresses the complement of this content in W . Even after accepting u in w , it is acceptable for interlocutors to sincerely and non-rhetorically ask \lceil Is it the case that ϕ ? \rceil . So, by (Correspondence), at least one of the points in CS_w^c is in the assertoric content of a corresponding utterance of $\lceil \neg\phi \rceil$ in w . Since such a point must be a member of ρ if it was not ruled out by accepting u , and since accepting u must involve either ruling out every member of ρ or no member of ρ , it follows that accepting u in w does not rule out any member of ρ . But then it must be that the weaker content of u is asserted at w . Since nothing specific about w has been assumed, this is also true at every other point w''' in CS_w^c such that $CS_{w'''}^c = CS_w^c$. But by (Independence), there is at least one such point where the stronger content of u is asserted. Contradiction.

So there can be no violation of (Uniformity). \square

1.5 The Second Argument

Suppose we associate an assertoric content with each assertoric utterance at each point where that utterance expresses a determinate content.¹⁹ Then we can define a function F from assertoric utterances u to sets of points such that, for each u , $F(u)$ is the set of points at which u expresses a truth — that is, the set of points

¹⁹An assertoric utterance may fail to express a determinate content at certain points where it does not exist, for example. I will follow Hawthorne and Magidor (2009, 389) in setting this complication aside for the purposes of my arguments below.

w such that w is an element of the assertoric content of u in w .²⁰ My second argument for (Uniformity) relies on two premises describing the relationship between the assertoric content of u and $F(u)$. The first of these premises formalizes the idea that accepting an assertion does not entitle interlocutors to rule out certain possibilities which, for all they presuppose, are compatible with its content:

(Stability): For any conversation c , assertoric utterance u , and point w , if the intersection of $F(u)$ and CS_w^c is σ , then applying (Assertion Rule) to the content of u at w does not result in eliminating any member of σ from CS_w^c .

The second premise formalizes the idea that accepting an assertion of a given sentence requires interlocutors to rule out all possibilities at which an assertion of the negation of that sentence would express a truth:

(Contraction): For any conversation c , assertoric utterance u of sentence ϕ , corresponding assertoric utterance u' of $\lceil \neg\phi \rceil$, and point w , if the intersection of $F(u')$ with CS_w^c is σ , then applying (Assertion Rule) to the content of u at w results in eliminating every member of σ from CS_w^c .

Like (Correspondence), (Stability) and (Contraction) can be defended on the basis of a fundamental methodological commitment to regarding the context set as a constraint on acceptable linguistic behavior. In the case of (Stability) and (Contraction), this commitment takes the form of a tight connection between what

²⁰ $F(u)$ is thus closely related to the Stalnakerian concept of a diagonal proposition. There is, however, an important difference: the Stalnakerian diagonal proposition of an utterance u is constructed out of the *semantically determined* contents of u at various points, whereas $F(u)$ is constructed out of the *assertoric* contents of u at various points. For a more precise characterization of the notion of the semantically determined content of an utterance at a point, see section 6.

interlocutors can acceptably assert and what effects their assertions would have on the context set — in particular, it takes the form of the principle that, *ceteris paribus*, it is acceptable to utter a sentence assertorically in the course of joint inquiry just in case that sentence expresses a content which, if accepted, would narrow but not empty the context set.

With this principle in hand, suppose (Stability) were violated — suppose, that is, that it were possible to perform an assertoric utterance u and thereby eliminate some point in $F(u)$ from the context set. Consider a case in which every point in the context set is a member of $F(u)$ (that is, a case in which it is presupposed that u would express a truth). If the goal of assertion is to narrow the context set, and if this can be done by uttering u , as is predicted given (Assertion Rule) and our assumption that (Stability) is violated, our verdict must be that u would be a felicitous contribution to joint inquiry.²¹

Considering a case with the required structure makes it clear that this prediction is false. We can modify [FIRE] to provide a familiar example: suppose it is common ground between Smith and Jones not only that either Ben or Bill is in the room but also that whoever is in the room is on fire. Assume, as before, that there are distinct points w', w'' in CS_w^c immediately after Jones's utterance u such that in w' , u expresses a content which is true just in case Bill is on fire, and in

²¹In principle, one could reject (Stability) while banning such cases by introducing some further principle to the effect that (Stability) must be satisfied whenever every point in the context set is a member of $F(u)$. But I can see no way of motivating such a principle. Indeed, there are positive reasons to think it is false. Suppose one rejected (Stability) in [FIRE]. If we hold the common ground of that case fixed but change Jones's utterance u to 'He is on fire or the person in the room is not on fire', then (given appropriate assumptions about the contents thereby asserted at different points) (i) every point in the context set is a member of $F(u)$, and (ii) at points where Bill is in the room, accepting u would rule out points in $F(u)$ at which Ben is in the room and on fire, whereas at points where Ben is in the room, accepting u would rule out points in $F(u)$ at which Bill is in the room and on fire, so that in either case, we would have a violation of (Stability). So selectively enforcing (Stability) is not as simple a strategy as it may at first seem. Thanks to an anonymous referee for bringing this sort of argument to my attention.

w'' , u expresses a content which is true just in case Ben is on fire. Then it is presupposed that Jones's utterance of 'He is on fire', where 'he' is intended to pick out the person in room, will express a truth. If (Stability) could be violated in this case, then Jones's utterance could nonetheless eliminate some point in the context set and thereby inform Smith about some feature of the world. But this is not the intuitive judgment about the case — without straining, one can only imagine Smith rejecting Jones's utterance as infelicitous ('Duh!' / 'That's *so* helpful.' / 'Who, damn it?!'). The most straightforward explanation for this fact is that (Stability) holds.²²

Similarly, suppose (Contraction) were violated. Then, given the connection between the context set and admissible linguistic behavior, there would be cases in which it was presupposed that an utterance of $\lceil \neg\phi \rceil$ would express a truth, and yet asserting ϕ would be felicitous because it would narrow the context set without eliminating it entirely.²³ Let us again modify [FIRE], holding fixed our assumption about the assertoric contents of Jones's utterance at various points in the context set of her conversation with Smith, but this time modifying the common ground so that what is presupposed is that whoever is in the room is *not* on fire. If Jones were to assertorically utter 'He is on fire', her utterance

²²One might attempt to resist this argument by giving up our assumption about the assertoric contents of Jones's utterance at various points in the context set of her conversation with Smith. The idea would be that, whereas in our modified version of [FIRE] the assertoric content of Jones's utterance includes $(F(u) \cap CS_w^c)$ at both w' and w'' , so that (Stability) is satisfied, this is an idiosyncratic feature of the case. But this strategy cannot ultimately succeed: either one must hold that, as a rule, for any utterance u in conversation c at point w , $(F(u) \cap CS_w^c)$ is a subset of the assertoric content of u at w , in which case (Stability) is trivially true, or one must hold that this condition is not always satisfied, in which case it will be straightforward to construct a case showing that our intuitions support (Stability). Similar considerations count against an analogous strategy for resisting the argument for (Contraction) below.

²³Again, one could in principle render problem cases consistent with the falsity of (Contraction) by stipulation, but I will disregard this possibility.

would be infelicitous ('What are you talking about?').²⁴ The most straightforward explanation for this fact is that (Contraction) holds.

In addition to (Stability) and (Contraction), our second argument will make use of a lemma concerning the interaction between the function F and the negation operator. The proposition to be established is that, for any assertoric utterance u of sentence ϕ and corresponding assertoric utterance u' of its negation $\lceil \neg\phi \rceil$, conversation c , and point w , if any point in CS_w^c is not in $F(u)$, then it is in $F(u')$. This can be shown as follows: Suppose our point, w' , is not in $F(u)$. Then, by the definition of F , u , as performed at w' , does not express a truth. So w' is not an element of the content p expressed by u at w' . By (Negation), u' expresses W/p at w' . So w' is an element of the content expressed by u' at w' . So u' as uttered in w' expresses a truth at w' . It follows that $w' \in F(u')$.²⁵

(Stability) and (Contraction) enable us to give an argument for the following

²⁴Andy Egan points out that it is possible to imagine Smith taking Jones's utterance as evidence that one of them has been mistaken about the context set of the conversation and then adjusting accordingly. But this kind of case is one in which it is a live possibility for Smith that the context set is defective; where Smith is certain that the context set is nondefective, Jones's utterance does seem to lack a felicitous literal interpretation. The easiest felicitous interpretation of the utterance to recover is metaphorical or idiomatic.

²⁵This argument tacitly appeals to an additional principle, namely:

(Propositionality): For any assertoric utterance u , conversation c , and point w , if w is in CS_w^c , then there is a content p such that u expresses p at w .

(Propositionality) is in tension with treatments of presupposition failure in terms of undefinedness, since such treatments explain the infelicity of a speaker asserting a sentence carrying a presupposition inconsistent with the common ground of her conversation in terms of its failure to express a content at any point in the context set. Those sympathetic to such treatments of presupposition failure need not be overly concerned, however, as my endorsement of (Propositionality) is a matter of convenience rather than necessity. My second argument for (Uniformity) could be reconstructed in an undefinedness-friendly way by adding to (Stability) and (Contraction) the further principle that applying (Assertion Rule) to the content of an utterance at a point results in eliminating from the context set of the conversation at that point every point at which the utterance does not express a content. This principle is defensible on much the same grounds as (Contraction); if it were false, one would expect the existence of felicitous utterances carrying presuppositions inconsistent with the conversational common ground.

principle:

(Diagonal): For any conversation c , assertoric utterance u , and point w , the CS_w^c -content of u in w is the intersection of CS_w^c with $F(u)$.

The argument is as follows:

Suppose we have a violation of (Diagonal): a case in which there is an assertion u of a sentence ϕ at a point w in a conversation c such that the CS_w^c -content of u in w is not the intersection of CS_w^c with $F(u)$ (call this set of points σ). Then the CS_w^c -content of u in w is either a proper subset of σ , or it contains points not in σ . If the former, then, by (Assertion Rule), accepting u in w would result in eliminating some points in $F(u)$ from CS_w^c in violation of (Stability). If the latter, then by our lemma, the CS_w^c -content of u in w contains at least one point in $F(u')$, where u' is a corresponding assertoric utterance of the sentence $\lceil \neg\phi \rceil$, and accepting u in w would fail to rule out this point, violating (Contraction).

So there can be no violation of (Diagonal). \square

For any assertoric utterance u , let us call the set $F(u)$ the *assertoric superdiagonal content of u* . Then the second argument for (Uniformity) establishes that the content expressed by an assertoric utterance at a point always agrees with the assertoric superdiagonal content of that utterance over the context set at that point. Though this conclusion does not entail the further claim that the assertoric content of an utterance at any point *is* its assertoric superdiagonal content, it strongly suggests it. There is a compelling inference to the best explanation from (Diagonal) to this stronger proposition; in the absence of countervailing evidence, we

ought to accept that the assertoric content of an utterance at any point is its assertoric superdiagonal content. But if this latter claim is true, (Uniformity) follows immediately.²⁶

Note that the thesis that the assertoric content of any utterance is its assertoric superdiagonal content is a constraint on a theory of assertoric content, not a theory of assertoric content in its own right. Though it guarantees the truth of (Uniformity), it is compatible with a wide range of total assignments of assertoric contents to utterances at points — including ones according to which assertoric content bears no discernible relation to the grammar of a language. Nevertheless, as we will see in the next section, the identification of assertoric content with superdiagonal assertoric content suggests a particular substantive theory of assertoric content.

1.6 Uniformity and Introspection

In addition to challenging Stalnaker's argument from (Assertion Rule) to (Uniformity), Hawthorne and Magidor (2009) offer a direct argument against the latter. They show that, given the Stalnakerian assumption that an assertoric utterance of a sentence at a point expresses that sentence's semantic content unless this would result in a violation of (Uniformity), and given that there can be failures of positive and negative introspection for presupposition, there will be violations of (Uniformity). Stalnaker's response (2009) is to defend positive and negative introspection. The conclusion of section 5, however, suggests an alternative way

²⁶Thus my second argument for (Uniformity), unlike my first, establishes the strong conclusion that no utterance expresses different contents at different points in the context set in addition to the weaker conclusion (which I labeled (Uniformity) in Section 2) that no utterance expresses different CS-contents at different points in the context set.

for the Stalnakerian to respond to this argument: to hold that an utterance expresses the same assertoric content at every point where it expresses any content at all. Choosing this response allows one to concede that there can be violations of both positive and negative introspection without having to concede that there are corresponding violations of (Uniformity).

We have seen that the claim that an utterance's assertoric content is its assertoric superdiagonal content is not yet a substantive theory of assertoric content. But the Stalnakerian has a substantive theory meeting this constraint to hand. If we can help ourselves to the notion of the grammatically determined content of an utterance u at a point w , then we can define the *semantic superdiagonal content* of u as the set of points w such that the grammatically determined content of u at w is true at w .²⁷ If we take the assertoric content of an utterance at any point to be its semantic superdiagonal content, the resulting theory will uniquely determine the assertoric content of every utterance while also satisfying the constraint that the assertoric content of an utterance be identical with its assertoric superdiagonal content.

Hawthorne and Magidor (2009) anticipate this superdiagonal proposal and offer two arguments against it. First, they argue that it fails to respect the intuitive connection between a sentence and its subject matter:

“In so far as we are permitted any intuitive fix at all on the asserted content, the proposal will lead to wildly unintuitive consequence[s]. If I assert ‘He is on fire’ there is some possible world where the sounds I produce have the semantic content that there are three goldfish in the room. According to the

²⁷As we will see shortly, there is a sense in which even the semantic superdiagonal content of an utterance is “grammatically determined”. The sort of grammatically determined content in play when semantic superdiagonal content is being defined is the *horizontal* content of an utterance at a point (see below).

superdiagonal proposal, my asserted content has no more to do with fire than it has to do with goldfish. Sure enough, the proposal can explain why I get fire-theoretic information—namely by intersecting the assertion with the context set. Nevertheless, our intuitive sense that the assertion itself has something to do with fire rather than goldfish is blatantly violated.” (2009, 389)

But the observation that the superdiagonal proposal assigns a content to ‘He is on fire’ that “has no more to do with fire than it has to do with goldfish” until the context set is considered need not trouble the theoretician concerned with modeling inquiry. A theory designed to capture the contribution of an utterance to inquiry should recognize that there are certain circumstances, radically unlike actual circumstances though they may be, under which the sentence ‘He is on fire’ *would* convey information about goldfish. This is a manifestation of the fact that the pairing between syntactic structures and meanings effected by any language is essentially arbitrary, being fixed by convention.

Of course, there is a sense in which the content of the sentence ‘He is on fire’ has to do with fire and not goldfish. Given the grammar of English, that sentence is used to convey information about fire and not about goldfish. A theory of content primarily concerned with capturing facts about the grammar of English would be mistaken if it paired ‘He is on fire’ with a fishy content. But a theory of the role of assertion in inquiry need not be (and, if my arguments are sound, cannot be) such a theory.

This brings us to Hawthorne and Magidor’s second argument against the superdiagonal proposal:

“Superdiagonal content is constructed out of semantic content. It thus assumes the reality of semantic content. But on the current proposal, semantic content is never asserted. It is surely rather troubling to insist on the reality

of semantic content but detach it altogether from the facts about what is asserted.” (2009, 389)

This argument hinges on the idea that associating each assertoric utterance with its semantic superdiagonal content in one’s model of inquiry precludes assigning semantic content an appropriate role in one’s broader theory of communication and linguistic competence. But to claim that there is reason to embrace a (Uniformity)-satisfying notion of assertoric content from within the possibility-carving framework is not to claim that there are not other, equally important notions of the contents of assertoric speech acts which do not satisfy (Uniformity). In fact, it is clear that there are. A brief reflection on the theoretical role of one such notion of content will clarify why Hawthorne and Magidor’s second argument against the superdiagonal proposal is not decisive.

The notion of content I have in mind is familiar: contents are paired with the sentences of a language by linguistic rules; mastery of these rules constitutes (or is at least a necessary condition on) speakers’ knowledge of that language. Following Hawthorne and Magidor (2009, 379), let us use the term *semantic content* to describe this species of content “determined by the standard compositional mechanisms of the language”. It is not obvious that the semantic content of an utterance is, on this construal, a form of content (in the sense of a set of points). Following Lewis (1980), we might instead regard the entity delivered by the compositional semantics when it is given an utterance as a function from ordered pairs of CP-contexts and indices to truth values, where CP-contexts are points centered on individuals and times and indices are ordered n -tuples containing more than simply a point parameter.²⁸

²⁸Though see King (2003), Cappelen and Hawthorne (2009), and Glanzberg (2011) for criticisms of this picture.

Even if standard compositional mechanisms fail to deliver genuine contents for utterances, however, such contents can easily be recovered from them. In fact, this can be done in multiple ways: for a given utterance u of sentence ϕ at a point w , we can recover the *horizontal content* expressed by u at w by building a CP-context out of w together with the speaker and time of u at w , feeding this CP-context into the semantic content of ϕ at w , using it again to fill in values for all the parameters in the index of the semantic content of ϕ except the point parameter, then abstracting over the point parameter. Alternatively, we can recover the semantic superdiagonal content expressed by u by: for each point w at which u expresses a content, using u to construct both a CP-context and its corresponding index, then feeding each such CP-context/index pair into the semantic content of ϕ at w , and finally including w in the semantic superdiagonal content just in case the result of this procedure is the True. It is the notion of Lewisian semantic content (or perhaps the related notion of horizontal utterance content) which Hawthorne and Magidor seem to have in mind when they argue that something has gone wrong if ‘He is on fire’ has no more to do with fire than it has to do with goldfish.

The idea that there are rules pairing syntactic forms (and thus, indirectly, utterances instantiating such forms) with semantic contents is of central importance to the project of explaining human linguistic competence. But a theory developed with the goal of capturing these rules need not assign sentences or utterances the same contents as the possibility-carving theory articulated above. It is plausible, for example, that the linguistic rule used to determine the referent of the demonstrative ‘that’ in context does not appeal to the context set at all; perhaps it takes into consideration only facts about the speaker’s intentions and gestures, or perhaps it takes into consideration only some abstract salience ranking among

possible referents, or perhaps it takes into consideration both of these as well as the capacities of an informed ideal audience — on any of these views, there can be violations of an analogue of (Uniformity) for horizontal utterance content when interlocutors are ignorant about the factors relevant to securing the referent of the demonstrative.

Indeed, as Hawthorne and Magidor emphasize, positing a notion of content which is tied to linguistic rules and does not satisfy (Uniformity) allows the theoretician to tell a story about how speakers are able to assign a (Uniformity)-satisfying variety of content to utterances — they do so by considering what horizontal content the semantic rules would assign to an utterance at each point in the context set, then treat the sentence as expressing the content that eliminates every point where the rules assign to the utterance a content which is false at that point. This is Stalnaker's process of diagonalization; Hawthorne and Magidor's idea is that it makes little sense to speak of superdiagonal content unless each utterance is associated with a horizontal content at each point in the context set. It is not clear that this idea is correct, since, if indices are suitably rich, both the semantic superdiagonal and the horizontal content of an utterance can be defined directly from its semantic content. Nevertheless, I think it should be conceded that the notion of horizontal utterance content is of independent theoretical interest.

But one can recognize the existence and theoretical interest of semantic content and horizontal utterance content without identifying them with assertoric content. In other words, the following three propositions are perfectly consistent, both logically and in terms of fitting naturally into a coherent theoretical outlook: (i) the assertoric content of an utterance is always its semantic superdiagonal content; (ii) explaining linguistic competence requires appealing to both a notion of semantic content and a notion of horizontal utterance content; (iii) the notion of

horizontal utterance content can be used to define the notion of the semantic superdiagonal content expressed by an utterance. So Hawthorne and Magidor's second argument against the superdiagonal proposal is not successful.²⁹

²⁹Thanks are due to Andy Egan, Simon Goldstein, John Hawthorne, John Mackay, Ofra Magidor, Taylor-Grey Miller, Meagan Phillips, and a number of anonymous referees for helpful comments on earlier drafts of this paper, as well as to audiences at Northern Illinois University, the University of Texas at Austin, the 29th European Summer School in Logic, Language, and Information, and the Central Division Meeting of the American Philosophical Association for illuminating discussion.

2 Do Not Diagonalize

2.1 Introduction

A woman is speaking in the next room. Smith turns to Jones and says, ‘That is either Zsa Zsa Gabor or Elizabeth Anscombe.’ The case is a familiar one in the literature on assertion, and so is the theoretical picture it has been taken to motivate.³⁰ For it seems that what Smith asserts is neither necessarily true nor necessarily false.³¹ And yet, if the woman in the next room is Gabor or Anscombe, there is pressure to think that Smith asserts the necessarily true proposition that that woman is either herself or someone else; correspondingly, if the woman in the next room is neither Gabor nor Anscombe, there is pressure to think that Smith asserts the necessarily false proposition that that woman is not herself. The familiar solution is to hold that Smith asserts neither the necessarily true proposition nor the necessarily false proposition, but rather a third contingent proposition, which is true if the woman in the next room is Gabor or Anscombe and false otherwise.³²

³⁰For the case (which appears here slightly modified), see Stalnaker (1978, p. 91). For the theoretical picture, see e.g. Stalnaker (1978, 2009, 2014).

³¹For the purposes of my argument, I assume in what follows that the objects of the speech act of assertion are propositions, referring to the propositional object of an assertion as its *assertoric content*. What is a proposition? I wish to take no stand on this issue. For simplicity of exposition, in what follows I treat propositions as sets of metaphysically possible worlds. Nevertheless, everything I claim is compatible with views of propositions which individuate them more finely. For even on such views, a given proposition will induce a partition on the set of possible worlds: there will be those in which the proposition is true, and those in which it is not. More generally, whatever theory of propositions one is inclined to endorse, one will need to give some account of their informational properties — that is, the way in which they exclude some possibilities but not others. My choice to theorize about propositions as sets of metaphysically possible worlds is a choice to focus on their informational properties. I take this choice to be consistent with a wide range of views about the nature of propositions.

³²Stalnaker’s case, though convenient as a starting point for our discussion because of its familiarity, has the disadvantage that it loses much of its intuitive force if we adopt a non-Millian

This third proposition has not been chosen arbitrarily. It has a special status: it is the proposition which characterizes the information Jones, ignorant though she is concerning the identity of the woman in the next room, can learn from Smith's utterance by assuming that the proposition it semantically determines is true.³³ For Jones is in a position to reason as follows: either the woman in the next room is Zsa Zsa Gabor or Elizabeth Anscombe, or the woman in the next room is neither Zsa Zsa Gabor nor Elizabeth Anscombe. If the woman in the next room is Zsa Zsa Gabor or Elizabeth Anscombe, then Smith's utterance semantically determines a necessary truth. If, on the other hand, the woman in the next room is neither Zsa Zsa Gabor nor Elizabeth Anscombe, then Smith's utterance semantically determines a necessary falsehood. So, assuming that the proposition Smith's utterance semantically determines is true, the woman in the next room must be either Zsa Zsa Gabor or Elizabeth Anscombe. Holding that Smith asserts the proposition which characterizes the information Jones is in a position to recover from her utterance fits naturally with the intuitive idea that our practice of asserting is intimately connected with our practice of communicating information.

The familiar solution, according to which Smith asserts this third proposition, has an obvious and theoretically appealing generalization: that speakers systematically assert propositions which characterize the information their interlocutors can learn from their utterances by assuming that the propositions they semantically determine are true.³⁴ According to such a view of assertion, the process

semantics for proper names. It is worth noting, then, that the same idea can just as easily be motivated using examples not containing proper names.

³³For a more precise characterization of the notion of the proposition semantically determined by an utterance, see section 2 below.

³⁴Importantly, this sort of view makes a distinction between the status of genuine interlocutors and the status of others, like eavesdroppers, who might happen to recover information from an

of recovering information from an utterance despite ignorance about contextual factors like the individuals referred to by demonstratives is an integral part of the explanation for successful communicative exchanges between fully rational interlocutors. For reasons which will become clear in the next section, let us call this generalization *the diagonal theory of assertoric content*, or, for the sake of brevity, **(Diagonalism)**.

(Diagonalism) is a claim about which propositions speakers assert, not a claim about the semantics of any natural language. The relationship between semantic theories and theories of assertoric content (which belong to the domain of speech act theory) is open to debate, and two theorists who agree about the grammar of English might nonetheless disagree about the truth of (Diagonalism). Moreover, the connection which (Diagonalism) enforces between the assertoric content of an utterance and the information it communicates to members of its audience is by no means inevitable. It could instead be that the information an utterance communicates is systematically related to, though not identical with, its assertoric content. Indeed, I will ultimately argue that such a picture is preferable to (Diagonalism). Nevertheless, it should be emphasized that (Diagonalism) is an appealing theory of assertoric content. On the one hand, in identifying the content of an assertion with the information audience members are in a position to recover from it, (Diagonalism) provides the most direct possible explanation of the platitude that assertion is a means of communicating information. On the other hand, there can be little doubt that speakers do often reason in the way we have imagined Jones reasoning about Smith's utterance.

utterance. The propositional content of an assertion is held to be intimately connected only with the information interlocutors can recover from an utterance, not with the information others can recover from it.

In a well-known example of the ubiquity of this phenomenon, John Perry (1986) asks us to imagine that he receives a postcard in the mail with no signature and no indication of the location from which it was posted. On the postcard are the words ‘I am having a good time’. In this situation, Perry observes, he is not in a position to recover any proposition of the form ‘ P is having a good time at t ’, though it is plausible that the proposition semantically determined by the sentence on the card, as inscribed by its author, is of this form. He can, however, recover the information that someone inscribed ‘I am having a good time’ on the postcard at some time t prior to its arrival in his mailbox, and that (if the author of the inscription was telling the truth), he or she was having a good time at t . Even those who ultimately reject (Diagonalism), then, must acknowledge the intuitive appeal of holding, on the basis of the familiarity of this sort of reasoning, that it should play a central role in our theory of assertion.

Notwithstanding its theoretical appeal, recent work on assertion has reached no clear consensus concerning (Diagonalism). Lewis (1980) is sensitive to the distinction between (Diagonalism) and competing theories of assertoric content, but does not take a stand concerning whether (Diagonalism) is correct. Rabern (2012) assumes that (Diagonalism) is false. Stojnić (2017) extends an argument against (Diagonalism) due to Soames (2002). Stalnaker (1978, 2014), meanwhile, argues in favor of (Diagonalism). His arguments are criticized by Hawthorne and Magidor (2009, 2010) and defended by Almotahari and Glick (2010).

My first goal in what follows is to show that (Diagonalism) systematically conflicts with our intuitive judgments about the conditions under which the contents of speakers’ assertions would be true or false. I believe this result eliminates (Diagonalism) as a tenable theory of assertoric content.

But there is more to be said: (Diagonalism) emerges naturally from an intuitive picture of assertion according to which assertoric content plays a certain privileged role in the theory of communication, and its failure as a theory of assertoric content gives us reason to reassess the relationship between assertion and communication. My second and third goals in what follows, then, are to show exactly how (Diagonalism) emerges from this intuitive picture and to suggest an alternative way of thinking about assertion which does not tempt us to embrace (Diagonalism). The alternative account of assertion I favor takes as its point of departure the platitude that speakers undertake commitments when they assert. On this picture of assertion, the content of an assertoric utterance is the proposition to which a speaker undertakes a distinctive sort of commitment in producing that utterance. Because it relaxes the connection between assertoric content and the theory of communication, my committal account of assertion is not subject to the problems which plague its competitor. In fact, given an independently plausible method for picking out the propositions to which speakers undertake commitments in asserting, it can be shown to make correct predictions about the cases which prove so problematic for (Diagonalism).

After characterizing (Diagonalism) more precisely in section 2, I present my argument against it in section 3 and respond to potential objections in section 4. Section 5 shows how (Diagonalism) follows from a number of plausible premises about the relationship between assertion and communication. Section 6 motivates and develops my committal account of assertion. Section 7 concludes.

2.2 (Horizontalism) and (Diagonalism)

The semantics of English is *Kaplanian* in the sense that the context in which a sentence is uttered determines the referents of any demonstratives, deictic pronouns, and similar expressions it contains, thereby partially determining what semantic content it expresses. Since the context sensitivity of subsentential expressions will concern us only in so far as it gives rise to the context sensitivity of full sentences, nothing will be lost if we think of this Kaplanian picture sententially, forgetting for the nonce how the contents of sentences systematically depend on the contents of their constituents: a sentence, uttered in a context, expresses a proposition; which proposition a sentence expresses is often determined by the context in which it is uttered. Let us call a way of pairing sentences in contexts with propositions a *grammar*.³⁵

If we have a grammar representing the English language, we can use it to associate any given utterance of a sentence at a world with the proposition obtained by using the context in which that utterance is produced at that world to fill in the semantic values of its context-sensitive vocabulary “in the familiar way”. Let us call this proposition — intuitively, the proposition semantically determined by the utterance at that world — the *horizontal proposition* of the utterance at the world.³⁶ In the case of Smith’s utterance concerning Gabor and Anscombe, the pressure to think that if Gabor is in the room then the content of Smith’s assertion is necessarily true arises from the fact that if Gabor is in the room then the hor-

³⁵Formally, I take a grammar to be a function from ordered sentence/context/world triples to truth-values.

³⁶Formally, I take the horizontal proposition of assertoric utterance u of sentence s (of some language with grammar G) in world w to be the set of worlds w' such that $G(\langle s, c_w, w' \rangle) = 1$, where c_w is the context in which u is produced at w .

horizontal proposition of Smith’s utterance is necessarily true (it is the proposition that Gabor is either Gabor or Anscombe). Corresponding to the notion of the horizontal proposition of an assertoric utterance at a world is a theory of assertoric content: the theory which holds that the content of an assertoric utterance at a world is always the horizontal proposition of that utterance at that world. Let us call this theory *the horizontal theory of assertoric content*, or, for the sake of brevity, **(Horizontalism)**.

(Horizontalism) is one systematic way of pairing assertions with propositional contents. As we have seen, however, it is not the only way: there is also (Diagonalism), which promises to dissolve the puzzle of the Gabor/Anscombe case by identifying what is asserted with the information Jones is in a position to recover from Smith’s utterance by assuming that it expresses a truth. I have suggested that this promise is empty: (Diagonalism) is false. Before I can present an argument for this conclusion, however, it is necessary to formulate (Diagonalism) more precisely. In order to do this, we require the formal tools to model how the horizontal proposition of an utterance depends on the world in which it is produced.

Consider an assertoric utterance (call it u) of a simple subject–predicate sentence containing a deictic pronoun — ‘She speaks Farsi’, for example. We can generate a matrix from u by creating one row and one column corresponding to each possible world which contains u and, for each ordered pair $\langle w, w' \rangle$ of such worlds, writing ‘T’ in the location where row w meets column w' if the assertoric content of u as uttered in w is true at w' and ‘F’ otherwise.³⁷ Call such a matrix a

³⁷Following Stalnaker (2004), I exclude from consideration possible worlds which do not contain the utterance. Since the proposition that the utterance occurs is always part of the common ground of the conversation, such worlds are never relevant to the dynamics of the context set. (See below for definitions of *common ground* and *context set*.)

propositional concept.

A propositional concept represents the way in which the the world in which an utterance takes place affects the proposition with which it is paired by our theory of assertion. The entries along the *row* corresponding to a given world characterize the proposition the utterance expresses at that world. The entries down the *column* corresponding to a given world record whether the propositions the utterance expresses at different worlds are true at that world.

\mathfrak{A}	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>a</i>	T	F	T	F
<i>b</i>	T	F	T	F
<i>c</i>	F	T	F	T
<i>d</i>	F	T	F	T

(*a*: Speaker demonstrates Smith; Smith speaks Farsi; *b*: Speaker demonstrates Smith; Jones speaks Farsi; *c*: Speaker demonstrates Jones; Smith speaks Farsi; *d*: Speaker demonstrates Jones; Jones speaks Farsi.)

For example, \mathfrak{A} is the propositional concept predicted by (Horizontalism) to correspond to our assertoric utterance of ‘She speaks Farsi’. For the sake of simplicity of exposition, we restrict our attention in \mathfrak{A} to a set of four worlds: we assume that the speaker is demonstrating either Smith or Jones, and that exactly one of the two speaks Farsi (but we make no assumption about which). In worlds *a* and *b*, the speaker is demonstrating Smith; in worlds *c* and *d*, she is demonstrating Jones. In worlds *a* and *c*, Smith but not Jones speaks Farsi; in worlds *b* and *d*, Jones but not Smith speaks Farsi.

As \mathfrak{A} shows, the horizontal proposition of ‘She speaks Farsi’ in *a* is the proposition that Smith speaks Farsi, whereas in *c* it is the proposition that Jones speaks Farsi. This corresponds to the fact that in *a* the utterance is produced while the

speaker is demonstrating Smith, whereas in c it is produced while the speaker is demonstrating Jones.

With the idea of a propositional concept in hand, we can now give (Diagonalism) a precise formulation. The information interlocutors are able to recover from an utterance by assuming that the proposition it semantically determines is true can be read off of the propositional concept assigned to it by (Horizontalism) in a straightforward way: it is the proposition uniquely characterized by the entries falling along the propositional concept's diagonal. Let us call this proposition the *superdiagonal proposition* of the utterance in question.³⁸ One option for defining (Diagonalism), then, would be to identify it with the thesis that speakers always assert the superdiagonal propositions of their utterances. There are reasons for thinking that this proposal is too restrictive, however, and I will pursue a different approach in what follows.³⁹

A more permissive way of understanding (Diagonalism) appeals to the *com-*

³⁸Formally, the superdiagonal proposition of an assertoric utterance u at world w is the proposition which is true at a world w' just in case the horizontal proposition of u in w' is true at w' .

³⁹The proposal that the assertoric content of an utterance is always its superdiagonal proposition is subject to a number of difficulties. Stojnić (2017, p. 175), for example, argues that "...if Jason says 'I am here now' and I follow up with 'That's necessarily true', or 'What Jason said is necessarily true', intuitively what I said is false." Assuming that 'that' and 'what Jason said' in the relevant sentences refer to what Jason asserts, the intuitive falsity of Stojnić's examples is difficult to reconcile with the claim that the assertoric content Jason's utterance is its superdiagonal proposition, since the superdiagonal proposition of an utterance of 'I am here now' is true at every world where the utterance exists. It is also difficult to make sense of the English counterfactual optative construction 'Would that ϕ !' on the superdiagonal proposal if we assume that, as uttered in a context, it expresses a speaker's desire that the assertoric content of ϕ as uttered in that context were true (though it is false). An utterance of 'Would that he were here!' seems to express a desire that some particular individual be present — that is, a desire which is satisfied in just those worlds where that individual is present — rather than a desire that whomever the speaker has in mind be present — that is, a desire which is satisfied in those worlds where some individual is both present and the referent of the speaker's 'he', regardless of whether that individual is the one the speaker has in mind at the world of utterance. Thanks to John Hawthorne for bringing this second problem with the superdiagonal proposal to my attention.

mon ground of a conversation at a world — that is, the set of propositions which interlocutors take for granted for the purposes of that conversation at that world.⁴⁰ The common ground of a conversation at a world determines a set of worlds — those which are compatible with every proposition in the common ground at that world, or, equivalently, those which are treated as live possibilities for the purposes of that conversation at that world. Let us call this set the *context set* of the conversation at that world. Since interlocutors have already ruled out for the purposes of the conversation all worlds not in the context set, we will not go wrong if we restrict our attention to the context set in characterizing the informal notion of the information they can recover from an utterance by assuming that it expresses a truth.

With this restriction in mind, let us refer to any proposition which agrees in truth-value with the superdiagonal proposition of an utterance over the context set of the conversation in which that utterance occurs as a *diagonal proposition* of that utterance. Each utterance has many diagonal propositions at a given world; these are equivalent over the context set at that world but diverge elsewhere. We can then understand (Diagonalism) as the claim that a speaker always asserts some diagonal proposition of her utterance.⁴¹ \mathfrak{B} below is the propositional concept predicted by (Diagonalism) to correspond to the assertoric utterance of ‘She speaks Farsi’ introduced above. Note that the entries along the diagonal of \mathfrak{A} (T

⁴⁰It is standard to define a more stringent notion of the common ground which is both iterated and collectivized (see, for example, Stalnaker (2014, pp. 4; 25)). Adopting such a stringent conception would not substantively affect my arguments in what follows, so I have opted for the simpler definition above.

⁴¹Formally, I take (Diagonalism) to be the thesis that, for any assertoric utterance u performed in conversation c at world w , if cg is the common ground of c at w and sd is the superdiagonal proposition of u , then the assertoric content of u at w is some p such that $p \cap cg = sd \cap cg$.

F F T) appear in each horizontal row of \mathfrak{B} .^{42/43}

\mathfrak{A}	a	b	c	d
a	T	F	T	F
b	T	F	T	F
c	F	T	F	T
d	F	T	F	T

\mathfrak{B}	a	b	c	d
a	T	F	F	T
b	T	F	F	T
c	T	F	F	T
d	T	F	F	T

(a : Speaker demonstrates Smith; Smith speaks Farsi; b : Speaker demonstrates Smith; Jones speaks Farsi; c : Speaker demonstrates Jones; Smith speaks Farsi; d : Speaker demonstrates Jones; Jones speaks Farsi.)

2.3 The Case Against (Diagonalism)

In a variety of circumstances, (Diagonalism) conflicts with the intuitive connection between the theory of assertoric content and our judgments about the conditions under which the contents of speakers' assertions would be true or false.⁴⁴ I will

⁴²(Diagonalism) is thus a *one-dimensional* theory of assertoric content in the sense that it does not permit the assertoric content of an utterance to vary across worlds in the context set of a conversation.

⁴³A methodological note: Some readers may worry that my appeal to notions like *propositional concept* and *context set* in my exposition of (Diagonalism) smuggles in certain assumptions about how to model the contents of assertoric utterances and propositional attitudes, and that the smuggling in of these assumptions undermines the interest of (Diagonalism) as a theory of assertoric content. In response to this sort of worry, let me comment briefly on what the framework in which I have chosen to formulate my arguments presupposes. First, propositional concepts are simply tools for visualizing the two ways in which possible worlds are commonly held to be relevant to the truth-values of utterances: on the one hand, in so far as the world at which an utterance is produced determines which proposition it expresses by resolving the semantic values of context-sensitive vocabulary; on the other hand, in so far as propositions differ in truth-value between possible worlds. Second, the idea that every conversation is associated with a set of possible worlds compatible with every presupposed proposition is simply one way of making formally tractable the platitude that there are some scenarios which are not treated as live possibilities during normal conversational exchanges (the possibility that one is radically mistaken about the meanings of words in one's own language, for example). The arguments in what follows could likely be recast in a framework which made this platitude formally tractable in a different way.

⁴⁴Importantly, I regard (Diagonalism) as a substantive rather than a definitional thesis about assertion. A theory of assertion is a development of our pre-theoretical beliefs about the social

introduce this conflict by appealing to a modified version of a case from King (2018) involving ‘ready’, a relational expression with an implicit argument, but the same point can be made without reference to such expressions.

Relational expressions with implicit arguments belong to a broader class of expressions which King calls *supplementives*: “contextually sensitive expressions whose context independent meanings do not by themselves suffice to secure semantic values for those expressions in contexts” (2018, p. 631). What exactly does a context supply which, in conjunction with the context-independent meaning of a supplementive, suffices to determine its semantic value in context? Borrowing terminology from King (2014), let us refer to this question as the question of the *metasemantics* of supplementives. In what follows, I will write as if I take for granted a *subjective intentionalist* metasemantics, according to which the semantic values of supplementives — and thus, to return to the sententialist perspective of the previous section, the propositions semantically associated with sentences containing them — are determined by what the speakers who utter them have in mind (i.e. intend to semantically contribute) at the time of utterance.⁴⁵ My argument against (Diagonalism), however, does not presuppose any particular metasemantics — all that it requires is that the features of the context which do the work of determining the semantic values of supplementives, whatever these

practice of making claims. As such, it is accountable to those pre-theoretical beliefs. It may turn out that there are compelling theoretical reasons to regard our intuitive judgments about assertion as only approximating the truth — but no theory of a social practice can float wholly free of untutored intuition. We must evaluate (Diagonalism) in the light of our everyday judgments about who has claimed what, and about whether what has been claimed is or would be true, and we must reject it if it cannot accommodate those judgments.

⁴⁵The intentions I have in mind should be distinguished from *communicative* intentions. A speaker can have a particular object in mind when uttering ‘That is my favorite object’ without intending to bring it about that members of her audience recognize which object she has in mind — thus, she can have a recognizably semantic intention without having a corresponding communicative intention.

may be, can be known to the speaker but not her audience.⁴⁶

What is interesting for our purposes, then, is what happens when we introduce a known asymmetry in the information about the context available to the speaker and her audience, so that all parties know that the speaker possesses information which uniquely determines the proposition semantically associated with her utterance, though her audience lacks this information. Consider the following case (modified from one in King (2018)):

[READY]: Susie and Matt are planning to go to lunch with baby Molly. The day is gloomy, but there is a chance that the weather will improve during their meal and they will have an opportunity to take a walk in the park before driving home. They go through their various preparations. Fully dressed, and eyeing baby Molly in her car seat, Susie proudly says, 'Molly is ready'. Matt knows that Susie is either thinking only about lunch (perhaps because she thinks it improbable that the weather will improve) or about both lunch and the walk, though he does not know which. Matt's evidential situation is common ground between the two.

I take it that [READY] is a case of successful assertion, in the sense that there is some propositional content Susie asserts in uttering 'Molly is ready'. The claim that [READY] is a case of successful assertion is thus weaker than the claim that Susie's utterance is wholly felicitous, on which I do not wish to take a stand. In support of the suggestion that [READY] is a case of successful assertion, consider that the practical import of Matt's uncertainty about what Susie has in mind is

⁴⁶Those favoring alternative accounts of the metasemantics of supplementives are invited to replace my talk of *what speakers have in mind at the time of utterance* in what follows with material describing whatever factor they think is responsible for fixing the relevant semantic values.

quite minimal: at most it means that he cannot learn from her utterance whether Molly is ready for the walk. Indeed, it is easy to imagine that [READY] is a low-stakes communicative situation in the sense that Matt is significantly more concerned about the answer to the question of whether Molly is ready for lunch than about the answer to the question of whether she is ready for the walk, and that for this reason there is a tacit understanding between him and Susie to the effect that expending the effort required to make what she has in mind fully explicit is not required.

Two sorts of readers may have reservations about following me in declaring [READY] a case of successful assertion: those who believe that a speaker cannot succeed in securing a semantic value for a suppletive if the value she secures is not suitably apparent to her audience, and those who believe that a speaker cannot succeed in asserting a proposition unless members of her audience are in a position to identify the proposition she has asserted. I will address these worries in sections 4.2 and 4.4 below.

Assuming for the time being that [READY] is a case of successful assertion, then, we have a counterexample to (Diagonalism). Setting aside worlds at which Molly is (i) ready for neither lunch nor the walk or (ii) ready for the walk and not ready for lunch — that is, worlds at which both of the possible horizontal propositions of Susie's utterance are false — we can represent the remaining portion of the context set of Susie and Matt's conversation immediately after Susie's assertion using four worlds: one (*a*) in which Susie has only lunch in mind at the time of her assertion and Molly is ready only for lunch, one (*b*) in which Susie has only lunch in mind at the time of her assertion and Molly is ready for both lunch and the walk, one (*c*) in which Susie has both lunch and the walk in mind at the time of her assertion and Molly is ready only for lunch, and one (*d*) in which Susie has

both lunch and the walk in mind at the time of her assertion and Molly is ready for both lunch and the walk. According to (Horizontalism), the propositional concept of Susie's assertion, restricted to this portion of the context set of [READY], is thus \mathfrak{C} :

\mathfrak{C}	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>a</i>	T	T	T	T
<i>b</i>	T	T	T	T
<i>c</i>	F	T	F	T
<i>d</i>	F	T	F	T

(*a*: Susie has only lunch in mind; Molly is ready only for lunch; *b*: Susie has only lunch in mind; Molly is ready for both lunch and the walk; *c*: Susie has both lunch and the walk in mind; Molly is ready only for lunch; *d*: Susie has both lunch and the walk in mind; Molly is ready for both lunch and the walk.)

According to (Diagonalism), on the other hand, the content of Susie's utterance in [READY] is uniform across the context set and true at *a*, *b*, and *d* (but not *c*), as shown in \mathfrak{D} :

\mathfrak{D}	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>a</i>	T	T	F	T
<i>b</i>	T	T	F	T
<i>c</i>	T	T	F	T
<i>d</i>	T	T	F	T

(*a*: Susie has only lunch in mind; Molly is ready only for lunch; *b*: Susie has only lunch in mind; Molly is ready for both lunch and the walk; *c*: Susie has both lunch and the walk in mind; Molly is ready only for lunch; *d*: Susie has both lunch and the walk in mind; Molly is ready for both lunch and the walk.)

Consider now the following claim about [READY]:

(1) Regardless of what Susie has in mind at the time of her assertion, what she asserts in [READY] is consistent with Molly not being ready for the walk.

(1) is intuitively false: if Susie has both lunch and the walk in mind at the time of her assertion, then what she asserts is false if Molly is not ready for the walk. But any plausible account of the truth-conditions of (1) will predict that it is true if there is a world w in the context set of Susie and Matt's conversation such that (i) w is in the content of Susie's assertion regardless of what she has in mind at the time of her assertion and (ii) Molly is not ready for the walk at w . According to (Diagonalism), regardless of what Susie has in mind at the time of her assertion, she expresses a proposition which is true at world a , where Molly is not ready for the walk. So (Diagonalism) incorrectly predicts that (1) is true.

Again, suppose that Susie and Molly are in fact in world d , so that Susie has both lunch and the walk in mind at the time of her assertion and Molly is ready for both lunch and the walk (it follows that Susie's assertion is true in d). Let us refer to this further specified version of [READY] as [READY*]. Consider:

(2) What Susie asserts in [READY*] could have been true even if Molly had not been ready for the walk.

(3) What Susie asserts in [READY*] would have been true if Molly had been ready for lunch but *not* the walk and Susie had not had the walk in mind at the time of her assertion.

(2) and (3) are intuitively false: in [READY*], Susie has the walk in mind at the time of her assertion, so what she asserts is false at every world where Susie

is not ready for the walk. But if the content of Susie's assertion is its diagonal proposition, then (2) and (3) should be true. For the diagonal proposition of Susie's assertion is true at a , where Molly is not ready for the walk and Susie does not have the walk in mind at the time of her assertion. Since the diagonal proposition of Susie's assertion is true at a , it follows that it *could* have been true if Molly had not been ready for the walk (since a is a world where Molly is not ready for the walk) — thus sentence (2) should be true — and that it *would* have been true if Molly had been ready for lunch but not the walk and Susie had not had the walk in mind (since a is the only world in the context set where this is the case) — thus sentence (3) should be true.

So (Diagonalism) yields incorrect predictions about the truth-values of (1-3). If, as I have suggested, the theory of assertion must be held accountable to our judgments about whether what has been asserted is or would be true, it follows that (Diagonalism) must be rejected.

Note that this problem for (Diagonalism) cannot be traced to idiosyncratic features of any particular type of expression. It is easy to construct structurally analogous cases involving other supplementive expressions as well as proper names. When it comes to other supplementive expressions, a general recipe for constructing such cases is to stipulate that the speaker has an intention which is specific enough to determine a unique semantic value for the supplementive, her audience knows this but is uncertain what this intention is, and this informational asymmetry is part of the common ground. I present a case involving a proper name below, leaving it to the reader to construct analogues of (1-3) and verify that the intuitive judgments about their truth-values are inconsistent with (Diagonalism).

[JUSTIN]: Eli is trying to convince Mark that she has too much to do to go to

trivia night at the local bar. In the middle of a long list of incomplete to-do items, she says, 'I have to draft an email to Justin.' As it happens, there are exactly two people named 'Justin' in their social circle: J. Trudeau and J. Theroux. Mark knows that Eli has exactly one Justin in mind but does not know which, and his evidential situation is common ground between the two.

2.4 Objections and Replies

2.4.1 The Dialectical Skeptic

I have used intuitive judgments about sentences describing the truth-values of asserted contents in various circumstances in my argument against (Diagonalism). A certain kind of critic may seek to challenge this choice. According to such a critic, the theory of assertion is the theory of a speech act which *essentially* and *by stipulation* has a certain characteristic kind of effect on the context set — one correctly described by (Diagonalism). From this perspective, it may seem a mistake to think that our intuitive judgments are at all relevant to the theory of assertion and an impossibility that (Diagonalism) should turn out to be false.

Let us refer to the stipulatively defined speech act about which this sort of critic wishes to theorize as *assertion in the broad sense*. Though I do not think that many who have written about assertion have had assertion in the broad sense in mind, I do not wish to deny that the concept has theoretical interest; perhaps, for example, we can use it to give a unified treatment of the way in which possibilities are ruled out in genuine inquiry, hypothetical reasoning, and plan formation. But the possibility of theorizing about assertion in the broad sense by no means

precludes theorizing about *assertion*, our social practice of making claims. Indeed, if my argument so far is correct, we must conclude that (Diagonalism) is false, and therefore, if what is asserted in the broad sense must always be some diagonal proposition of an utterance, that assertion is not a species of assertion in the broad sense. Though it is rarely worth engaging in terminological disputes, this result suggests that some confusion could be avoided if assertion in the broad sense were renamed.⁴⁷

2.4.2 The Gricean

My argument depends on the claim that Susie’s assertion in [READY]/[READY*] is successful — that is, that there is some content she asserts. A certain kind of critic, influenced by the Gricean idea that communication is cooperative and involves mutual recognition of intentions between speaker and listener, might seek to deny this claim. Perhaps [READY] is not a case of successful assertion, since successful assertion requires that a speaker’s audience be in a position to identify which proposition was asserted.

⁴⁷It is sometimes suggested that Stalnaker, the most persistent advocate of (Diagonalism), uses the word ‘assertion’ to pick out assertion in the broad sense, so that arguments like the one I have offered above do not threaten his position. While it is true that Stalnaker does sometimes use the word in this fashion, he is aware that doing so is somewhat misleading. For example, he writes, “‘Assertion’ is probably too narrow a term for [assertion in the broad sense], since it has a connotation of seriousness not shared by all declarative speech acts with this [effect on the context set]” (Stalnaker 2018, p. 386). Moreover, he often writes as if he does *not* have assertion in the broad sense in mind: “I should emphasize that I am not claiming that one can *define* assertion in terms of a context-change rule, since that rule will govern speech acts that fall under a more generic concept. A full characterization of what an assertion is would also involve norms and commitments.” (Stalnaker 2014, p. 89). A natural interpretation, then, is that Stalnaker recognizes both assertion and assertion in the broad sense, and endorses (Diagonalism) as it pertains to both. Readers who regard this textual evidence as inconclusive will, I hope, be satisfied upon learning that Stalnaker has confirmed in conversation with me that he recognizes both assertion and assertion in the broad sense.

While it is undoubtedly true that an assertion can in certain cases be infelicitous (and even, we may assume, unsuccessful) because its audience cannot identify what the speaker has in mind — as when, out of the blue, I utter ‘He should have seen it coming!’ — many have thought that there are also *feliculously underspecified* uses of a broad range of supplementive expressions: cases in which such expressions are used felicitously (and thus successfully) even though context does not assign them a unique semantic value. King (2014) offers the following example of a case in which a supplementive appears to be felicitously underspecified (see also the examples in King (2018)):

[LOST WINDS]: Suppose I am surfing at Lost Winds beach. South of me are a bunch of other surfers. I keep seeing different surfers from the pack getting incredible waves. I comment to my buddy nodding to the south ‘Those guys are good.’ It seems doubtful that my intentions pick out a unique group of guys and it seems doubtful that my buddy takes them to. I think the reason is that for the purposes of our conversation, all that I am trying to or need to convey is that there are numerous good surfers among the group to the south. So here my intentions, and the fact that they would be recognized by an idealized hearer, determine merely a range of groups, and hence potential semantic values, for the demonstrative. (King 2014, pp. 112-113)

[LOST WINDS] is a case in which it is not clear how to identify a unique proposition corresponding to a speaker’s utterance. Nevertheless, it seems quite clearly to be a case of successful assertion. If this is so, it is difficult to see why a case like [READY], which differs from a case of felicitous underspecification only in that the speaker has an intention which fixes a unique semantic value for her supplementive, should be different.

Of course, the existence of felicitously underspecified uses of supplementives is contested; it is open to the Gricean to deny that [LOST WINDS] and related cases involve successful assertion (though this seems to me to be a serious price to pay

for maintaining her position). Fortunately, there is more to be said in response to the Gricean. First, if the Gricean were correct in claiming that [READY] is not a case of successful assertion, we would be left with no explanation for the intuitive falsehood of (1-3) and the intuitive truth of:

(4) What Susie asserts in [READY*] is true if Molly is ready for both lunch and the walk.

For if no content is asserted in [READY*], our prediction must be that the existence presuppositions generated by the pseudo-clefts in (1-3) and (4) are not satisfied, leading to a truth-value gap (and thus to infelicity).

Second, even if the Gricean were correct in claiming that Susie's assertion in [READY]/[READY*] is not successful, to construe this response as an objection to my position would be dialectically inappropriate in the following sense. My argument is an argument against (Diagonalism). In offering it, I presuppose that ruling out (Diagonalism) requires an argument; that is, that neither (Diagonalism) nor its negation is obviously false. But if (Diagonalism) is true, Susie and Matt *can* easily coordinate on the assertoric content of Susie's utterance in [READY], at least as it bears on the context set: it is a diagonal proposition of Susie's utterance, true in worlds *a*, *b*, and *d*, and false in world *c*. This means that, insofar as the Gricean thinks it is intuitively clear that [READY] is not a case of successful assertion, she must also think that it is intuitively clear that (Diagonalism) is false. But then the Gricean can at best charge me with believing what is true — namely, that (Diagonalism) should be rejected — for the wrong reasons.

2.4.3 The Methodological Skeptic

Our judgments about the truth-values of constructions of the form ‘ x said that ϕ ’ are notoriously context-sensitive. If constructions of the form ‘What S asserts in W (could/)would have been true (even) if Q ’ (henceforth *truth-conditional pseudo-clefts*) are context sensitive in the same way, this fact might be thought to cast doubt on my conclusions.

In fact, I believe there is little reason to doubt the theoretical utility of truth-conditional pseudo-clefts. For truth-conditional pseudo-clefts have two important features which mitigate the extent to which they are context-sensitive. First, they are pseudo-clefts. Since pseudo-cleft constructions carry an exclusiveness implicature, they do not share the tendency of ‘said-that’ constructions to pick up on proper parts of an utterance, as when a speaker utters a conjunction and I felicitously report her as having said the first conjunct.⁴⁸ Second, they are modal; this makes them more sensitive than ‘said-that’ constructions to differences between materially equivalent but distinct contents.

It may help to allay fears about truth-conditional pseudo-clefts to consider a few examples. Cappelen and Lepore (1997) provide a range of cases in which the truth of ‘said-that’ constructions does not appear to require that the indirectly reported proposition match the proposition that was originally expressed. I will consider two here, showing that no analogous problem arises for truth-conditional pseudo-clefts. I believe much the same can be said about other cases.

(5a) A: I bought a pair of Bruno Magli shoes and then I ate lunch.

⁴⁸For discussion of the exclusiveness of pseudo-clefts and a suggestion that they be analyzed in terms of conventional implicature, see Collins (1991).

(5b) B: A said that he bought a pair of Bruno Magli shoes.

(5c) B: What A asserts in (5a) would have been true if he had bought a pair of Bruno Magli shoes but not eaten lunch.

Whereas the 'said-that' construction in (5b) is able to embed a sentence expressing only a single conjunct of the conjunction A asserts in (5a), the infelicity of the truth-conditional pseudo-cleft in (5c) suggests that it is tracking the full conjunctive proposition rather than only its first conjunct.

(6a) François: Chartreuse is Maria's favorite color.

(6b) Speaker (demonstrating a chartreuse dress): François said that the color of that dress is Maria's favorite color.

(6c) Speaker (demonstrating a chartreuse dress): What François asserts in (6a) would have been true if the color of that dress had been avocado and avocado had been Maria's favorite color.

Whereas the 'said-that' construction in (6b) is able to embed a sentence expressing a proposition which is materially equivalent to but distinct from the proposition François asserts in (6a), the infelicity of the truth-conditional pseudo-cleft in (6c) suggests that it is tracking the original proposition rather than the one which is embedded under 'said-that' in (6b).

I conclude that there is no special reason to worry about treating truth-conditional pseudo-clefts as evidence in developing a theory of assertion. Indeed, even if truth-conditional pseudo-clefts were demonstrably context-sensitive, this fact would undermine my argument against (Diagonalism) only if it could be shown that I

have mischaracterized their truth conditions *in sentences (1-3)* — the onus is on the critic of my argument to show that this is so.⁴⁹

2.4.4 The Coordinationist

Certain metasemantic accounts of supplementives, including King's (2014) coordination account, require that a speaker manifest her intention to fix the value of a context-sensitive parameter openly enough that a reasonable and minimally informed member of her audience would be able to follow along. If we accept the coordination account or something similar, our verdict concerning [READY] must be that Susie fails to determine a unique semantic value for the supplementive expression in her utterance.

There is reason, however, to be discontent with this prediction, as it is difficult to see an appealing way to reconcile it with both the intuitive falsehood of (1-3) and the intuitive truth of (4). For if (1-3) are false but (4) is true, we cannot hold either that all utterances during the production of which a speaker fails to secure a semantic value for a supplementive expression are false, or that they are all true, or that they are neither true nor false.⁵⁰

⁴⁹Thanks to Ofra Magidor for emphasizing this point.

⁵⁰Even if we complicate our picture of assertion so that, for example, Susie is assertion-related in [READY] to two propositions corresponding to the two contextually relevant implicit arguments for 'ready', it is difficult to see how to generate intuitive predictions: either claims about what Susie asserts must be satisfied by one or more propositions to which she is assertion-related, in which case (2) is true, or they must be satisfied by all, in which case we mistakenly predict that, if Susie has only lunch in mind at the time of her assertion, then what she asserts is false if Molly is ready for lunch but not the walk.

2.5 Assertion and Communication

The failure of (Diagonalism) is theoretically interesting in its own right. But, as we have seen, it is also interesting because (Diagonalism) is bound up with a common and natural way of thinking about assertion. Any pressure to reject (Diagonalism), then, is pressure to reject this common and natural way of thinking. In this section, I describe the perspective on assertion in question and clarify its connection with (Diagonalism).

In asserting, speakers communicate information to their interlocutors. When all goes well, speaker and audience leave a communicative exchange with more in common doxastically than when it began. Asserting thus facilitates our efforts to coordinate on a shared conception of the world.

The propositional contents of assertoric utterances distinguish between possibilities in just the way required for coordination on a shared conception of the world. There is therefore a natural theoretical impulse to hold that the propositional content of an assertion is the content which a speaker communicates to her audience. If we understand communication of a content in terms of excluding possibilities in which it is false, the theory which results from this identification of assertoric content and communicative upshot is as follows:

(Communication): When a speaker asserts a proposition, her interlocutors rule out for the purposes of the conversation all possibilities in which that proposition is false.⁵¹

⁵¹Provided, of course, that the speaker's assertion is not rejected by her interlocutors. Note that (Communication) is consistent with the claim that interlocutors sometimes also rule out additional possibilities, as when they come to accept an implicated content in addition to an asserted one.

Let us refer to a systematic way of pairing assertoric utterances with propositions as a *species of utterance content*, and let us say that a species of utterance content *plays the communicative role* just in case it pairs any given assertoric utterance with the proposition which characterizes its communicative upshot. Then (Communication) is the claim that assertoric content plays the communicative role.⁵²

Note that (Communication) is not a claim about what is distinctive or individuating of the speech act of assertion. It requires only that, however assertion is individuated, asserting a content involves communicating it to one's audience. For this reason, (Communication) is consistent with all major accounts of the nature of assertion.

If communication can be understood in terms of the ruling out of possibilities, (Communication) offers the most direct possible explanation of the platitude that we assert in order to communicate information to one another: we are able to do so because what is asserted *is* what is communicated. But there is also a further motivation for adopting (Communication). For (Communication) seems able to play an important role in explaining the dynamics of conversation. If an assertion of some content p in the context of an ongoing conversation is not rejected, it generally becomes infelicitous for interlocutors to perform various speech acts: to assert a content inconsistent with p , to ask whether p , to hypothesize that $\neg p$, and so forth. It is attractive to explain this observation in terms of (Communication) and the principle that the set of speech acts which are felicitous continuations of a conversation at a time is constrained by which possibilities are live for the purposes of that conversation at that time. For our purposes, we can limit our

⁵²See Stalnaker (1978, p. 86) and Dummett (1991, p. 47) for endorsements of (Communication).

attention to one such constraint:

(Connection): *Ceteris paribus*, it is acceptable to perform an assertoric utterance in the course of joint inquiry just in case the assertoric content of that utterance, if accepted, would eliminate some but not all of the worlds in the context set.

To endorse (Communication) while rejecting (Connection) would be to decide to use the context set to model how assertion allows interlocutors to coordinate on a shared conception of the world while at the same time denying that the shared conception on which they coordinate shapes the future course of their project of inquiry — a logically possible but unappealing theoretical picture. I will therefore assume in what follows that friends of (Communication) should also endorse (Connection). For clarity, let us refer to the conjunction of (Communication) and (Connection) as *The Possibility-Carving Picture*.

The Possibility-Carving Picture is appealingly simple, intuitive, and explanatorily powerful.⁵³ It can be shown, however, that the conjunction of (Communication) and (Connection) requires us to adopt (Diagonalism). It follows, given the argument of the preceding sections, that The Possibility-Carving Picture is false.

Suppose we pair assertoric utterances with propositional contents by associating any utterance u with the set of worlds w such that the assertoric content of u is

⁵³It is not, however, universally accepted. Some who prefer to think about belief in terms of guised propositions may hold that when a speaker asserts, her interlocutors come to accept the content of her assertion for the purposes of the conversation, but deny that this involves ruling out for the purposes of the conversation all possibilities in which the content of her assertion is false (thus rejecting (Communication)). In general, they may think, coming to accept a content under some guise does not enable hearers to rule out all possibilities in which it is false; some such possibilities could only be ruled out if they came to accept the same content under a different guise.

true at w . Call the proposition thus associated with an utterance its *proto-diagonal*. The construction of an utterance's proto-diagonal proposition is thus very similar to the construction of its superdiagonal proposition, which was introduced in section 2. The two differ, however, in that the superdiagonal proposition of an utterance is defined in terms of its horizontal proposition at various worlds (which is in turn grammatically determined), whereas its proto-diagonal proposition is defined in terms of its assertoric content at various worlds, which could in principle fail to be grammatically determined.

I have shown in Chapter 1 that (Communication) and (Connection), together with a few plausible ancillary premises, suffice to establish the following principle:

(Proto-Diagonalism): For any conversation c , assertoric utterance u , and world w , the assertoric content of u at w agrees with its proto-diagonal proposition in truth-value at every world in the context set of c at w .⁵⁴

Informally presented, the argument for this conclusion is that, if (Proto-Diagonalism) could be violated, then, given (Communication), accepting an assertoric utterance of a sentence at a world could involve either ruling out some worlds in the context set at which that utterance expresses a truth or failing to rule out some worlds in the context set at which an utterance of the negation of that sentence expresses a truth. But, given (Connection), this would mean that it could be felicitous to assert a sentence in a conversation where it is presupposed that that sentence expresses a truth, and that it could be felicitous to assert a sentence in a conversation

⁵⁴In the chapter 1, I call this principle '(Diagonal)'; I have renamed it here to avoid confusion with (Diagonalism).

where it is presupposed that an assertion of its negation would express a truth. But neither of these sorts of utterances are in fact felicitous. So (Communication) and (Connection) together secure (Proto-Diagonalism).

(Proto-Diagonalism) is not (Diagonalism), since (Proto-Diagonalism) but not (Diagonalism) could be true in a world where the assertoric contents of speakers' utterances bore no relation to the grammar of their language. Nevertheless, the connection between The Possibility-Carving Picture and (Proto-Diagonalism) is dialectically significant in that (Diagonalism) is by far the most plausible theory of assertoric content according to which (Proto-Diagonalism) is true — among theories of assertoric content which have been proposed in the literature, it alone both satisfies (Proto-Diagonalism) and shows how the assertoric content of an utterance depends in a straightforward way on the grammar of the community in which it is produced. There is thus a compelling abductive case from (Proto-Diagonalism) to (Diagonalism): an argument for (Proto-Diagonalism) is *ipso facto* an argument for (Diagonalism). Our argument against (Diagonalism), then, suggests that at least one of (Communication) and (Connection), and thus also The Possibility-Carving Picture, must be rejected.

Rejecting The Possibility-Carving Picture raises two questions: first, the question of how to explain how speakers coordinate doxastically via assertion without accepting (Communication) and (Connection); second, the question of what theoretical interest remains in the theory of assertion if we give up (Communication) or (Connection). I will consider the former question in this section and the latter in the next.

Having given up (Diagonalism), we must decide what to make of (Communication) and (Connection). Some may be tempted at this juncture to abandon

The Possibility-Carving Picture's underlying model of communication entirely — perhaps communicating is more like getting the right sentence into one's interlocutors' heads than causing them to rule out the right possibilities. But I would like to suggest a less dramatic departure from The Possibility-Carving Picture. In [READY], it is clear that the communicative effect of Susie's utterance is to update the common ground with one of its diagonal propositions: to rule out world *c*, at which Susie has both lunch and the walk in mind and Molly is ready only for lunch. This suggests an appealingly minimal revision of The Possibility-Carving Picture: in cases like [READY], though no diagonal proposition of an utterance is its assertoric content, some diagonal proposition is the content which is intersected with the context set. Though (Communication) is false, then, it is *almost* true — assertions *do* characteristically cause interlocutors to rule out possibilities for the purposes of the conversation, though (Communication) incorrectly characterizes the set of possibilities ruled out in this way. Similarly, though (Connection) mistakenly identifies the proposition acceptance of which must rule out some but not all worlds in the context set as the assertoric content of an utterance, a modified version of the principle which replaces talk of an utterance's assertoric content with talk of its diagonal propositions remains plausible.

The lesson I propose to draw from the failure of (Diagonalism), then, is that it is not the assertoric content of an utterance but rather some diagonal proposition of that utterance which plays the communicative role. Since all diagonal propositions of an utterance are equivalent over the context set of a conversation, it does not matter which one we choose. For the sake of simplicity, however, no harm will be done if we hold that it is the superdiagonal proposition of an utterance which plays the communicative role.⁵⁵

⁵⁵I have so far said nothing about the status of (Horizontalism). In [READY], we have seen that

2.6 Assertion and Commitment

It may seem natural at this point to conclude that the failure of The Possibility-Carving Picture suggests that assertoric content is not particularly interesting: if (Communication) and (Connection) are false, then what role does it play apart from explaining our judgments about certain sentences describing cases like [READY]? This is the second of our two questions from the previous section. In what follows, I suggest that this dismissive attitude toward assertoric content is wrongheaded: though The Possibility-Carving Picture is false, the theory of assertoric content retains its interest because it is intimately related to the theory of assertoric *commitment*.

It is a platitude that we undertake commitments when we assert, becoming liable to criticism for speaking falsely, misleadingly, or otherwise inappropriately — accountability is the price of communication. Indeed, intuition delivers more than the platitude: intuitively, there is a distinctive sort of commitment which a speaker undertakes in asserting, this commitment is associated with a unique propositional content, and any other commitments incurred by a speaker in asserting are derivative of her commitment to this content.

Let us refer to the unique propositional content to which a speaker undertakes this distinctive sort of commitment in asserting as the *principal commitment* associated with her utterance. Then we can say that a species of utterance content *plays the committal role* just in case it pairs any given assertoric utterance with the

the horizontal proposition is asserted rather than any diagonal proposition. Having distinguished between assertoric content and other species of utterance content, then, it is tempting to embrace the horizontal theory of the former. This is, however, a matter on which I wish to take no stand. Perhaps there are other cases (the Gabor/Anscombe case, for example) in which some diagonal proposition of an utterance is asserted; perhaps not. All that is required for my purposes here is that there are some cases in which no diagonal proposition is asserted.

principal commitment its speaker undertakes in producing it. Just as the thesis that assertoric content plays the communicative role is the most straightforward possible explanation of the platitude that speakers assert in order to communicate, the thesis that assertoric content plays the committal role is the most straightforward possible explanation of the platitude that speakers take on commitments when they assert. Thus the following principle will be familiar, similar as it is to principles endorsed by a wide range of theorists, including Searle (1976, p. 10), Brandom (1984, p. 640), and Williamson (2000, pp. 268-269):

(Commitment): When a speaker asserts a proposition, the principal commitment she incurs is to that proposition.

Though it is both familiar and plausible, (Commitment) is by no means trivial. I have insisted that the theory of assertion be held accountable to our intuitive judgments about the truth and falsity of what speakers have asserted. (Commitment) is the substantive thesis that the propositional objects which explain these judgments are also the objects to which speakers undertake principal commitments in asserting. In the remainder of this section, I will motivate this thesis by proposing a heuristic for identifying the principal commitments speakers undertake in asserting and showing that this heuristic suggests that, even in cases where assertoric content does not play the communicative role, it nevertheless continues to play the committal role.

How can we single out the principal commitment associated with a given assertion? An initial impulse is to appeal to liability to blame: a proposition p is the principal commitment a speaker incurs in asserting just in case she is liable to blame in the conversational state resulting from her assertion if p is false or she

does not have appropriate evidence for p (and she lacks a compelling excuse for committing to p).⁵⁶

Unfortunately, a blame-only account of principal commitment does not distinguish between the assertoric content of an utterance and certain obvious entailments of that content. If a speaker is liable to blame if the proposition that Mark is in the kitchen is false, then she is liable to blame if the proposition that Mark is in the house is false. Likewise, if a speaker is liable to blame if it is false that John knows that Mark is in the kitchen, then she is liable to blame if it is false that Mark is in the kitchen. But, as we have seen, the notion of the principal commitment associated with an assertion is intuitively not indiscriminate in this way: a speaker who asserts that Mark is in the kitchen is liable to blame if the proposition that Mark is in the house is false *because* she is liable to blame if the proposition that Mark is in the kitchen is false, not vice versa.

A natural move to make at this point would be to introduce a notion of *non-derivative* liability to blame and characterize the principal commitment associated with an assertoric utterance u as the unique proposition p such that, in performing u , a speaker becomes non-derivatively liable to blame if p is false or she does not have appropriate evidence for p . If liability to blame for the falsity of a proposition because it is entailed by another proposition for the falsity of which one is liable is a species of derivative liability, this move would allow us to circumvent the overgeneration problem just introduced.

Unfortunately, even if it is granted that there is always a unique non-derivative

⁵⁶Note that I do not wish to commit myself to any particular account of the epistemic norm of assertion or of the connections between liability, blame, and excuse. I invite the reader to reformulate my arguments in this section using her favored accounts of these topics — the conclusions should be unaffected.

commitment associated with any given assertion, intuition is not always clear concerning which of a speaker's commitments is non-derivative. In cases where assertoric content does not play the communicative role, which commitment of the speaker is non-derivative? Consequently, I wish to explore the possibility that the principal commitment associated with an assertion can be singled out as involving a distinctive combination of liability to blame and *vulnerability to criticism*, where this latter notion is understood in terms of targetability by certain expressions—most notably 'that'—which can be used to refer to propositions.⁵⁷

It is worth briefly pointing out that, like blameability, targetability with 'that' alone cannot be used to pick out the principal commitment associated with a speaker's assertion. For the contents of other speech acts are targetable in the same way. If I conjecturally utter, "Perhaps the train is late because of track work," you can felicitously respond, "That's unlikely — there's no announcement on the website," with your use of 'that' targeting the proposition that the train is late because of track work, to which I have not committed.

I therefore propose that the principal commitment associated with an assertoric utterance can be picked out terms of two factors: non-derivative liability to blame and targetability by propositional anaphors. Undertaking a principal commitment to p , on this picture, involves both introducing p as a candidate for subsequent reference and exposing oneself to blame if p is false or one does not have appropriate evidence for p :

(Commitment Heuristic): Proposition p is the principal commitment a speaker undertakes in performing assertoric utterance u just in case (i) in virtue of

⁵⁷That 'that' can be used to refer to propositions is suggested by the semantic non-defectiveness of the sentences 'That is true' and 'That is false' in certain contexts.

performing u , she is non-derivatively liable to blame if p is false or she does not have appropriate evidence for p , and (ii) p is targetable by ‘that’ in the conversational state resulting from her performance of u .

Note that (Commitment Heuristic) is not a theory of the nature of assertoric commitment. No theory of the nature of assertoric commitment could appeal to targetability by ‘that’, since there can be assertions in languages other than English — perhaps even languages lacking propositional anaphors entirely. (Commitment Heuristic) is intended merely to state conditions which single out the principal commitment associated with a speaker’s assertion in a useful range of cases. This more modest achievement is all that is required for the purpose of distinguishing between an utterance’s superdiagonal proposition and its assertoric content when it comes to assertoric commitment.

One might worry at this point that, in appealing to availability for anaphoric reference using ‘that’, (Commitment Heuristic) stacks the deck in favor my preferred view that the species of content which plays the committal role is the same as the species of content which explains our intuitions about cases like [READY]. This concern is misplaced, however, for two reasons. First, it is nontrivial to claim that ‘what is asserted’ or ‘that’ must pick out a proposition for which the speaker becomes *non-derivatively* liable to blame when she asserts. Second, it is independently plausible that the species of utterance content which plays the committal role should be targetable with ‘that’ in the conversational state resulting from an assertion. The alternative would be for a speaker’s principal commitment to play a role in discourse similar to the role played by presupposed or conventionally implicated content — to be targetable only by interrupting the conversation with ‘Hey, wait a minute!’ or similar devices. This would be a strange situation: the commitments speakers explicitly take up in asserting are not peripheral like these

other forms of content.

Granting (Commitment Heuristic), our question becomes whether, when the assertoric content of an utterance does not play the communicative role, it nevertheless plays the committal role. I claim that this is the case. My argument for this claim has the following structure: Consider an arbitrary case in which the assertoric content of a speaker's utterance u does not play the communicative role and a sentence ϕ which expresses a proposition that entails the superdiagonal proposition of u , but not its assertoric content, in the context set. Suppose a member of the speaker's audience learns the proposition expressed by ϕ prior to the conversation. If, after the speaker asserts u , this audience member cannot felicitously assert the sentence 'I already knew that', 'that' must be unable to target any diagonal proposition of u . Given that an utterance's assertoric content and superdiagonal proposition are the only contents for which its speaker plausibly assumes non-derivative liability to blame in producing it, and given that 'that' does seem able to target an utterance's assertoric content, establishing that 'that' cannot in general target an utterance's superdiagonal proposition suffices to establish that assertoric content plays the committal role even when it does not play the communicative role.

Consider, for example, [READY*], and suppose we fill in the case to include the following backstory: the day prior, Matt encounters his prophetic friend Sibyl, who informs him that Susie will soon assert the sentence 'Molly is ready' and thereby express a truth. Then Sibyl's claim contextually entails every diagonal proposition of Susie's utterance in [READY*] but not its assertoric content. But, in response to Susie's assertion of 'Molly is ready,' Matt cannot felicitously say:

(7) I already knew that.

Indeed, (7) seems not only infelicitous but false. After his conversation with Sibyl but before his lunch plans are made the following day, Matt has no idea what implicit argument Susie will have in mind when she utters ‘Molly is ready’. Since Matt *does* already know a diagonal proposition of Susie’s utterance, the infelicity (and seeming falsehood) of (7) suggests that ‘that’ cannot target this diagonal proposition. But the assertoric content of Susie’s utterance is the only other natural candidate for the referent of ‘that’. And [READY*] was an arbitrary example of the relevant kind of case; the same intuitive verdict applies to other cases, as well.

A second example will clarify both the structure and the intuitive force of the test. You have volunteered to participate in a psychological experiment designed to investigate whether subjects can reliably guess your preferences. Franz is running the experiment; he has chosen you because he has great faith in your honesty. The experiment will work as follows: at noon tomorrow, you and one of Franz’s subjects will be presented with an array of objects. Both you and the subject will have a clear view of the objects; the subject’s view of your upper body, however, will be obscured. Your task will be to point to an object and say to the subject ‘That is my favorite object’. The subject will then be asked to make a prediction about which of the objects is your favorite, and Franz will note whether he is correct or incorrect. This process will be repeated until Franz has a satisfactory idea of the subject’s ability to guess your preferences regarding objects.

Supposing that Franz’s faith in your honesty is not misplaced, and that he has successfully communicated it to the subject, he and the subject can know today that at noon tomorrow you will point to an object and, speaking truthfully, say ‘That is my favorite object’. Now imagine that, at noon tomorrow, just after you point at an object and say ‘That is my favorite object’, the subject says ‘I already

knew that'. If 'that' can target a diagonal proposition of your utterance, then the subject's utterance should be felicitous. But it is not felicitous — indeed, what the subject says is intuitively false. The natural explanation is that, in assertorically uttering 'I already knew that', the subject can only mean that he already knew the horizontal proposition associated with your utterance, which is its assertoric content — knew, that is, that the particular object you demonstrated was your favorite. But, of course, the subject cannot know this proposition because he does not know which demonstration accompanied your utterance. So (Commitment Heuristic) once again suggests that, when the assertoric content of an utterance diverges from its superdiagonal proposition, the principal commitment a speaker incurs in producing that utterance is the former rather than the latter.

2.7 Conclusion

Combining (Commitment) with the thesis that superdiagonal content rather than assertoric content plays the communicative role yields a substantive theory of conversation which is able to capture the insights of The Possibility-Carving Picture while also accommodating our intuitive judgments about [READY] and related cases. The Possibility-Carving Picture is superseded by one which recognizes greater complexity in the structure of conversation: where The Possibility-Carving Picture posits a single species of utterance content which is both asserted and communicated, the view I have advocated has room enough for two distinct species, one—assertoric content—which plays the committal role, and one—superdiagonal content—which is communicated.⁵⁸

⁵⁸I would like to thank Liz Camp, Andy Egan, Thony Gillies, John Hawthorne, Ethan Jerzak, Jeff King, Harvey Lederman, Ernie Lepore, Ofra Magidor, Eliot Michaelson, Paul Pietroski, Jonathan Schaffer, and Susanna Schellenberg for helpful comments on earlier drafts of this paper, as well as

3 Why Horizontalism

3.1 Introduction

Death is an occasion for philosophical reflection. Consider:

[SHERLOCK]: *A terrible scream — a prolonged yell of horror and anguish burst[s] out of the silence of the moor.*⁵⁹ Holmes runs towards the source of the sound. Watson, never quite his companion's equal, trails behind. There has been foul play, both men are convinced, and the victim is either Baskerville or Selden. Reaching the scene of the incident well before Watson, Holmes discovers a body he recognizes as Selden's at the base of a cliff. "He's fallen quite a distance!", he calls over his shoulder to Watson, who has not yet come close enough to observe the body himself.

[SHERLOCK] contains two puzzles—one for the detective, and one for the philosopher. The detective's puzzle is, of course, the identity of the criminal who set his hound on Selden, causing him to plunge to his death as he fled in terror. The philosopher's puzzle, which is equally difficult though perhaps less obvious, concerns Holmes's utterance. For when Holmes utters "He's fallen quite a distance!", he does so knowing that Watson is in no position to identify the referent of the pronoun 'He' (though Watson knows it is either Selden or Baskerville). And Watson knows that Holmes knows this; in fact, it is common ground between the two

audiences at Rutgers, the New York Philosophy of Language Workshop, and the Pacific Division of the American Philosophical Association for helpful discussion.

⁵⁹Conan Doyle (1902, 189). Devoted fans will, I hope, excuse any philosophically motivated alterations of the plot.

that, for all Watson knows, Holmes's 'He' might refer to Baskerville and might refer to Selden. So [SHERLOCK] has the following general structure: a sentence containing a context-sensitive lexical item is uttered assertorically when it is common ground between the utterer and his interlocutor that the latter is ignorant of features of the context relevant to the determination of the semantic value of that lexical item.

What, then, is the propositional content Holmes asserts in uttering "He's fallen quite a distance!?" According to one kind of answer, to which we may pre-theoretically refer as *horizontalist*, since Holmes sees that the body is Selden's and intends in uttering 'He' to refer to Selden, he asserts the proposition *that Selden has fallen quite a distance*. According to a competing kind of answer, to which we may pre-theoretically refer as *diagonalist*, since it is common ground between Holmes and Watson that Watson is not in a position to determine which individual Holmes intends to refer to in uttering 'He', Holmes does not assert the proposition *that Selden has fallen quite a distance*; instead, he asserts something more like the proposition *that the individual to which he intends to refer has fallen quite a distance*. These are different propositions: the second, but not the first, is true in possibilities where the body on the moor which prompts Holmes's utterance is Baskerville's rather than Selden's.⁶⁰

The question of which proposition Holmes asserts in uttering "He's fallen quite a distance!" should be distinguished from the question of which proposition

⁶⁰To say that Holmes *asserts* a given proposition is not to say that Watson *learns* that proposition from Holmes's utterance. The proposition Watson learns from Holmes's utterance may or may not be the same as the proposition Holmes asserts — I do not wish to take a stand on this issue. For those to whom assertion-talk does not come naturally, the difference between the horizontalist and diagonalist answers can to an extent be appreciated by considering the question of what Holmes *tells* Watson in [SHERLOCK], though, again, this question must be distinguished from the question of what Watson comes to believe on the basis of this telling.

the grammar of English assigns to the sentence “He’s fallen quite a distance!” as uttered in the context described in [SHERLOCK]. Fixing an answer to the latter question — for example, by holding that context determines the semantic value of ‘He’ so that it refers to whomever the speaker has in mind while speaking — does not yet fix an answer to the former. It *may* be that Holmes asserts whatever proposition the grammar of English assigns to the sentence he utters relative to the context in which he utters it. But it may also be that what Holmes asserts is more loosely connected to the grammar of English; in what follows, I will consider a number of ways in which a looser connection might be realized.⁶¹ For this reason, two theorists who agree about the grammar of English might nonetheless disagree about the correct answer to the philosopher’s puzzle.⁶²

Indeed, it may be that, as Lewis (1980) and others have argued, the grammars of natural languages do not in fact determine unique propositional contents for sentences in contexts — it may be that they determine only functions from indices to truth values. The intelligibility of this possibility renders especially vivid the point, made variously by Dummett (1991), Ninan (2010), and Rabern (2012), that

⁶¹Indeed, with the exception of what in sections 2 and 3 I will call (Nondefective Objective Horizontalism) and (Objective Horizontalism), every way of associating utterances with assertoric contents considered below is one according to which the assertoric content of an utterance can diverge from its grammatically determined content.

⁶²This is not to say that any two theorists who disagree about the correct answer to the philosopher’s puzzle must agree about the grammar of English. Disagreement about the correct answer to the philosopher’s puzzle could be rooted in disagreement at any of a number of levels: disagreement about the meanings of the context-insensitive lexical items which occur in Holmes’s utterance, for example, or disagreement about how the context of Holmes’s utterance determines the semantic value of the pronoun ‘he’ (that is, disagreement about the correct *metasemantics* for deictic pronouns in English — see footnote 68 below). The point I wish to register is simply that two theorists might disagree about the correct answer to the philosopher’s puzzle solely in virtue of disagreeing about the correct theory of assertion; that is, that agreement about grammar (construed so as to incorporate both semantics and metasemantics) does not suffice for agreement about the correct answer to the philosopher’s puzzle. This more modest claim is all I need to motivate the discussion which follows. Thanks to an anonymous referee for pressing me to clarify this point.

a theory of the grammatical dependence of semantic content on context (a theory which belongs to the domain of semantics) is not ipso facto a theory of assertoric content (a theory which belongs to the domain of pragmatics).⁶³

In the years since Stalnaker (1978) first popularized the idea that the theory of assertoric content could be developed in relative isolation from semantic theory, work on assertion has reached no clear consensus favoring either the horizontalist or the diagonalist answer to the philosopher's puzzle. Lewis (1980) is sensitive to the distinction between the two, but does not take a stand concerning which of them (if either) is correct. Rabern (2012) assumes that the horizontalist answer is correct. Stojnić (fc) extends an argument due to Soames (2002) favoring the horizontalist answer over the diagonalist. Stalnaker (2014), meanwhile, argues in favor of the diagonalist answer. His arguments are criticized by Hawthorne and Magidor (2009, 2010) and defended by Almotahari and Glick (2010).

In what follows, I present a number of arguments which bear on the debate between horizontalists and diagonalists in a novel way. It is a platitude that, in asserting a proposition, one thereby undertakes a commitment to its truth.⁶⁴ My strategy, then, is to hold fixed the idea that the assertoric content of an utterance is the content to which a speaker undertakes a commitment in producing that utterance, considering the horizontalist and diagonalist answers to the philosopher's

⁶³Thanks to an anonymous referee for pressing me to clarify this point.

⁶⁴It is sometimes suggested that this is not a platitude, and that Stalnaker rejects the view that assertion is a committal speech act, on the grounds that he occasionally theorizes about non-committal speech acts, such as hypothetical reasoning and plan formation, along with assertion. Though the Stalnakerian framework is useful for modeling these other speech acts, and though Stalnaker himself seldom writes about commitment, it is a misreading of Stalnaker to attribute this view to him. He makes this clear in a number of places. For example: "I should emphasize that I am not claiming that one can *define* assertion in terms of a context-change rule, since that rule will govern speech acts that fall under a more generic concept. A full characterization of what an assertion is would also involve norms and commitments." (2014, 89).

puzzle in terms of the communal practices of undertaking assertoric commitments which they suggest. For if assertion is a device for undertaking commitments, then competing accounts of assertoric content entail competing claims about the propositions to which speakers undertake commitments by asserting.

When I speak of the communal practice of undertaking assertoric commitments associated with a particular theory of assertoric content, I mean the systematic way in which speakers in a community would take one another to have undertaken assertoric commitments to particular propositions if the theory of assertoric content in question correctly described that community. Thus the communal practice of undertaking assertoric commitments associated with the horizontalist answer to the philosopher's puzzle has Holmes undertaking a commitment to the proposition *that Selden has fallen quite a distance* when he asserts in [SHERLOCK], while the communal practice of undertaking assertoric commitments associated with the diagonalist answer has him undertaking a commitment to the proposition *that the individual to which he intends to refer has fallen quite a distance*. The idea is that a theory of assertoric content, though stutable simply as a rule for associating propositions with utterances, makes predictions about certain aspects of the the behavior of communities of speakers. In light of this connection, I will move freely in what follows between talking about theories of assertion as rules for associating propositions with utterances and talking about them in terms of the committal practices communities would enact if they were governed by those rules.⁶⁵

If certain possible committal practices offer a linguistic community advantages as compared to others, this gives us some reason to expect that existing linguistic

⁶⁵Thanks to an anonymous referee for pressing me to clarify this point.

communities implement those practices. So, if it can be shown that a horizontalist committal practice offers advantages over a diagonalist one, we have some reason to expect that the horizontalist answer to the philosopher's puzzle is correct. Matters are complicated by the fact that, once one embarks on the project of assessing the advantages and disadvantages of possible committal practices, it quickly emerges that the menu of options is rather more extensive than my preliminary exposition of the philosopher's puzzle in terms of the horizontalist and diagonalist answers suggests. Nevertheless, I will attempt to show that a committal practice which corresponds to the horizontalist answer has important advantages over its competitors.

3.2 Committal Practices Introduced

To solve the philosopher's puzzle is to identify the content of Holmes's assertion in [SHERLOCK]. I have proposed that one strategy for doing this is to assess the relative merits of the various committal practices a linguistic community might adopt. Before such comparisons can be made, however, we must have some understanding of what a committal practice is, as well as of the space of possible committal practices. The purpose of this section is to introduce the theoretical notions required to define various committal practices, as well as to present four simple practices, three of which will then be generalized in the next section.

As I will understand them, committal practices are ways of associating a *grammar* with a *social practice of asserting*.

A grammar — for our purposes, a function from sentence/context/world

triples to truth values — is a useful thing.⁶⁶ For suppose a linguistic community has settled on a grammar. Because every utterance is an utterance of a particular sentence in a particular extralinguistic context, a grammar determines a function from utterances to the propositions (functions from possible worlds to truth values) they grammatically determine.⁶⁷ This means that members of the community can exploit shared information about their grammar and about relevant features of the extralinguistic contexts in which utterances are produced to raise to salience certain propositions which would otherwise have been extremely difficult to coordinate on doxastically.⁶⁸

A social practice of asserting — of producing an utterance in order to undertake a certain kind of commitment to the truth of a proposition — is a useful thing. For suppose a linguistic community has adopted such a practice. Then interlocutors can do more than merely raise to salience certain propositions: they can *describe* the world to one another.

⁶⁶I thus set aside for the sake of simplicity proposals, like that of Lewis (1980), on which the grammar of language is a function from ordered n -tuples (with n greater than 3) to truth values. Everything I say in what follows is compatible with such proposals, except that the procedures for recovering the horizontal and superdiagonal propositions of an utterance given below must be revised to account for the additional complexity of the underlying grammar.

⁶⁷There is a small literature concerning cases in which the passage from utterance to extralinguistic context is unusual, as in the so-called “answering machine paradox” (see, for example, Sidelle (1991) and Predelli (1998)). We need not be overly concerned about such cases in what follows, however, since it is a criterion of adequacy for any theory of them that it show how, in everyday cases like the ones at issue here, we can move from a given utterance to the context in which it was produced.

⁶⁸How exactly does a context determine which proposition is grammatically associated with a given sentence? Borrowing terminology from King (2014), let us refer to this question as the question of the *metasemantics* of context sensitivity. Though in what follows I will sometimes write as if I take for granted an intentionalist metasemantics according to which the propositions grammatically associated with sentences containing demonstratives and deictic pronouns are determined by the referential intentions (and perhaps also gestures) of the speakers who utter them, this is an issue on which I do not wish to take a stand. Because they concern assertoric content rather than semantics, my arguments in what follows are compatible with any plausible metasemantics for context-sensitive vocabulary. Thanks to an anonymous referee for pressing me to clarify this point.

A complete account of a social practice of asserting must consist of (at least) two parts: first, a specification of the speech acts available to participants in the practice for undertaking assertoric commitments to various propositions; second, a description of the normative status of being assertorically committed to the truth of a proposition. The general shape of the first of these parts is familiar: the acts in question are sincere utterances of sentences in the indicative mood. The work of a committal practice is to pair such sincere utterances with the propositions to which speakers commit themselves by performing them. Any interesting committal practice will be such that the proposition associated with an utterance is systematically related to the sentence uttered, the context in which the utterance is produced, and the grammar the community has adopted; what distinguishes various practices is what they take this systematic relation to be, and whether any additional factors are relevant.⁶⁹

About the normative status of being assertorically committed to the truth of a proposition, I hope I will be excused for saying little. A satisfying account of such a normative status would characterize the conditions under which one who assumes it is deserving of praise and blame, and might also show how its existence is grounded in facts about what speakers collectively believe and intend, or about how they are disposed to reward and punish one another. Answering these questions is beyond the scope of my discussion here. For my purposes in what follows, it will suffice to point out that the status of being committed to the truth of a proposition in the manner peculiar to assertion is not the same as the status of being responsible for having intentionally produced belief in that

⁶⁹In what follows I will assume, along with most others who have written on assertion, that an utterance is associated with at most one assertoric content. For an alternative picture, according to which assertion must be modeled as a relation between utterances and contents rather than as a function from utterances to contents, see Soames (2005).

proposition in some audience. I understand this as a descriptive claim; I will shortly present my reasons for endorsing it. It is worth pointing out, however, that there is a corresponding theoretical question: the question of what it is in virtue of which the two levels of commitment differ. Answering this question would require giving a metaphysical account not only of the normative status of being assertorically committed to the truth of a proposition but also of the normative status of being responsible for having intentionally produced belief in a proposition. Again, though this is an interesting question, it is beyond the scope of my discussion here to answer it. Let us turn now to the descriptive question of whether assertoric commitment outstrips mere responsibility for having produced a belief.⁷⁰

To see that the kind of commitment associated with assertion goes beyond the kind of commitment associated with mere intentional communication, consider the intuitive contrast between cases in which a speaker deliberately asserts a false proposition and cases in which she merely intentionally communicates the same proposition. If I yawn and thereby deliberately produce in you a false belief that I am tired, I am perhaps deserving of blame. But if I *testify* to you that I am tired, in addition to whatever blame I deserve for intentionally communicating a falsehood, I am also blameworthy for lying, speaking falsely, and so forth. Similarly, if I praise a candidate's handwriting and thereby deliberately produce in you a false belief that he or she lacks philosophical promise, I am deserving of blame. But if I actually *testify* to you that the candidate lacks philosophical promise, in addition to whatever blame I deserve for intentionally communicating this falsehood, I am also blameworthy for defaming the candidate. With greater commitment comes

⁷⁰Thanks to an anonymous referee for pressing me to distinguish between this descriptive claim and the corresponding theoretical question.

greater liability to punishment: as a moral and linguistic community, we respond to defamation with much more serious censure than we do to the implication or insinuation of false and harmful propositions, and, more generally, to lying with much more serious censure than to merely misleading. Correspondingly, the availability to speakers of certain defense strategies in the face of criticism (e.g. “Don’t accuse me of lying — I never *said* that *p*; at most I suggested it!”) tracks the difference in seriousness between genuine assertoric commitment and mere intentional communication.^{71,72}

A linguistic community which has settled on a grammar and wishes to institute a social practice of asserting is faced with an important decision. For, while we may assume that members of the community are certain about which grammar they have settled on, they will often *not* be certain about the features of particular contexts of utterance to which their grammar is sensitive. Moreover, this uncertainty will often be obvious to all interlocutors in a conversation: it will often be *common ground* that certain interlocutors are uncertain about relevant features of the context in which an utterance is produced.⁷³

In such cases, a speaker’s utterance at a given world is associated with mul-

⁷¹See section 6 for an application of this observation.

⁷²Note that the claim that the commitment involved in genuine assertion goes beyond the commitment involved in mere intentional communication is consistent with attempts to analyze assertion at least partly in terms of intentional communication. In the Gricean tradition, for example, what a speaker means in performing a communicative act is defined partly in terms of what she thereby intends to communicate, and what a speaker says in producing an utterance (which for our purposes we may understand to be equivalent to what she asserts) is defined partly in terms of what she means in producing that utterance. This sort of view can be reconciled with the intuitive attractiveness of holding that speakers are more strongly committed to what they assert than to what they (for example) conversationally implicate as long as we think that the extra conditions which must be met for a proposition to be said rather than merely meant can be relevant to the level of commitment a speaker undertakes in putting it forward.

⁷³I will understand the common ground of a conversation at a time to be the set of propositions all interlocutors take for granted for the purposes of the conversation at that time.

tuple salient propositions. One is the proposition the grammar assigns to the utterance when we feed it the sentence uttered and the context in which it was uttered at the world in question. Let us call this the *horizontal proposition* of the utterance.⁷⁴ Another is the proposition we get by feeding the grammar the sentence uttered and then: for each world at which the utterance exists, feeding in that world and the context in which the utterance is produced at that world, and collecting up the worlds at which the result of this process is the True. Let us call this the *superdiagonal proposition* of the utterance.^{75,76} A third is the intersection of the superdiagonal proposition with the *context set* of the conversation — that is, the set of worlds which are taken by all interlocutors to be live possibilities for the purposes of the conversation (equivalently, the set of worlds at which every proposition in the common ground is true). Let us call this the *contextual diagonal proposition* of the utterance. The decision our imagined linguistic community must make is whether, in such circumstances, to hold speakers primarily responsible for the truth of the horizontal propositions, the superdiagonal propositions, or the contextual diagonal propositions expressed by their utterances. In other words, the decision our linguistic community must make is between:

⁷⁴More precisely, if G is the grammar, s is the sentence uttered, and c_w is the context in which it is uttered in w , the horizontal proposition of the utterance at w is the set of worlds w' such that $G(\langle s, c_w, w' \rangle) = 1$.

⁷⁵More precisely, if f is the function which maps each utterance/world pair $\langle u, w \rangle$ to the context c_w in which u is uttered at w , then the superdiagonal proposition of an utterance u of a sentence s is the set of worlds w such that $G(\langle s, f(\langle u, w \rangle), w \rangle) = 1$.

⁷⁶The truth-values of the superdiagonal proposition of an utterance at different worlds are always calculated with reference to the grammar G actually in use by the community. A distinct proposition, which we may call the *hyperdiagonal*, is definable by feeding each world, sentence, and utterance context into the grammar in use by the community in which the utterance is produced at that world. The superdiagonal and hyperdiagonal propositions of an utterance u at a world w may differ in truth-value at worlds where that utterance is produced in a community with a grammar distinct from the grammar of the community in which it is produced at w . Nevertheless, since the distinction between the superdiagonal proposition of an utterance and its hyperdiagonal proposition does not affect the plausibility of any of the arguments in what follows, I will suppress it for the purposes of this chapter.

(Nondefective Objective Horizontalism): For all conversations c , utterances u , speakers s , and worlds w : If c is nondefective at w ,⁷⁷ and if s assertively utters u at w , then the object of s 's assertoric commitment in uttering u is the horizontal proposition expressed by u at w .

(Nondefective Superdiagonalism): For all conversations c , utterances u , speakers s , and worlds w : If c is nondefective at w , and if s assertively utters u at w , then the object of s 's assertoric commitment in uttering u is the superdiagonal proposition expressed by u .

(Nondefective Contextual Diagonalism): For all conversations c , utterances u , speakers s , and worlds w : If c is nondefective at w , and if s assertively utters u at w , then the object of s 's assertoric commitment in uttering u is intersection of the superdiagonal proposition expressed by u and the context set of c at w .

If we allow that the context set of a conversation at a world may fail to include that world, the possibility arises that it could be common ground at a world w that an utterance expresses a horizontal proposition which it does not in fact express at w . Thus we should also acknowledge the coherence of a further committal practice:

(Nondefective Subjective Horizontalism): For all conversations c , utterances u , speakers s , and worlds w : If c is nondefective at w , and if s assertively utters u at w , and if there is a unique proposition p such that it is common ground at w that p is the horizontal proposition expressed by u ,

⁷⁷A conversation is *nondefective* at a world iff at that world all interlocutors take the same propositions for granted for the purposes of the conversation.

then p is the object of s 's assertoric commitment in uttering u at w .⁷⁸

Except in special cases, (Nondefective Objective Horizontalism) and (Nondefective Superdiagonalism) will assign different assertoric contents to a given utterance.⁷⁹ Similarly, unless the common ground of the conversation in which an utterance occurs is empty (so that its context set is the set of all worlds), (Nondefective Contextual Diagonalism) and (Nondefective Superdiagonalism) will always assign different assertoric contents to a given utterance. Analogous remarks apply to (Nondefective Subjective Horizontalism) vis-à-vis our three other committal practices.

Nevertheless, when a true proposition characterizing the horizontal proposition expressed by an utterance is part of the common ground of the conversation in which that utterance occurs — so that, for example, it is both true and presupposed that when the speaker utters 'He is from Argentina', her 'He' refers

⁷⁸If, following Lewis (1975), we think of a community as having coordinated on a grammar just in case it obeys a convention of truthfulness and trust in the deliverances of that grammar, then we must regard facts about whether sentences are taken to express truths in contexts at worlds as explanatorily prior to facts about the grammars of communities. In that case, it may seem to make little sense to assume that a community has coordinated on a grammar and then ask how members of that community use that grammar to assign assertoric contents to utterances.

But one can maintain that conventions of truthfulness and trust suffice to determine the grammar of a community without conceding the incoherence of questions about committal practices. This is because all of the committal practices described here are ones in which interlocutors engaged in serious inquiry will (i) try to produce utterances the horizontal propositions of which are true at the worlds in which they are uttered and (ii) expect others to do the same — that is, committal practices in which interlocutors will obey conventions of truthfulness and trust in the same grammar. This is because the horizontal and superdiagonal propositions determined by any utterance are materially equivalent in the world in which it is produced. Even in cases where speaker or audience mistakenly rule out the world of utterance as a live possibility or have false beliefs about the horizontal proposition determined by an utterance, they are still *trying* to obey a convention of truthfulness and trust in a particular grammar. So we can, if we want, agree with the committed Lewisian that facts about conventions of truthfulness and trust uniquely determine the grammar on which a community coordinates while also maintaining that they fail to determine which committal practice the community employs.

⁷⁹Special cases include, for example, utterances of sentences like 'That is non-self-identical', which express necessary falsehoods however the context sensitivity of their lexical items is resolved.

to Smith — all four of the committal practices described above agree on which worlds in the context set of the conversation are compatible with the utterance’s assertoric content. For this reason, we must distinguish between the question of which proposition a committal practice pairs with a given utterance, on the one hand, and the question of how that committal practice predicts the content of the utterance in question would change the context set if accepted, on the other. My arguments in what follows aim to show that communities can have reason to favor one committal practice over another even though, under certain circumstances, both answer this second question in the same way.⁸⁰

3.3 Six Practices

For reasons which will become clear in the next section, the fact that the committal practices introduced so far are defined only when the conversation is nondefective is undesirable. Setting aside this idealization, then, our imagined linguistic community must choose between a larger family of possibilities. First, there are three generalizations of (Nondefective Objective Horizontalism):

(Objective Horizontalism): For all utterances u , speakers s , and worlds w : If s assertively utters u at w , then the object of s ’s assertoric commitment in uttering u is the horizontal proposition expressed by u at w .

(Speaker-centered Horizontalism): For all utterances u , speakers s , and worlds w : If s assertively utters u at w , and if there is a unique proposition p such that s takes for granted for the purposes of the conversation at w that p is the horizontal proposition expressed by u , then p is the object of s ’s

⁸⁰Thanks to an anonymous referee for pressing me to clarify this point.

assertoric commitment in uttering u .

(Audience-centered Horizontalism): For all conversations c , utterances u , speakers s , and worlds w : If s assertively utters u at w , and if there is a unique proposition p such that every member of c other than s takes for granted for the purposes of the conversation at w that p is the horizontal proposition expressed by u , then p is the object of s 's assertoric commitment in uttering u .

Note that (Speaker-centered Horizontalism) and (Audience-centered Horizontalism) are also the corresponding generalizations of (Nondefective Subjective Horizontalism).

Then there is the corresponding generalization of (Nondefective Superdiagonalism):

(Superdiagonalism): For all utterances u , speakers s , and worlds w : If s assertively utters u at w , then the object of s 's assertoric commitment in uttering u is the superdiagonal proposition expressed by u .

Since in normal cases it will be transparent what the superdiagonal proposition expressed by an utterance is, the speaker- and audience-centered versions of superdiagonalism are not of sufficient interest to merit independent consideration.⁸¹ Finally, there are two generalizations of (Nondefective Contextual Diagonalism):

⁸¹Of course, in cases where there is ignorance about what property is denoted by a predicate, or about the Kaplanian character of a context-sensitive expression, there will be uncertainty about which superdiagonal proposition is expressed. For example, an utterance of 'ophthalmologists are eye doctors' in English expresses a necessarily true superdiagonal but a contingent hyperdiagonal. So there is some reason to distinguish between speaker- and audience-centered versions superdiagonalism. Nevertheless, since we have assumed that interlocutors have coordinated on a

(Speaker-centered Contextual Diagonalism): For all conversations c , utterances u , speakers s , and worlds w : If s assertively utters u at w , then the object of s 's assertoric commitment in uttering u is the intersection of the superdiagonal proposition expressed by u with what, at w , s takes the context set of c to be for the purposes of the conversation.

(Audience-centered Contextual Diagonalism): For all conversations c , utterances u , speakers s , and worlds w : If s assertively utters u at w , and if there is a unique set of worlds p such that every member of c other than s takes for granted for the purposes of the conversation at w that p is the context set of c , then the object of s 's assertoric commitment in uttering u is the intersection of the superdiagonal proposition expressed by u with p .

Because there is no objective fact of the matter about the context set of a conversation when it is defective, there are only speaker- and audience-oriented versions of contextual diagonalism.

Each of the six labeled propositions just introduced describes a committal practice on which a linguistic community could conceivably coordinate. They do not, of course, jointly exhaust the theoretically possible committal practices a community could adopt. Speakers could, for example, undertake commitments to the horizontal propositions of their utterances at arbitrary worlds outside the context set, or to any of the various diagonal propositions of their utterances (that is, the propositions which agree with the superdiagonal over the context set at the world of utterance but potentially diverge from it elsewhere). But they do capture what

grammar, and since the speaker- and audience-centered versions of (Superdiagonalism) are subject to many of the same criticisms as (Superdiagonalism), I omit further discussion of them in what follows.

I take to be the most theoretically appealing options.⁸² In what follows, I argue that (Objective Horizontalism) has distinctive advantages over its competitors. In particular, I argue that the committal practice described by (Objective Horizontalism) is *general* and *useful*, and that it can naturally be extended into a practice governing the use of optative constructions.

3.4 Generality

A practice of holding speakers accountable for the contents of their utterances is *general* to the extent that it issues verdicts about a variety of cases — that is, to the extent that the rule which characterizes that practice assigns assertoric contents to a wide variety of utterances. One such practice is *strictly less* general than another just in case the latter (i) issues verdicts about every case about which the former issues verdicts, and (ii) issues verdicts about some cases about which the former does not issue verdicts. Similarly, one committal practice is *loosely less* general than another just in case the latter issues verdicts about a wider range of practically significant cases than the former.

Committal practices which assume a nondefective conversation are less general than the others we have considered so far, since the rules which characterize them do not assign assertoric contents to utterances produced in defective conver-

⁸²It may seem that I have neglected to mention two important further options: first, that a speaker is committed to the conjunction of the propositions *newly entailed* by the context set when it is updated with the superdiagonal proposition of her utterance; second, that a speaker's assertoric commitment is determined by some disjunctive rule (for example, Stalnaker's (1978) proposal that one asserts the horizontal proposition of one's utterance in some circumstances and a diagonal proposition in others). The first of these options is, however, illusory. Given that the conjunction of the superdiagonal and the context set will itself be newly entailed by the context set after updating, and given that it is the strongest such proposition, what seems at first to be a further option is in reality equivalent to (Nondefective Contextual Diagonalism). The possibility of a disjunctive committal practice will be discussed in section 9 below.

sations. In particular, (Nondefective Objective Horizontalism) and (Nondefective Superdiagonalism) are strictly less general than (Objective Horizontalism) and (Superdiagonalism), respectively, and (Nondefective Contextual Diagonalism) is strictly less general than its speaker- and audience-centered generalizations. (Nondefective Objective Horizontalism) is also loosely less general than its speaker- and audience-centered generalizations, in so far as cases in which a conversation is defective are considerably more common than cases in which there is uncertainty across the context set about which horizontal proposition is expressed by an utterance.

Similarly, audience-centered committal practices are strictly less general than (Objective Horizontalism) and (Superdiagonalism), as well as loosely less general than speaker-centered practices, in so far as cases where the audience consists of two or more individuals with different beliefs about the horizontal proposition expressed by an utterance are considerably more common than cases in which a speaker's own beliefs fail to determine a unique horizontal proposition for her utterance.

(Speaker-centered Horizontalism) fails to be general, in so far as there may not be a unique possible-worlds proposition which the speaker takes to be the horizontal proposition of her utterance. If, for example, a speaker has formed the mistaken belief that Diana, Princess of Wales and Catherine, Duchess of Cambridge are one and the same, and if this speaker points at Catherine and exclaims "She's a national treasure!", then there is no unique possible-worlds proposition which she takes to be the horizontal proposition of her utterance: instead, there are two equally good candidates between which she fails to distinguish.

(Speaker-centered Contextual Diagonalism) may be subject to the same kind

of difficulty, since a speaker who is, for example, struck by lightning, and who forgets the recent history of an ongoing conversation in which she is participating, may suspend judgment about the common ground of that conversation. But intuitions about such cases are mixed: perhaps such a speaker takes for granted after the lightning strike that the common ground consists of whatever propositions she still takes for granted, given that these can be expected to also be taken for granted by her interlocutors. More generally, perhaps the attitude of *taking for granted for the purposes of the conversation* satisfies principles of positive and negative introspection. If it does, there will always be something a speaker takes to be the common ground of any conversation she is participating in. Since I do not wish to take sides in this debate, I will simply assume that (Speaker-centered Contextual Diagonalism) fares better with respect to generality than (Speaker-centered Horizontalism).

3.5 Utility I: Non-Eliminativity and Unintentional Liability

Assertion is an indispensable part of inquiry. At the same time, assertoric commitment is a serious business, and interlocutors are rational to avoid undertaking commitments for which they may later be censured. One way for a committal practice to fail to be useful, then, is for it to associate propositional contents with utterances in such a way as systematically to discourage interlocutors from asserting. In this section, I will consider two versions of this kind of problem.

At times, the evolution of the context set of a conversation is *non-eliminative*: this occurs when the context set of a conversation at some time t' is not a subset of the context set of the conversation at some earlier time t . Non-eliminative context-set evolution is a normal part of inquiry. Interlocutors who have been

presupposing a proposition may discover that it is false and come to presuppose its negation, or they may question their grounds for presupposing it and come to presuppose neither it nor its negation. In the first kind of case, the context set after the change and the context set prior to the change are disjoint; in the second kind of case, the context set after the change is a proper superset of the context set prior to the change.

The former kind of non-eliminative evolution interacts in problematic ways with contextual diagonalism. If, for example, a community has adopted (Speaker-centered Contextual Diagonalism) as its committal practice, and if a speaker in this community has made any assertion at all, then she has undertaken a commitment to a proposition at least as strong as what she takes to be the context set of her conversation. If she and her interlocutors subsequently come to presuppose the negation of any proposition which was previously presupposed, she will have undertaken an assertoric commitment to a proposition which is false everywhere in the new context set of the conversation, and will therefore be liable to criticism for having asserted falsely. So when the context set evolves in this way, both versions of contextual diagonalism predict that *every* assertion made in the conversation before the non-eliminative update is false. A committal practice with this feature overgenerates the kind of blameworthiness which arises from asserting a falsehood: speakers can perform utterances which express true horizontal and diagonal propositions and nonetheless be liable to criticism for having asserted falsely.

Just as the problem of non-eliminative context-set evolution affects contextual diagonalism, the problem of unintentional liability affects audience-centered committal practices. For any audience-centered committal practice holds assertors hostage to the beliefs of other interlocutors. A speaker in a community which has

adopted (Audience-centered Horizontalism), for example, can be held accountable for asserting a proposition she could not have reasonably believed she would assert. This will happen whenever she has the misfortune of having false but justified beliefs about what her interlocutors take for granted about her referential intentions. If Smith points to Jones and says “He is a thief,” and if Jones is standing next to Johnson, then if Smith’s audience mistakenly believes that he is pointing at Johnson, Smith has committed to the proposition that Johnson is a thief, not the proposition that Jones is a thief. Nor will it help for Smith to clarify that he intended to demonstrate Jones and justifiably believed that this intention would be obvious to all involved — his intention is not relevant to determining the object of his assertoric commitment. Analogous cases can be constructed involving (Audience-centered Contextual Diagonalism) and a speaker’s false but justified beliefs about the context set of her conversation.

3.6 Utility II: Mere Intentional Communication

We have seen that assertoric commitment is a different and more serious affair than the sort of commitment one undertakes when one merely intentionally communicates a proposition. This is a useful state of affairs — it allows us to maintain a fruitful distinction between asserting and speech acts like hinting, insinuating, implicating, and so forth. Speakers are able to exploit the difference between asserting and performing these less committal speech acts to a variety of ends: to get a point across politely, to maintain plausible deniability, and so forth. In order for the distinction between asserting and merely communicating to be a useful one, however, it must not turn out that all or most of the propositions a speaker is primarily interested in communicating systematically fall into the category of the merely communicated. One way for a committal practice to fail to be useful, then,

is for it to classify too many propositions as merely communicated; this will occur when a committal practice pairs utterances with propositional contents which are, in an intuitive sense, *too weak*. In this section and the next, I argue that the practice described by (Superdiagonalism) fails to be useful in this way.

Suppose Smith assertively utters “He is an embezzler,” that it is common ground between him and his audience that the person he is pointing to is Jones, and that the person he is pointing to is indeed Jones (call this scenario [PRONOUN]). On (Superdiagonalism), Smith is assertorically committed in [PRONOUN] to the proposition that whomever he is demonstrating is an embezzler; it is only because it is common ground that Smith is pointing to Jones that the proposition that Jones is an embezzler is communicated. But surely a committal practice which treats Smith’s commitment to the proposition that Jones is an embezzler in [PRONOUN] as no stronger than the commitment he would have incurred by getting the same proposition across by hinting or implicating is guilty of assimilating too much to the category of what is merely intentionally communicated.

The friend of (Superdiagonalism) will be quick to point out that even a proponent of (Objective Horizontalism) must seemingly acknowledge that, if I am talking to Jones, and if Smith approaches me and says “The person you are talking to is an embezzler,” and if it is common ground between me and Smith that the person I am talking to is Jones (call this scenario [DEFINITE]), then Smith is committed in some strong sense not only to the horizontal proposition of his utterance, but also to the proposition that Jones is an embezzler. Whatever this sort of commitment amounts to, the friend of (Superdiagonalism) will continue, why can’t she help herself to it in explaining how speakers come to be committed to the horizontal propositions of their utterances? Why not think, that is, that [DEFINITE] suggests that there is a form of commitment which is stronger than that

associated with mere intentional communication, and which is not limited to the propositions speakers actually assert?

Unfortunately for the friend of (Superdiagonalism), however, [DEFINITE] cannot do the work she needs it to do. What she needs to demonstrate is that the kind of commitment Smith undertakes in [DEFINITE] vis-a-vis the proposition that Jones is an embezzler is stronger than the commitment associated with mere intentional communication, in the sense that it does not permit speakers correctly to deny that they are liable to criticism more serious than the criticism appropriate for having intentionally communicated a falsehood. This, I argue, she cannot do.

[DEFINITE] is, of course, a case in which Smith intentionally communicates the proposition that Jones is an embezzler. The question is what feature of [DEFINITE] accounts for our intuition to the effect that Smith is *more* strongly committed to this proposition than he would have been had he communicated it by, for example, implicating it. In this connection, it is important to note that [DEFINITE] specifies that the person I am talking to is in fact Jones, so that the truth of the horizontal proposition of Smith's utterance depends on the truth of the proposition that Jones is an embezzler. Given that the person I am talking to is Jones, we might say, Smith's utterance in [DEFINITE] *de facto commits* him to the proposition that Jones is an embezzler.

To discern the contribution of this fact about de facto commitment to our intuitions concerning [DEFINITE], it will be helpful to consider a structurally related case in which the common ground between me and Smith leaves the question of who I am talking to wholly unresolved (call this scenario [DEFINITE (NO PRESUPPOSITION)]).

In [DEFINITE (NO PRESUPPOSITION)], Smith neither asserts nor intentionally

communicates the proposition that Jones is an embezzler. Nevertheless, in virtue of the fact that Jones is in fact the person I am talking to, his utterance de facto commits him to the proposition that Jones is an embezzler. Notably, much of the intuitive force of the judgment that Smith is committed to this proposition in [DEFINITE] carries over to [DEFINITE (NO PRESUPPOSITION)]: given that I am talking to Jones, Smith has nailed his flag, in some sense, to the proposition that Jones is an embezzler. This suggests that our intuitions about the two cases are tracking speakers' de facto commitments.

Yet de facto commitment has little to do with the normative status of being assertorically committed to a proposition — indeed, it is not clear that it is a normative status at all. Speakers need not even be justified in believing the propositions to which they become de facto committed in asserting. Smith, for example, might have excellent justification for the descriptive proposition he asserts in [DEFINITE (NO PRESUPPOSITION)] without knowing anything at all about Jones, and this is sufficient to render his assertion beyond reproach (unless, of course, it is false). Moreover, to the extent that speakers are blameworthy if the propositions to which they are de facto committed are false, this can be explained with reference to the fact that the propositions to which they are assertorically committed are also false; no independent normative status speakers bear to their de facto commitments needs to be invoked. Most tellingly, Smith can convincingly deny any sort of assertion-like commitment to the proposition that Jones is an embezzler: he can deny that he said, claimed, or asserted it (“I would never do such a thing without having better evidence!”), and also that he intended to communicate it (“How could I have known that the truth of what I claimed would have anything to do with Jones?”). Insofar as the friend of (Superdiagonalism) is searching for a species of commitment which does not afford speakers these

sorts of committal exit strategies, then, de facto commitment cannot serve her purposes.

Having considered the nature of de facto commitment in [DEFINITE (NO PRESUPPOSITION)], let us return to [DEFINITE]. It might be thought that the fact that in [DEFINITE] Smith both intentionally communicates and de facto commits himself to the proposition that Jones is embezzler results in a more assertion-like commitment than would result from either factor in isolation. But it is easy to show that Smith's commitment to the proposition that Jones is an embezzler in [DEFINITE] affords him some of the same committal exit strategies we encountered in our discussion of [DEFINITE (NO PRESUPPOSITION)].

Suppose we flesh out the description of [DEFINITE] so that Smith has good evidence that Jones has a doppelgänger named Schmones, that Schmones is an embezzler, and that Jones is away on a vacation. He also knows that I have no inkling of the complexities of the situation, and will interpret his utterance as concerning Jones. Smith thus intends a common sort of linguistic trick — to assert a truth while communicating a falsehood, thereby manipulating my beliefs without actually lying. In fact, however, despite his excellent evidence, Smith is wrong, and it is Jones to whom I am talking, so that what he asserts is false. If, in such a scenario, on learning that Jones is not an embezzler, I charge Smith with having claimed otherwise, he might justifiedly reply as follows: “It is true that what I claimed turned out to be false, that the truth of what I claimed turned out to depend on the actions of Jones, and indeed that I knew that Jones was no embezzler at the time of my utterance. Nevertheless, I can hardly be said to have *claimed* that Jones was an embezzler. I had evidence of the most compelling sort that the man you were speaking to was not Jones but a clever impersonator, and that this impersonator was an embezzler. I confess that I intentionally led

you to believe that Jones was an embezzler, and for that I perhaps deserve some slight recrimination, but my claim was that *the person you were talking to* was an embezzler, and nothing stronger. Indeed, I chose my words carefully so as *not* to claim that Jones was an embezzler, for I knew that proposition to be false."

The case of Smith's utterance in [DEFINITE], therefore, fails to furnish a community adopting (Superdiagonalism) with a solution to the problem of [PRONOUN]: in such a community, the genuinely assertoric commitments undertaken by speakers using context-sensitive vocabulary are exceedingly weak — weak enough that Smith, upon pointing to Jones and saying "He is an embezzler," could justifiably protest that it was unfair to hold him responsible for having claimed that Jones was an embezzler.⁸³ Whereas a community which adopts (Objective Horizontalism) can draw a useful distinction between the kind of commitment associated with the horizontal proposition of an utterance and the kind of commitment associated with propositions which are merely implicated (namely, the distinction between assertoric commitment and intentional communication), a community which adopts (Superdiagonalism) cannot: for such a community, both the horizontal proposition of an utterance and other merely implicated propositions are contents which a speaker intentionally communicates but is not otherwise com-

⁸³Of course, even on (Objective Horizontalism), Smith could argue that he should be *excused* for having claimed that Jones was an embezzler, given that he had evidence that he was pointing at Schmones. But this is a different sort of defense: one in which Smith concedes that he has committed to the horizontal proposition of his utterance.

mitted to.⁸⁴

3.7 Utility II: Secondary Commitment

Discontent with analogizing a speaker's commitment to the horizontal proposition of her utterance with the commitment involved in hinting or implicating, the friend of (Superdiagonalism) might seek to introduce a special-purpose notion of *secondary commitment*. The idea would be to associate with each assertoric utterance a set of propositions to which its speaker undertakes a commitment which, while not quite as strong as genuine assertoric commitment, is nonetheless stronger than the commitment associated with mere intentional communication.

What might a theory of secondary commitment look like? As it turns out, it is surprisingly difficult to articulate one which is compatible with (Superdiagonalism) and does not succumb to defects of the sort described in previous sections. Suppose, for example, that a speaker is secondarily committed to any propositions which come to be entailed by the context set of the conversation after it is

⁸⁴A question remains about why, in a less fanciful version of [DEFINITE] (that is, one in which there are no complexities involving Schmones), Smith's speech about claiming *that the person I was talking to is an embezzler* rather than *that Jones is an embezzler* seems less convincing. My suggestion here is that this is because we take Smith to *know* that the person I am talking to is an embezzler just in case Jones is an embezzler, and we also assume that this biconditional is in the common ground. Given these two assumptions (and a plausible closure principle for knowledge), Smith is blameless for having asserted that the person you are talking to is an embezzler if and only if he is in a position to blamelessly assert that Jones is an embezzler — he meets the evidential standard for asserting one just in case he meets the evidential standard for asserting the other, and so on. Since, given the common ground, he also communicates the same propositions regardless of what he asserts, there is a sense in which it is beside the point for him to argue that he claimed one but not the other: if he would have deserved criticism for asserting that Jones is an embezzler, he actually deserves criticism for asserting that the person I was talking to is an embezzler. This is perhaps a second sense in which Smith's commitment to the proposition that Jones is an embezzler is stronger than the commitment usually associated with hinting or implicating, though not one which will be of much comfort to the friend of (Superdiagonalism), since speakers often lack the kind of knowledge it requires, as when they are mistaken about, or suspend judgment concerning, which propositions are materially equivalent to the assertoric contents of their utterances.

updated with the assertoric content of her utterance (call this the *naive incremental* theory of secondary commitment). Now suppose that some third party (an Oracle, we can imagine) has told us that Jones is an embezzler if and only if Johnson is an arsonist, and this biconditional has become common ground. Then, in uttering “He is an embezzler,” Smith becomes secondarily responsible not only for the proposition that Jones is an embezzler, but also for the proposition that Johnson is an arsonist.

The naive incremental theory of secondary commitment thus has the following undesirable feature: even if he knows that Jones is an embezzler, Smith cannot point to Jones and assertively utter “He is an embezzler” without undertaking a secondary commitment to the proposition that Johnson is an arsonist. If he does so assert, moreover, and it is false that Johnson is an arsonist, then the Oracle can at best be censured for the falsehood of his biconditional — he escapes commitment to the proposition that Johnson is an arsonist altogether. By contributing his biconditional to the common ground, the Oracle effectively raises the normative stakes for Smith: either Smith must think it rational to undertake two secondary commitments, or he must remain silent. The Oracle, in contrast, despite playing just as important a role in the introduction of the proposition that Johnson is an arsonist into the common ground, enjoys immunity from secondary commitment to it.

A naive incremental practice of secondary commitment, then, is subject to the same kind of criticism as the practices of assertoric commitment discussed in section 5: it has structural features which discourage interlocutors from asserting and thereby impede inquiry. It is, moreover, subject to the additional criticism that it distributes secondary commitments among interlocutors in an essentially arbitrary way. For surely, if Smith has good evidence that Jones is an embezzler

but the Oracle has no evidence for his biconditional, Smith should escape criticism entirely — whereas the naive incremental theory predicts that both Smith and the Oracle should be blamed for contributing false propositions to the common ground.

Intuitively, what has gone wrong with the naive incremental theory is that it commits Smith to a contextual entailment of the content of his assertion which arises only because of the presence in the common ground of a proposition contributed by the Oracle. Perhaps, then, we could amend the naive incremental theory of secondary commitment along the following lines. In calculating the secondary commitments associated with an assertion, we follow a two-step process. First, we remove from the common ground all propositions to which individuals other than the speaker are already assertorically or secondarily committed. Then we see which propositions are newly entailed by this weakened common ground when it is updated with the content of the speaker's assertion; these are the speaker's secondary commitments. Let us call this the *sophisticated incremental* theory of secondary commitment.

The sophisticated incremental theory of secondary commitment avoids predicting that Smith undertakes a secondary commitment to the proposition that Johnson is an arsonist, since the weakened common ground relative to which Smith's assertion is evaluated does not contain the biconditional contributed by the Oracle. Unfortunately, the problem recurs in more complex cases. For now imagine that, instead of one Oracle, we have two (let us call them 'the Oracle' and 'the Moracle'). If the Oracle asserts that if Jefferson is a forger, then Jones is an embezzler if and only if Johnson is an arsonist, and the Moracle asserts that Jefferson is a forger, then the proposition that Jones is an embezzler if and only if Johnson is an arsonist enters the common ground without being an assertoric or

secondary commitment of either the Oracle or the Moracle. It follows, according to the sophisticated incremental theory, that Smith will once again become secondarily committed to the proposition that Johnson is an arsonist if he subsequently asserts that Jones is an embezzler.

It is also worth pointing out that the two theories of secondary commitment just surveyed both involve the friend of (Superdiagonalism) in problems very similar to those faced by the contextual diagonalist. Since the conjunction of the propositions in the common ground with the content of a speaker's assertion is always among the propositions newly entailed by the common ground, the naive incremental theory has speakers always undertaking secondary commitments which are at least as strong as the context set. This is problematic given how commonly the context set evolves non-eliminatively. The situation is less extreme for the sophisticated incremental theory, but it still has speakers undertaking secondary commitments to, for example, the whole body of background beliefs brought to the conversation before any interlocutor has uttered anything at all.

So appealing to a notion of secondary commitment, at least in one of the forms just described, results in the overgeneration of predictions of assertoric commitment and the corresponding overgeneration of blame when things go wrong. This leaves the friend of (Superdiagonalism) in a difficult position: either stick with assertoric commitment, in which case speakers are committed to too little, or opt for secondary commitment, in which case speakers are committed to too much.⁸⁵

⁸⁵Especially committed proponents of (Superdiagonalism) may maintain that, even if what I have claimed so far is true, there must be *some* other notion of commitment which will serve their purposes. The burden is on them, however, to develop a theory of the relevant notion of commitment. In the absence of such a theory, it is difficult to resist the conclusion that a superdiagonalist committal practice struggles with respect to utility.

3.8 Modal Coherence

If the propositional contents of utterances can be associated with distinct forces (conventionally associated, for example, with distinct sentential moods), a practice of mapping utterances to propositional contents will be easier for interlocutors to use if it can be applied regardless of the mood in which a sentence is uttered. Versions of horizontalism fare better in this regard than their competitors. This is especially clear when we consider moods which require speakers to look outside the context set, such as the optative. For suppose there is a distinctive force corresponding to optative constructions, which we may call the *desirous force*. In English, the optative construction ‘Would that ϕ !’ is plausibly associated with just such a desirous force: it is used not to assert that the speaker desires that ϕ , but rather to express this desire directly.⁸⁶ But if we adopt a plausible theory of the felicity conditions for such constructions and assume that utterances in the desirous mood are mapped to propositions in the same way as utterances in the indicative mood, (Superdiagonalism) and the two forms of contextual diagonalism lead to significant problems. Suppose for illustration that ϕ is ‘He is here’. I will consider two theories of the felicity conditions for an utterance of ‘Would that ϕ !’ and show that both make unintuitive predictions.

⁸⁶In treating the English optative construction as a force marker, I assume that it does not interact compositionally with its prejacent to change which proposition is expressed; that is, I assume that the proposition which the grammar assigns to ϕ is the same as the proposition which the grammar assigns to ‘Would that ϕ !’, the difference between the two being entirely a matter of what speech act is conventionally performed by uttering them. I consider this assumption plausible. Nevertheless, even if theoretical considerations ultimately suggest that it is false, it suffices for my purposes to note, first, that there may be other languages in which optative constructions are genuine force markers in the sense described, and, second, that a community *could* introduce some construction which worked in the suggested way. The argument would then be that the introduction of such a construction would be useful only for a community with a horizontalist committal practice, and that it is an advantage for a committal practice to be easily extensible in this way.

First, consider the theory that $\lceil \text{Would that } \phi! \rceil$ is felicitous for a speaker just in case (i) it is common ground that the assertoric content of ϕ is false, and (ii) all of the best possible worlds (according to that speaker) are worlds where the assertoric content of ϕ is true.⁸⁷ Then, if ϕ is ‘He is here’ and (Superdiagonalism) is true, we predict that ‘Would that he were here!’ is felicitous just in case (i) it is common ground that the person the speaker has in mind is not at the location in which the utterance is produced, and (ii) all of the best possible worlds according to the speaker are worlds in which the person she has in mind when producing the utterance is at the location in which the utterance is produced. But for most natural preference orderings on worlds, this account will yield the unappealing consequence that it is *almost never* appropriate for speakers to utter the optative construction under consideration. Suppose Devin loves Kevin more than any other and loathes Evan. Neither Kevin nor Evan is present. If the conversation turns to Kevin, and Devin exclaims “Would that he were here!” intending to refer to Kevin, then her utterance is felicitous, according to our first theory, just in case (i) it is presupposed that the superdiagonal proposition of ‘He is here’ is false, and (ii) all of the best worlds according to Devin are ones in which the superdiagonal proposition of ‘He is here’ is true. The first of these conditions is satisfied because, we may imagine, it is common ground that Devin intends to refer to Kevin and that Kevin is not present. But the second condition is not plausibly satisfied because, although the worlds which are subjectively best for Devin are all ones in which Kevin is present, they are not all ones in which the superdiagonal proposition of ‘He is here’ is true: witness the world where ‘He’

⁸⁷What if, according to the speaker, every world is such that there is a world which is better than it? Then instead of speaking of “all of the best possible worlds,” we can require that every world w be such that (i) there is a better world at which the assertoric content of ϕ is true, and (ii) every better world is such that the assertoric content of ϕ is true. (Cf. Kratzer 2012, 39.)

refers to Evan and Evan is not near Devin.⁸⁸ So Devin's utterance of "Would that he were here!" comes out as inappropriate despite her intention to refer to Kevin and sincere desire for his proximity.

Second, consider the theory that $\lceil \text{Would that } \phi! \rceil$ is felicitous for a speaker just in case (i) it is common ground that the assertoric content of ϕ is false, and (ii) the assertoric content of ϕ exceeds some threshold of desirability according to the speaker. We can either take the relevant level of desirability to be context-invariant, for example by requiring that the expected utility of the assertoric content of ϕ exceeds the expected utility of the assertoric content of $\lceil \neg\phi \rceil$ when calculated using something like the speaker's Bayesian ur-priors, or, following Grosz (2012), we can take it to be contextually determined. Either way, if (Superdiagonalism) is true, the felicity of an utterance of "Would that he were here!" will depend not only on the utility of various possible scenarios in which he is here, but also on the utility of various possible scenarios in which the speaker has a different individual in mind and that individual is here. Thus if the speaker intends to refer to Kevin in uttering "Would that he were here!", then even if she assigns great utility to the worlds in which Kevin is nearby, her utterance may nonetheless be infelicitous because she assigns great disutility to certain worlds in which the referent of 'He' is Evan and Evan is nearby. It is difficult to see what purpose would be served by a grammaticalized optative construction if this were its meaning.

In addition to the problems just mentioned, both theories, when combined with (Superdiagonalism), have the unfortunate consequence that the referential

⁸⁸There is a question about whether all the worlds which are subjectively best for Devin are ones at which she utters "He is here"; if not, then the relevant superdiagonal proposition will fail to be defined over the set of best worlds, yielding once again the prediction that sentences of the form $\lceil \text{Would that } \phi! \rceil$ are *almost never* felicitously assertible.

intentions of a speaker who utters $\lceil \text{Would that } \phi! \rceil$ contribute only to its presuppositional profile and not to its at-issue content. Thus all utterances of “Would that he were here!” express a positive attitude toward the *same* superdiagonal proposition. Whatever benefits accrue to a linguistic community when it incorporates indexicals into its language, (Superdiagonalism) thus ensures that those benefits are not available when it comes to optative constructions.⁸⁹

The problems faced by contextual diagonalism when it comes to modal coherence are even more severe. For, according to contextual diagonalism, the only worlds where the assertoric content of an utterance is true are worlds in what the relevant individual or individuals take to be the context set. But, on either theory of the semantics of $\lceil \text{Would that } \phi! \rceil$, a speaker may felicitously employ that construction only if it is presupposed that the assertoric content of ϕ is false — that is, only if there are no worlds in the context set at which the assertoric content of ϕ is true. So, on the first semantics of $\lceil \text{Would that } \phi! \rceil$, as long as the set of best possible worlds is nonempty, $\lceil \text{Would that } \phi! \rceil$ will be unassertable. For each best world will either be outside the context set, or it will be inside the context set. The nature of contextual diagonalism ensures that the assertoric content of ϕ is false at all worlds in the former category. With respect to the latter category, either the assertoric content of ϕ is true at some best worlds in the context set, in which case the presupposition that the assertoric content of ϕ is false is not satisfied, and the utterance will fail condition (i) for felicitous assertion, or the assertoric content of

⁸⁹Indeed, this observation follows from the more general point that, on (Superdiagonalism), the assertoric content of an utterance is sensitive only to the sentence uttered and not to the extralinguistic context in which it is uttered: an utterance of a given sentence always counts as assertion of the same proposition, regardless of context. This is not to say that (Superdiagonalism) is incompatible with grammatical context sensitivity — the total pattern of the dependence of grammatical content on extralinguistic context is what determines the superdiagonal proposition associated with a given utterance, so that metasemantic differences in the underlying grammar correspond to the assignment of different superdiagonal propositions to utterances.

ϕ is false at all best worlds in the context set, in which case the utterance will fail condition (ii) for felicitous assertion.

Similarly, on the second semantics for \ulcorner Would that $\phi!$ \urcorner , either the assertoric content of ϕ will be true at some worlds inside the context set, in which case the presupposition that it is false will not be satisfied, or it will be false throughout the context set, in which case, given the nature of contextual diagonalism, it must be false at all possible worlds; in other words, it must be the contradictory proposition. Even granting that the contradictory proposition can exceed the relevant threshold of desirability, we get the unpalatable consequence that \ulcorner Would that $\phi!$ \urcorner is assertable just in case \ulcorner Would that $\psi!$ \urcorner is assertable for arbitrary ϕ and ψ .

3.9 Partiality

The committal practices discussed so far have the following feature in common: the contents they assign to assertoric utterances determine *total* functions from the set of worlds where those utterances exist to the set of truth values.⁹⁰ There are possible committal practices, however, which are not helpfully modeled using contents which determine total functions of this sort. For example, we can imagine:

(Nondefective Partial Diagonalism): For all conversations c , utterances u , speakers s , and worlds w : If c is nondefective at w , and if s assertively utters u at w , then the object of s 's assertoric commitment in uttering u is the partial function from worlds to truth-values which agrees with the superdiagonal

⁹⁰This is not to say that they assign contents to all utterances at all worlds where those utterances exist—we saw in section 4 that this is not the case.

proposition expressed by u over the context set of c in w and is undefined elsewhere.

(Nondefective Partial Diagonalism) differs from (Superdiagonalism) in that, according to the former but not the latter, a speaker undertakes no commitment concerning the truth of the superdiagonal proposition of her utterance at worlds outside the context set of her conversation. (Nondefective Partial Diagonalism) resembles a position defended by Stalnaker (2014):

“...the definition I have given for the diagonal proposition (the assertoric content) determines only a partial proposition: as specified so far, the function from possible worlds to truth-values is defined only relative to the domain of possible worlds in the context set.” (Stalnaker 2014, 221)

(Nondefective Partial Diagonalism) arguably fares better than (Superdiagonalism) with respect to considerations of utility. If Smith utters ‘He is an embezzler’ in a conversation the context set of which determines that ‘He’ refers to Jones, (Nondefective Partial Diagonalism) does not erroneously predict that Smith is primarily committed to the proposition that whomever he has in mind is an embezzler — at least, not *as opposed to* the proposition that Jones is an embezzler, since these two propositions are true at the same worlds in the context set, and Smith’s commitment is undefined elsewhere.

(Nondefective Partial Diagonalism) is, however, problematic as a possible committal practice. Like the other nondefective views discussed above, it needs to be generalized so as to be either speaker-centered or audience-centered. Since its audience-centered generalization fails to be general in much the same way as (Audience-centered Contextual Diagonalism), we may restrict our attention to its speaker-centered generalization:

(Speaker-centered Partial Diagonalism): For all conversations c , utterances u , speakers s , and worlds w : If s assertively utters u at w , then the object of s 's assertoric commitment in uttering u is the partial function from worlds to truth-values which agrees with the superdiagonal proposition expressed by u over what, at w , s takes the context set of c to be for the purposes of the conversation, and which is undefined elsewhere.

Even (Speaker-centered Partial Diagonalism) is problematic, however. For in cases of non-eliminative context-set evolution where interlocutors who have been presupposing a proposition discover that it is false and come to presuppose its negation, (Speaker-centered Partial Diagonalism) liberates speakers from all prior assertoric commitments. Even when speakers merely cease to presuppose a proposition without presupposing its negation, since the context set after the change is a proper superset of the context set prior to the change, speakers have undefined assertoric commitments at some worlds in the context set.

Stalnaker is aware of this problem. A natural solution would, of course, be to fall back on (Superdiagonalism). The solution Stalnaker proposes, however, is to allow the assertoric content of utterances to be "extended" outside the context set in cases where "the context (the common ground)... provide[s] a natural extension" (Stalnaker 2014, 221). Stalnaker provides no systematic theory of the circumstances under which this will be the case or the ways in which these circumstances determine extensions, but the cases he offers as examples suggest that extension is possible if and only if it is common ground that an utterance expresses a particular horizontal proposition, in which case its assertoric content is identified with this horizontal proposition. Stalnaker's proposal (modified to reflect the preferability of speaker-centered views) thus amounts to:

(Stalnakerian Disjunctivism): For all conversations c , utterances u , speakers s , and worlds w : If s assertively utters u at w , then the object of s 's assertoric commitment in uttering u is the partial function from worlds to truth-values which agrees with the superdiagonal proposition expressed by u over what, at w , s takes the context set of c to be for the purposes of the conversation, and which is undefined elsewhere, unless there is a unique proposition p such that it is common ground in c at w that p is the horizontal proposition expressed by u , in which case the object of s 's assertoric commitment is p .⁹¹

(Stalnakerian Disjunctivism) goes some way toward solving the problem of non-eliminative evolution of the context set, but it does not go far enough. Suppose Smith points at a man and assertively utters "He is an embezzler." Suppose also that it is common ground that the man in question is either Jones or Johnson. Then there is no proposition p such that it is presupposed that p is the horizontal proposition of Smith's utterance, and he asserts the partial diagonal proposition of his utterance. But now if we discover that we have been presupposing some trivial falsehood — that Pluto is a planet, say — and the context set evolves non-eliminatively, Smith will be entirely free from assertoric commitments. Given the commonness of uncertainty about the horizontal propositions expressed by utterances containing context-sensitive vocabulary, the possibility of Stalnakerian extension to a non-partial content will go unrealized too often for (Stalnakerian

⁹¹John Hawthorne points out that it seems in spirit of (Stalnakerian Disjunctivism) to allow for a certain kind of extension even when there fails to be a unique proposition which it is common ground is the horizontal proposition expressed by an utterance. For example, an utterance of a conjunction could be such that both conjuncts contain context-sensitive expressions, but it is common ground that the expression in the first conjunct takes a certain value, whereas the value taken by the expression in the second differs across the context set. In such a case, it seems in the spirit of (Stalnakerian Disjunctivism) to say that the assertoric content of the utterance is defined at worlds outside the context set, and that it is at least as strong as the first conjunct. We may, however, safely ignore this complication in what follows.

Disjunctivism) to describe a useful committal practice.

3.10 Conclusion

Our exploration of the considerations favoring various committal practices has revealed a clear frontrunner: (Objective Horizontalism). Considerations of generality tell against practices which assign contents to utterances only when the conversation is nondefective, as well as against audience-centered practices and, to a lesser extent, certain speaker-centered practices. The problem of non-eliminative context-set evolution affects contextual diagonalism as well as Stalnaker's partial committal practices. Unintentional liability is a further problem for audience-centered practices. (Superdiagonalism), meanwhile, struggles with overly weak assertoric commitments and the problem of modal coherence. (Objective Horizontalism) thus emerges as the least problematic, and correspondingly most useful, committal practice. This result gives us some reason to expect that we ourselves have adopted a horizontalist practice, thus adding some weight to the case against the diagonalist answer to the philosopher's puzzle.⁹²

⁹²Thanks are due to Andy Egan, John Hawthorne, Jeff King, Jeff Russell, and two anonymous referees for their helpful comments on earlier drafts of this paper.

Bibliography

- Almotahari, M., and Glick, E. (2010). 'Context, Content, and Epistemic Transparency'. *Mind* 119: 1067–1086.
- Brandom, R. (1983). 'Asserting'. *Noûs* 17(4): 637–650.
- Cappelen, H. & Hawthorne, J. (2009). *Relativism and Monadic Truth*. Oxford University Press.
- Cappelen, H. & Lepore, E. (1997). 'On an Alleged Connection Between Indirect Speech and the Theory of Meaning'. *Mind & Language* 12: 278–296.
- Collins, P.C. (1991). *Cleft and Pseudo-Cleft Constructions in English*. Routledge.
- Conan Doyle, A. (1902). *The Hound of the Baskervilles*. McClure, Phillips, & Company.
- Dummett, M. (1991). *The Logical Basis of Metaphysics*. Harvard University Press.
- Gibbard, A. (Unpublished Manuscript). 'Propositions: Not What We Believe?'
- Glanzberg, M. (2011). 'More on Operators and Tense'. *Analysis* 71: 112–123.
- Grosz, P. G. (2012). *On the Grammar of Optative Constructions*. John Benjamins Publishing Company.
- Hawthorne, J., & Magidor, O. (2009). 'Assertion, Context, and Epistemic Accessibility'. *Mind* 118: 377–397.
- Hawthorne, J., & Magidor, O. (2010). 'Assertion and Epistemic Opacity'. *Mind*, 119: 1087–1105.
- Kanger, S., & Öhman, S. (eds) (1980). *Philosophy and Grammar*. D. Reidel Publishing Company.
- King, J. C. (2003). 'Tense, Modality, and Semantic Values'. *Philosophical Perspectives* 17: 195–245.
- King, J. C. (2014). 'The Metasemantics of Contextual Sensitivity'. In Burgess, A. & Sherman, B. (eds) *Metasemantics: New Essays on the Foundations of Meaning*, pp. 97–118. Oxford University Press.

- King, J. C. (2018). 'Strong Contextual Felicity and Felicitous Underspecification'. *Philosophy and Phenomenological Research* 97(3): 631–657.
- Kirk-Giannini, C. D. (2018). 'Uniformity Motivated'. *Linguistics and Philosophy* 41: 665–684.
- Kratzer, A. (2012). *Modals and Conditionals*. Oxford University Press.
- Lewis, D. (1975). 'Languages and Language'. In Keith Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*, Vol. 7, pp. 3–35. University of Minnesota Press.
- Lewis, D. (1980). 'Index, Context, and Content'. In Kanger and Öhman 1980, pp. 79–100.
- Ninan, D. (2010). 'Semantics and the Objects of Assertion'. *Linguistics and Philosophy* 33: 355–380.
- Perry, J. (1986). 'Perception, Action, and the Structure of Believing'. In R.E. Grandy & R. Warner (eds.) *Philosophical Grounds of Rationality: Intentions, Categories, Ends*, Clarendon Press, pp. 333–61.
- Predelli, S. (1998). 'I Am Not Here Now'. *Analysis* 58: 107–115.
- Quine, W. V. (1940). *Mathematical Logic*. Harvard University Press.
- Rabern, B. (2012). 'Against the Identification of Assertoric Content with Compositional Value'. *Synthese* 189(1): 75–96.
- Schlenker, P. (2011). 'Indexicality and *De Se* Reports'. In K. von Stechow, C. Maienborn, & P. Portner (Eds.), *Semantics: An International Handbook of Natural Language Meaning*, Vol. 2, pp. 1561–1604. De Gruyter Mouton.
- Searle, J. (1976). 'A Classification of Illocutionary Acts'. *Language in Society* 5(1): 1–23.
- Sidelle, A. (1991). 'The Answering Machine Paradox'. *Canadian Journal of Philosophy* 21: 525–539.
- Soames, S. (2002). *Beyond Rigidity*. Oxford University Press.
- Soames, S. (2005). 'Naming and Asserting.' In Szabo, Z. G. (ed.) *Semantics vs Pragmatics*, pp. 356–382. Oxford University Press.

- Stalnaker, R. (1970). 'Pragmatics'. *Synthese* 22: 272–289.
- Stalnaker, R. (1978). 'Assertion'. In his *Context and Content*, Oxford University Press, 1999, pp. 78–95. Originally published in *Syntax and Semantics* 9: 315–332.
- Stalnaker, R. (1999). *Context and Content*. Oxford: Oxford University Press.
- Stalnaker, R. (2004). 'Assertion Revisited'. *Philosophical Studies* 118: 299–322.
- Stalnaker, R. (2009). 'On Hawthorne and Magidor on Assertion, Context, and Epistemic Accessibility'. *Mind* 118: 399–409.
- Stalnaker, R. (2014). *Context*. Oxford University Press.
- Stalnaker, R. (2018). 'Dynamic Pragmatics, Static Semantics'. In D. Fogal, D. Harris, & M. Moss (eds.) *New Work on Speech Acts*, pp. 384–399. Oxford University Press.
- von Stechow, A. & Zimmermann, T. E. (2005). 'A Problem for a Compositional Treatment of *De Re* Attitudes'. In G. Carlson & F. J. Pelletier (Eds.), *Reference and Quantification: The Partee Effect*, pp. 207–228. CSLI Publications.
- Stojnić, U. (2017). 'On the Connection between Semantic Content and the Objects of Assertion'. *Philosophical Topics* 45(2): 163–179.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.