


**PERSPECTIVE IN CONTEXT
RELATIVE TRUTH, KNOWLEDGE, AND THE FIRST PERSON**

Dirk Kindermann

**A Thesis Submitted for the Degree of PhD
at the
University of St Andrews**



2012

**Full metadata for this item is available in
Research@StAndrews:FullText
at:**

<http://research-repository.st-andrews.ac.uk/>

**Please use this identifier to cite or link to this item:
<http://hdl.handle.net/10023/3164>**

This item is protected by original copyright

Perspective in Context

Relative Truth, Knowledge, and The First Person

Dirk Kindermann

*A Thesis Submitted for the Degree of PhD
University of St. Andrews, August 2012*

I, Dirk Kindermann, hereby certify that this thesis, which is approximately 80.000 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in February 2008 and as a candidate for the degree of Doctor of Philosophy in February 2008; the higher study for which this is a record was carried out in the University of St Andrews between 2008 and 2012.

27 August 2012

Signature of candidate

I, Patrick Greenough, hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Doctor of Philosophy in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

27 August 2012

Signature of supervisor

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any *bona fide* library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the electronic publication of this thesis: Access to the printed copy but embargo of the electronic publication of the thesis for a period of 3 years on the ground that publication would preclude future publication.

27 August 2012

Signature of supervisor

Signature of candidate

ABSTRACT

This dissertation is about the nature of perspectival thoughts and the context-sensitivity of the language used to express them. It focuses on two kinds of perspectival thoughts: ‘subjective’ evaluative thoughts about matters of personal taste, such as *Beetroot is delicious* or *Skydiving is fun*, and first-personal or *de se* thoughts about oneself, such as *I am hungry* or *I have been fooled*. The dissertation is a defence of a novel form of relativism about truth – the idea that the truth of some (but not all) perspectival thought and talk is relative to the perspective of an evaluating subject or group.

In Part I, I argue that the realm of ‘subjective’ evaluative thought and talk whose truth is perspective-relative includes attributions of knowledge of the form ‘S knows that *p*.’ Following a brief introduction (chapter 1), chapter 2 presents a new, error-theoretic objection against relativism about knowledge attributions. The case for relativism regarding knowledge attributions rests on the claim that relativism is the only view that explains all of the empirical data from speakers’ use of the word ‘know’ without recourse to an error theory. In chapter 2, I show that the relativist can only account for sceptical paradoxes and ordinary epistemic closure puzzles if she attributes a problematic form of semantic blindness to speakers. However, in chapter 3 I show that all major competitor theories – forms of invariantism and contextualism – are subject to equally serious error-theoretic objections. This raises the following fundamental question for empirical theorising about the meaning of natural language expressions: If error attributions are ubiquitous, by which criteria do we evaluate and compare the force of error-theoretic objections and the plausibility of error attributions? I provide a number of criteria and argue that they give us reason to think that relativism’s error attributions are more plausible than those of its competitors.

In Part II, I develop a novel unified account of the content and communication of perspectival thoughts. Many relativists regarding ‘subjective’ thoughts and Lewisians about *de se* thoughts endorse a view of belief as self-location. In chapter 4, I argue that the self-location view of belief is in conflict with the received picture of linguistic communication, which understands communication as the transmission of information from speaker’s head to hearer’s head. I argue that understanding mental content and speech act content in terms of *sequenced worlds* allows a reconciliation of these views. On the view I advocate, content is modelled as a set of sequenced worlds – possible worlds ‘centered’ on a group of individuals inhabiting the world at some time. Intuit-

ively, a sequenced world is a way a group of people may be. I develop a Stalnakerian model of communication based on sequenced worlds content, and I provide a suitable semantics for personal pronouns and predicates of personal taste. In chapter 5, I show that one of the advantages of this model is its compatibility with both nonindexical contextualism and truth relativism about taste. I argue in chapters 5 and 6 that the empirical data from eavesdropping, retraction, and disagreement cases supports a relativist completion of the model, and I show in detail how to account for these phenomena on the sequenced worlds view.

ACKNOWLEDGEMENTS

My greatest debt in writing this dissertation is to my main supervisor Patrick Greenough. Always encouraging, patient, and generous with his intellectual energy, Patrick has been an ideal supervisor and the model of a constructive and unprejudiced philosophical interlocutor. The traces of countless hours of inspiring conversation can be found on every page of this dissertation.

I would like to express my deepest gratitude to Jonathan Schaffer. His advice and comments, before, during and after acting as my second supervisor, have been tremendously important. I would also like to thank my other supervisors: Crispin Wright for early guidance, and Brian Weatherson for invaluable feedback especially at the late stages of the project.

During my time at the Arché Philosophical Research Centre in St Andrews, I benefitted immensely from exchanges with a great number of people. In particular, I would like to thank Derek Ball, Mark Bowker, Jochen Briesen, Björn Brodowski, Jessica Brown, Herman Cappelen, Colin Caret, Yuri Cath, Laura Celani, Josh Clarkson, Michael De, Paul Dimmock, Dylan Dodd, Noah Friedman-Biglin, Ephraim Glick, Steven Hall, Katherine Hawley, Andri Hjálmarsson, Thomas Hodgson, Nick Hughes, Ole Hjortland, Jonathan Ichikawa, Frederique Janssen-Lauret, Bruno Jacinto, Spencer Johnston, Ira Kiourti, Julia Langkau, Martin Lipman, Federico Luzzi, Toby Meadows, Julien Murzi, Andrea Onofri, Nikolaj Pedersen, Walter Pedriali, Laura Porro, Graham Priest, Simon Prosser, Stephen Read, François Recanati, Marcus Rossberg, Anders Schoubye, Daniele Sgaravatti, Stewart Shapiro, Martin Smith, Jason Stanley, Rachel Sterken, Andreas Stokke, Margot Strohminger, Paul Sweeney, Josh Thorpe, and Elia Zardini. Special thanks go to the members of Arché's Contextualism & Relativism research project and the Arché Semantics Group. Torfinn Huvenes, in particular, has been an invaluable philosophical companion.

Some of the material in this dissertation was presented at conferences and workshops in Edinburgh and Miami, at the Australasian Association of Philosophy Conference 2010 in Sydney, the Wolheim Society at the University of California at Berkeley, the conference 'Meaning, Context and Implicit Content' in Cerisy, 'Contexts, Perspectives and Relative Truth' in Bonn, at SPE4 in Bochum, the LOGOS workshop 'Semantic Content and Context Dependence' in Barcelona, and at Arché's 'Contextualism and Relativism Workshop II.' I thank the audiences at all these occasions. I

am also indebted to Josh Dever, Andy Egan, Chris Kelp, Martin Montminy, Hazel Pearson, and Stephan Torre for their helpful comments and feedback.

In 2010, I visited the philosophy departments at the Australian National University and at UC Berkeley. At the former, I benefitted greatly from conversations with Brian Rabern and Clas Weber, amongst others. At UC Berkeley, I had the fortune to meet a number of exceptional people, who made my visit extremely productive and enjoyable. I would like to thank John MacFarlane and Seth Yalcin, who were remarkably generous with their time, both in personal meetings and in the weekly sessions of the newly formed semantics reading group. I also profited from many interactions with Justin Bledin, Mike Caie, Jessica Gelber, and Mike Rieppel.

Dilip Ninan deserves very special thanks. I learned more from talking to Dilip than I could possibly acknowledge, and his influence on the ideas in Part II of this dissertation will be obvious to anyone who knows his work.

I gratefully acknowledge the financial support of the German Academic Exchange Service (DAAD), the Arts and Humanities Research Council, and the University of St Andrews.

I would like to thank Liz Shannon for her love and friendship, and for brightening up my life.

I am very fortunate to have spent the last decade thinking about philosophy. Without the inspiration and encouragement of Michael Eskin, I probably would not have started. Without the unquestioning and loving support of my parents, Irmgard and Günther Kindermann, none of it would have been possible. This dissertation is dedicated to my parents.

Dirk Kindermann
St Andrews
May 2012

CONTENTS

Abstract	v
Acknowledgements	vii
Contents	ix
I Introduction	I
I Knowledge	5
2 Relativism and Semantic Blindness	7
2.1 Introduction	7
2.2 Contextualism and the Case for Relativism	8
2.3 An Error-theoretic Objection against Relativism	12
2.4 Relativist Replies	14
2.5 Sceptical Paradox and Index-Blindness	20
2.6 Context Confusion without Index-Blindness?	25
2.7 Epistemic Closure Puzzles	26
2.8 Conclusion	28
3 Knowledge Attributions and Speaker Error	29
3.1 Introduction	29
3.2 Nonindexical Contextualism and Use-Index-Blindness	31
3.3 Moderate Insensitive Invariantism and Pragmatic Blindness	35
3.4 Sceptical Insensitive Invariantism and Pragmatic Blindness	43
3.5 Interest-Relative Invariantism and Supervenience-Blindness	49
3.6 The Views on the Scoreboard	57
3.7 On the Evaluation and Comparison of Blindness Attributions	58
3.8 Relativism Rescued?	63
3.9 Conclusion	66
II <i>De Se</i> and Subjective Attitudes	67
4 Perspective in Context: Sequenced Worlds	69
4.1 Introduction	69
4.2 Centered Content	70

4.2.1	Attitudes <i>De Se</i>	70
4.2.2	Attitudes Towards ‘Subjective’ Matters	73
4.2.3	Egocentricity and Other Similarities	73
4.3	The Lockean Picture of Communication	75
4.4	The Conflict	76
4.5	Sequenced Worlds Content	77
4.6	Belief in Context	80
4.7	Assertion and the Common Ground	82
4.7.1	Common Ground	82
4.7.2	Assertion and Acceptance	84
4.8	A Semantic Proposal	88
4.9	Speech Act Content	92
4.10	Competing Solutions	97
4.10.1	Centered Belief Content without Lockean Communication	98
4.10.2	Lockean Communication without Centered Belief Content	100
4.11	Objections	101
4.11.1	Belief and Sequenced Worlds Content	101
4.11.2	On Believing in Context	102
4.11.3	Objecting on Someone Else’s Behalf	103
4.11.4	Sameness of Content	103
4.11.5	‘Tasty to us’	104
4.11.6	Negated Taste Claims	105
4.12	Conclusion	107
4.A	Diagonalisation and Context-Sensitivity	108
5	Sequenced Worlds, Eavesdropping, Retraction	111
5.1	Introduction	111
5.2	Contextualism and Relativism	112
5.3	On Two Implementations of Relativism	116
5.4	Eavesdroppers	118
5.5	Norms of Assertion	122
5.6	The Objection from ‘tasty to us’	128
5.7	Retraction	130
5.7.1	The Data and Its Interpretation	131
5.7.2	The Common Ground and the Passage of Time	134
5.7.3	Retraction and the Common Ground	137
5.7.4	Norms for Retraction	146
5.8	Conclusion	151
6	Sequenced Worlds, Disagreement, Truth Ascriptions	153
6.1	Introduction	153
6.2	Disagreement	153
6.2.1	Introduction	153
6.2.2	Disagreement on the Sequenced Worlds Model	156

6.2.3	A Note on Transworld Disagreement	162
6.2.4	Faultlessness	164
6.2.5	Nonindexical Contextualism and Disagreement	168
6.3	On the Dynamics of (Dis)agreement	169
6.4	Believing Alike, Agreement, and Cognitive Significance	174
6.5	Truth Ascriptions, Tense, and Predicates of Personal Taste	179
6.5.1	Monadic 'true'	180
6.5.2	Cross-temporal Truth Ascriptions	183
6.5.3	Tense and Predicates of Personal Taste	187
6.6	Modality, Transworld Assessments, and Predicates of Personal Taste .	194
6.7	On Content	196
6.8	Conclusion	198
Bibliography		201

INTRODUCTION

Truth is a virtue we want our thought and talk to possess. It is good to believe what is true. It helps our actions if we get the facts right. What we believe and assert thus trades in truth and falsity – it is truth-apt. This assumption of truth-aptness pervades contemporary philosophy, and it is the heart and soul of one of the most fruitful projects in the study of meaning – the project of *truth-conditional semantics*.

Yet some of the things we believe and assert seem to resist classification into the true and the false. Is it a truth that twelve-tone music is bizarre, that skydiving is fun, or that beetroot is tasty? If this question has an odd ring, it is perhaps because Truth with a capital ‘T’ seems to be the wrong dimension of evaluation for these claims. They seem to be ‘subjective’ – the assessments of a subject rather than judgments of objective fact. At the same time, giving up the idea that truth and falsity are central features of the objects of our beliefs and assertions is a radical departure from well-established theorising. Moreover, in some sense we do want to call even our most subjective claims true or false, for instance when we express our agreement by remarking, ‘That’s true.’

Relativism about truth promises to give us both. It secures a place for evaluative sentences and their contents among the truth-apt, declarative sentences of our language for which truth-conditional semantics can be given, and it guarantees that evaluative thought and talk is subjective rather than objectively true or false. As I understand it, *relativism about truth*, for a piece of language and thought, is roughly the idea that a single assertion or belief of this kind can be true relative to one perspective and false relative to another.

The relativism that is the topic of this dissertation is an empirical thesis about how we *in fact* think and talk when we make subjective, evaluative judgments. The core of the view is a semantic and pragmatic account of our use of expressions such as ‘bizarre,’ ‘fun,’ and ‘tasty,’ which allow us to talk about matters of personal taste. The main motivation for relativism comes from facts about our use of such expressions, notably from cases involving disagreement, eavesdropping, or retraction. Such facts about use, the relativist maintains, are straightforwardly predicted by relativism. In contrast, none of the competing theories regarding the semantics and pragmatics of these expressions makes all and only correct predictions in these cases. Hence, relativism should be preferred on empirical grounds.

In what follows, I will largely take it for granted that the empirical data gives us suf-

ficient motivation to take seriously the idea that the truth of expressions about matters of personal taste is relative to the evaluator's perspective. Not everyone will agree with this starting point. Many think that alternative views such as contextualism, objectivism, or expressivism regarding expressions of personal taste receive enough empirical support to spare us the need for an idea as radical as the relativisation of truth. Others object that the very idea of relative truth is incoherent. Although I put these reasonable concerns to one side, I do not intend to dismiss them. But I submit that the recent wave of pioneering work on relative truth by philosophers and linguists such as Andy Egan, Max Kölbel, Peter Laserson, John MacFarlane, Mark Richard, Tamina Stephenson and others has done enough to justify the dual assumption of conceptual coherence and significant empirical support as the point of departure for the present investigation.

I hope to convince the reader of two points. First, the prospects are bright for extending this relativism to other areas of thought and talk, areas whose subject matter has a less 'subjective' semblance; in particular to our ascriptions of knowledge of the form '*S* knows [does not know] that *p*'. Second, the prospects are bright for integrating relativism with an account of first-personal thought and talk, which crucially also depends on the subject's perspective.

The dissertation falls into two parts. Part I explores the empirical prospects for providing a relativist semantics for knowledge attributions, an area of thought and talk in which dependence on the evaluator's perspective is far from obvious. Chapter 3 presents a problem for relativism. The case for relativism about knowledge attributions as championed by MacFarlane (2005a, 2011b) rests on the argument that relativism is the only theory that makes the correct predictions for all of the empirical data and is to be preferred on these grounds to other views, all of which must resort either to unsatisfactory pragmatic explanations of the troublesome data or to the attribution of systematic error to speakers. Chapter 3 argues that this way of motivating relativism fails. Sceptical paradoxes and epistemic closure puzzles present data that the relativist can only explain in a satisfactory way if she attributes an implausible form of semantic blindness to ordinary speakers. Relativism is thus subject to an error-theoretic objection just like its contextualist and invariantist competitors.¹

Error-theoretic objections are ubiquitous in the debate on the semantics and pragmatics of knowledge attributions. Their shape is simple. First, it is shown that a theory makes the wrong predictions for a particular set of data from speakers' use of 'know.' Then it is claimed that the theory's best account of the data involves the attribution of systematic error to speakers. Finally, it is claimed that this error attribution is implausible. Given their ubiquity and simplicity, it is surprising that it is far from clear what

¹A version of chapter 3 is forthcoming as 'Relativism, Sceptical Paradox, and Semantic Blindness' in *Philosophical Studies*; cf. Kindermann (forthcoming).

makes an error attribution implausible. But distinguishing different error-theoretic objections in terms of their force is crucial in the dialectical situation the debate is facing. If every view must have recourse to attributions of systematic error, we need to evaluate and compare the given error theories if we want to come to an adequately reasoned decision between rival views. Chapter 3 does just this. I start by establishing in some detail for which data points invariantist and contextualist theories need to posit error, even where a pragmatic account of the data is provided. In so doing, I distinguish between different kinds of semantic, pragmatic, and metaphysical error. This allows us to put the views on the scoreboard, where we can read off and compare their success in empirical predictions. I then suggest adequacy conditions for error theories by identifying a number of criteria for the evaluation of kinds of error and their attribution. I end by using these criteria to suggest a defence strategy for relativists. The moral of Part I is that while the empirical case for relativism about ‘know’ is less straightforward than relativists may have hoped, reflection on the force of error-theoretic objections and the adequacy conditions of error theoretic explanations paints a bright picture for relativism about knowledge attributions. There is thus good reason to think that the dependence of truth on the perspective of the evaluating subject is not restricted to intuitively subjective parts of our thought and talk.

Part II explores the connection between subjective attitudes and *de se* or first-personal attitudes – attitudes about oneself, which we typically express using the first-personal pronouns ‘I’, ‘me’, ‘my.’ On a widespread version of relativism about subjective attitudes, the content of these attitudes is understood in similar ways to a Lewisian account of the content of *de se* attitudes. On the Lewisian account, a *de se* thought such as *I am hungry* has *centered content*, which is modelled by a set of centered worlds. Centered worlds are possible worlds ‘centered’ on an individual at a time. Thus, the thought *I am hungry* has a content modelled by the set of centered worlds such that the center is hungry in the world at the time. A subjective thought such as *Beetroot is tasty* has a content modelled by the set of centered worlds such that beetroot is tasty to the center at the time. But while the centered contents of subjective and *de se* attitudes share many similarities, their roles in linguistic communication are crucially different. In Part II, I argue that this gives us reason to modify our account of centered content. In particular, I argue for what I call sequenced worlds relativism: a novel unified account of *de se* and subjective attitudes and their communication, on which contents are understood as sets of *sequenced worlds* – sets of possible worlds ‘centered’ on a sequence of individuals.

Chapter 4 motivates and develops the bare bones of the sequenced worlds view. I begin with a conflict between the centered content approach to subjective and *de se* attitudes and the received picture of linguistic communication, which understands communication as the transmission of information from speaker’s head to hearer’s head. I

then argue that sequenced worlds content provides a solution to the conflict that reconciles the insights of both views. In the rest of the chapter, I implement the sequenced worlds view in a broadly Stalnakerian model of communication and give a semantics for predicates of personal taste on which sentences like ‘Beetroot is tasty’ express sequenced worlds contents.

Chapter 5 and 6 develop the details of the sequenced worlds view. I show that the view itself forces no decision between contextualism and relativism upon its supporters because it is neutral between a nonindexical contextualist and a truth relativist view of predicates of personal taste. However, subsequent discussion of key pragmatic phenomena that the views are designed to explain shows that the relativist version makes better predictions. Chapters 5 and 6 thus present an argument in favour of sequenced worlds relativism.

In chapter 5, I start by presenting the nonindexical contextualist and truth relativist versions of the sequenced worlds view. I then extend the basic framework to explain the empirical data from eavesdropping cases and retraction cases. Crucial for the explanation of eavesdropping cases is an account of the relationship between a discourse and subjects who are not participants in the discourse but nonetheless receive information from it. In order to explain retraction cases on the broadly Stalnakerian model of discourse, I provide an account of how the information accumulated in the course of the conversation must be updated with the passage of time. I introduce a taxonomy for retraction cases and show how to account for the withdrawal of information in each of these cases.

Chapter 6 continues the development of the sequenced world view. I give an account of disagreement and explain how we can make sense of the notion of faultless disagreement. I show that the sequenced worlds model of common ground provides the resources for an account of the conversational dynamics of agreement and disagreement about matters of personal taste. Next, I discuss some implications of the view for our practice of ascribing monadic truth, and I address the complex interaction of predicates of personal taste with tense as well as with modality. This allows us to locate sequenced worlds relativism among the family of relativist views and to compare it with its main alternatives.

Part I

Knowledge

RELATIVISM AND SEMANTIC BLINDNESS

2.1 Introduction

The case for relativism about knowledge attributions, as made by John MacFarlane (2005a, 2011b), rests on the following line of thought. Indexical epistemic contextualism – the view that the content of sentences of the form ‘*S* knows that *p*’, and ‘*S* does not know that *p*’ may vary with the epistemic standards salient in the context of use – is supported by appeal to linguistic data from ordinary speakers’ use of knowledge sentences: the same speakers tend to accept ‘*S* knows that *p*’ as true when uttered in one context but not when uttered in another context. But contextualism makes incorrect predictions about, for instance, speakers’ inter-contextual truth ascriptions and their retraction of knowledge claims. The best explanation of this data available to contextualists involves the attribution of semantic blindness to speakers. Similarly, classical and subject-sensitive invariantists face troubling data, whose explanation commits them to some error theory or other. In contrast, MacFarlane claims, relativism can account for all of the data. Relativism is the only view that avoids the ‘double-edged sword’ of attributing systematic error to speakers (MacFarlane, 2005a, 215). That is its main virtue.

The relativist has a *prima facie* strong empirical case. *Pace* MacFarlane, however, I argue in this chapter that relativism cannot explain all of the data from speakers’ use of knowledge sentences without the attribution of semantic blindness to speakers. In section 2.2, I introduce the relativist’s case against contextualism and in favour of relativism.¹ In section 2.3, I present a recent error-theoretic objection against relativism that appeals to speakers’ inter-contextual judgments of relativised truth ascriptions to knowledge claims. In section 2.4, I show why the strategy behind this objection fails. However, there is a simpler semantic blindness objection to relativism that is immune to the given replies, as I argue in section 2.5. Speakers tend to be puzzled by sceptical paradoxes, and relativist explanations of this phenomenon are bound to appeal to speakers’ semantic blindness. I identify two different kinds of semantic blindness

¹In this chapter, I will not be concerned with MacFarlane’s arguments against traditional and subject-sensitive versions of invariantism that complete his case for relativism. The presentation of his argument against contextualism is intended to exemplify his argument-by-elimination strategy and to introduce the semantic framework.

involved in the debate: content-blindness and index-blindness. In section 2.6, I consider and reject a reply on behalf of the relativist. Finally, in section 2.7 I show that the objection generalises to ordinary cases of epistemic closure from a variety of data on which relativists have rested their case.

This chapter thus presents an argument against relativism based on the error-theoretic objection from sceptical paradox. But to relieve the anxious reader: this is not the final word. As I will argue in chapter 3, the prospects for relativism look bright once all the views' commitments to blindness theses are evaluated and compared.

2.2 Contextualism and the Case for Relativism

On contextualist semantics, the content expressed by an occurrence of '*S* knows that *p*', and '*S* does not know that *p*', depends in part on the epistemic standards salient in the conversational context, where this dependence can be traced to the occurrence of 'know' (Cohen (1987); DeRose (1995); Lewis (1996); cf. Schaffer (2004a)). When John uses the sentence 'Bill knows that he has hands' in an everyday context with low epistemic standards (LOW) he expresses, very roughly, the content *Bill knows relative to low standards that he has hands*. When Mary is in a context with high epistemic standards (HIGH), for instance in an epistemology class on scepticism or in a courtroom, she uses the sentence to express the content *Bill knows relative to high standards that he has hands*.² The sentence is true as used by John in LOW, but false as used by Mary in HIGH. The truth value of knowledge sentences can vary across contexts of use even when the facts about the knowing subject's situation do not change. Contextualists motivate the variability of content by appeal to ordinary speakers' use of 'know'. While speakers accept many knowledge attributions as true in mundane contexts of use, they tend to give in to, e.g., sceptical considerations that raise the epistemic standards and reject these attributions as false in such contexts.

Trouble for contextualism comes, among other things, from speakers' inter-contextual truth ascriptions to and retraction of knowledge claims, and from disagreement. When Mary is in HIGH, say in a conversation about brains-in-vats, she will judge knowledge attributions expressed by the sentence 'Bill knows that he has hands' false, even when the attribution is made by John in LOW. Likewise, speakers in LOW will judge knowledge denials, e.g. 'Bill does not know that he has hands', false even when they are made in HIGH. But contextualism predicts that in judging knowledge attributions

²This characterisation is simplified in two respects. First, epistemic standards – the strength of epistemic position one must be in to count as knowing – come in degrees. There are more different standards and corresponding contexts than HIGH and LOW. Second, the propositional form *S knows relative to high/low standards that p* is supposed to be neutral between different contextualist proposals for the structure of the propositions expressed.

and denials uttered in other contexts, speakers are sensitive to the epistemic standards at play in the context of use.

What is more, speakers will retract earlier knowledge claims when their context changes in relevant respects. Suppose John sincerely asserts ‘I know that my car is parked in the driveway.’ Mary points out that car thieves could be roaming John’s neighbourhood and that his car might have been stolen. She thereby raises the standards, and it is natural for John to retract his earlier assertion by saying ‘I guess I was wrong. I did not know that my car was parked in the driveway.’ But according to contextualism, John is mistaken in his retraction. After all, his earlier assertion is true in the context of its use. Contextualism predicts that it would be appropriate for John to reply: ‘I did not say that I *know* that my car is parked in the driveway. I only meant that I know by low epistemic standards that my car is parked in the driveway. And that is still true.’ This, however, is not a natural reply for John.³

Finally, speakers are reported to disagree when one of them, in Low, asserts ‘S knows that *p*’ and the other one, in HIGH, asserts ‘S doesn’t know that *p*.’ On the contextualist semantics, however, the first speaker expresses the content that S knows that *p* relative to low standards, and the second expresses the content that S does not know that *p* relative to high standards. Since these contents are perfectly compatible, speakers should not be inclined to judge that they disagree. That is, speakers are mistaken in their disagreement judgments.⁴

Contextualists’ best response to the recalcitrant data is to adopt an error theory: Speakers are systematically mistaken in their inter-contextual truth ascriptions and retraction of knowledge claims. These mistakes are explained by speakers’ *semantic blindness*: ‘users of the word ‘know’ are blind to the semantic workings of their language.’ (Hawthorne (2004, 107); cf. DeRose (2006, 321)) The kind of semantic blindness contextualists need to ascribe is what I propose to call content-blindness:

CONTENT-BLINDNESS

Speakers are blind to the fact that particular sentences (sentences of the form ‘S knows that *p*’/‘S does not know that *p*’) can express different contents in different contexts (bar indexical expressions in the substitution instances of ‘S’ or ‘p’).

Contextualists take the content of knowledge sentences to depend on the context of use, so the kind of semantic blindness they must attribute can be further specified as

³See MacFarlane (2005a) and Williamson (2005). Further objections to epistemic contextualist semantics concern, e.g., disagreement judgments (MacFarlane, 2007b), belief reports (Cappelen and Lepore, 2005; Hawthorne, 2004), the analogy between ‘know’ and gradable adjectives (Stanley, 2005), the analogy with indexicals and with quantificational determiners like ‘all’ and ‘every’ (Schaffer and Szabó, forthcoming).

⁴Regarding (dis)agreement reports as evidence for shared content, see Cappelen and Hawthorne (2009). Further objections to contextualism concern, e.g., belief reports (Cappelen and Lepore, 2005; Hawthorne, 2004) and the comparison of ‘know’ with gradable adjectives (Stanley, 2005).

use-content-blindness.

Let me settle some terminology. An expression is *indexical* iff its content at a context depends on features of the context. A sentence is (semantically) *context-sensitive* iff either its content or its truth value (or both) at a context depends on features of the context.⁵ Thus, contextualism about knowledge is the view that sentences of the form ‘*S* knows that *p*’ are indexical – specifically, that the content of ‘*S* knows that *p*’ depends on epistemologically significant features of the context. It is one among other views that hold that sentences of the form ‘*S* knows that *p*’ are context-sensitive.⁶ Content-blindness is ignorance of an expression’s indexicality.^{7 8}

According to MacFarlane, the attribution of semantic blindness undermines the contextualist’s cause. Contextualists appeal to the variability of speakers’ intra-contextual truth ascriptions in support of their view. But the more data they explain away by appeal to speakers’ ignorance of the semantic workings of ‘know’, the weaker is their appeal to speakers’ usage of ‘know’ as evidence in favour of their semantics.⁹ For MacFarlane, the point generalises:

[A] *general* problem with positing speaker error to explain away facts about use is that such explanations tend to undermine the evidential basis for the semantic theories they are intended to support. All of these semantic theories are justified indirectly on the basis of facts about speakers’ use of sentences, and the more error we attribute to speakers, the less we can conclude from these facts. (MacFarlane, 2005a, 215)

Relativists piggyback on the case against contextualism. On MacFarlane’s relativist semantics, the truth value of knowledge sentences can vary from one context of assessment to another: ‘*S* knows that *p*’ may be true at (context of use) C_U and (context of assessment) C_{A1} and false at C_U and C_{A2} .¹⁰ Underlying the relativist semantics is a roughly Kaplanian picture (Kaplan, 1989). Sentences in contexts of use express

⁵For the above notion of indexicality and a different notion of context-sensitivity, see MacFarlane (2009).

⁶These views include nonindexical contextualism (Brogaard, 2008; Kompa, 2002; MacFarlane, 2009) and relativism (MacFarlane, 2005a, 2011b).

⁷On the above notion of context-sensitivity, contingent sentences and tensed sentences count as context-sensitive. This broad notion of context-sensitivity is useful as an umbrella term for specific kinds of context-sensitivity, among which is the sensitivity to the epistemic standards salient at a context.

⁸Akerman and Greenough (2010) propose an alternative way of distinguishing between kinds of context-sensitivity and corresponding kinds of blindness. They provide a helpful discussion and comparison of these kinds of blindness for the case of vague expressions.

⁹For a defence of contextualism from semantic blindness objections see Blome-Tillmann (2008), Cohen (2004a), DeRose (2006), and Schaffer and Szabó (forthcoming).

¹⁰In this chapter, I will focus on MacFarlane’s version of relativism. Similar relativist semantics for knowledge attributions have been proposed by Richard (2004, 2008) and Kölbel (2009). All of the arguments in favour and against relativism discussed in this chapter apply, *mutatis mutandis*, to all of these versions.

contents. Contents are evaluated at an index (Lewis, 1980), what Kaplan called a circumstance of evaluation, to yield truth values.¹¹ Kaplan took the index (circumstances of evaluation) to include at least a world and a time coordinate, determined by the context of use. Relativism can be located in this picture by the following three theses:

1. INVARIANT CONTENT

‘S knows that *p*’ expresses a content that is invariant across contexts of use (bar indexical expressions in the substitution instances of ‘S’ and ‘p’).

2. INDEX SENSITIVITY

The index contains an epistemic standards coordinate, to which the truth value of ‘S knows that *p*’ is sensitive. More generally, call a sentence *index-sensitive* iff its truth value depends on (some coordinate in) the index.

3. ASSESSMENT SENSITIVITY

The truth value of ‘S knows that *p*’ depends on the epistemic standards salient at the context of assessment, which may be different from the context in which the sentence is uttered. Put together with 2, the epistemic standards-coordinate in the index is determined by the context in which a knowledge claim is assessed.

Relativism about knowledge attributions can thus be characterised as the view that while the semantic contribution of ‘knows’ to a sentence ‘S knows that *p*’ does not make the sentence indexical, the sentence is assessment-index-sensitive: Its truth value depends on the epistemic standards coordinate in the index, which is determined by the context of assessment.¹²

Relativism yields an elegant account of the data that troubles contextualism. An assessor in HIGH is correct in ascribing falsity to ‘Bill knows that he has hands’ as uttered in LOW, since it is the standards at play in the assessor’s context that matter to her truth value ascription. Moreover, relativism predicts that John will retract his earlier assertion ‘I know that my car is parked in the driveway’ when the context is shifted to HIGH by Mary in conversation. When John in HIGH assesses his past assertion, it is appropriate for him to use the standards salient in his present context to judge his assertion false and to correct himself. Relativism also explains intra-contextual truth ascriptions. When I judge knowledge sentences uttered in my own context as true (false), context of assessment and context of use are identical. Relativism’s predictions in these cases coincide with those of contextualism. Finally, relativists can explain the

¹¹I will here ignore the differences between Lewisian index and Kaplanian circumstance of evaluation. They have no bearing on any of arguments discussed in this chapter. I will henceforth use the shorter ‘index’.

¹²Contextualism can also be situated in this picture of semantics: It locates the relevant epistemic standards in the context of use that combines with a sentence to deliver a content; contents then are evaluated at an index, which does not include a standards coordinate.

felicity of disagreement reports of two speakers in relevantly different contexts of use, one of whom grants and the other denies that a subject knows that p , by the fact that, on account of INVARIANT CONTENT, one speaker expresses the negation of the other speaker's knowledge claim.

MacFarlane completes the case for relativism by making similar points against traditional invariantism and interest-relative invariantism. These views run into trouble with data, e.g., from intra-contextual truth ascriptions and temporal and modal embeddings of 'know', respectively. And again, relativism makes correct empirical predictions where invariantist views must wield error theories to their defence. MacFarlane concludes that since relativism is the only view that respects all of the data without requiring the attribution of systematic speaker error, it is superior to all other views.

2.3 An Error-theoretic Objection against Relativism

Contextualism and all forms of invariantism each face trouble from cases in which their view does not predict the assessor's egocentric sensitivity to epistemic standards (or stakes) in their inter-contextual judgments of knowledge attributions. One error-theoretic strategy against relativism, exemplified by Montminy (2009), presents cases for which relativism supposedly cannot predict this egocentric focus either. As we will see in section 2.4, this particular strategy is not promising.

According to the first variant of the strategy that Montminy pursues, relativism is committed to the attribution of systematic speaker error when speakers make judgments about explicitly relativised truth claims, made in other contexts, concerning knowledge attributions. Montminy has us consider, first, the following dialogue in Low:

- (2.1) John: We both know that Neil Armstrong was the first man to set foot on the moon.
 Bob: That's true.

Suppose Mary is in HIGH and is presented with (2.1). It is natural for her to judge false not only John's knowledge attribution but also Bob's assessment of John's claim. This is what relativism predicts. Mary is right in taking the standards salient in her context as relevant for the assessment of John's knowledge attribution as well as Bob's truth ascription to John's claim. Mary thinks that John and Bob do not know that Neil Armstrong was the first man to set foot on the moon, and that it is false that they know this, so Bob's claim that it is true that they know is false in Mary's context.

Trouble for relativism is supposed to arise when speakers explicitly relativise their truth ascriptions. Consider a similar dialogue in Low:

- (2.2) John: We both know that Neil Armstrong was the first man to set foot on the moon.
 Bob: That's true relative to this context.

According to Montminy, Mary in HIGH will judge John's knowledge attribution false and she will also judge Bob's truth ascription false, *regardless of the relativisation*. This runs against relativism's prediction. On relativist semantics, the explicitly relativised truth ascription made by Bob in Low is true at every context of assessment, hence also at Mary's.

Relativists hold that there are (at least) two different truth predicates. The first is the ordinary English monadic predicate 'true'. Its application to propositions yields sentences that are themselves assessment-sensitive.¹³ That is, sentences of the form 'It is true that p ' or 'The proposition that p is true' have truth values that vary with the context of use and context of assessment. Importantly, monadic 'true' is disquotational in the following sense: Whenever, in a context of assessment, we correctly judge that S knows that p , we can also judge correctly that it is true that S knows that p , and vice versa. The equivalence schema *It is true that p iff p* is true at any C_U and C_A . Since 'true' is assessment-sensitive, the left-hand side will be true (false) at a C_U and C_A just in case the right-hand side is true (false) at that C_U and C_A .

In dialogue (2.1), Bob is ascribing ordinary monadic 'true' to the proposition expressed by John's knowledge attribution. In the context of assessment of the dialogue, Low, this truth ascription is correct; 'That's true' is true at $C_{U(J\&B)}$ and $C_{A(Low)}$. However, 'That's true' is false at $C_{U(J\&B)}$ and $C_{A(High)}$. Relativists predict correctly that it is appropriate for Mary in her context of assessment to reject Bob's assertion as false.

The second truth predicate is relational. 'True relative to context C' relates truth bearers to contexts in which they are evaluated.¹⁴ It is assessment-insensitive. That is, the truth value of ascriptions of this truth predicate to sentences or propositions, as in 'It is true relative to this context that p ', does not vary with contexts of assessment. Thus, the proposition expressed by the following assertion is true relative to any context of assessment: 'It is true relative to Low that John and Bob know that Neil Armstrong was the first man to set foot on the moon.' Since this is plausibly what Bob is saying in (2.2), relativism predicts that Mary in HIGH will judge the proposition expressed by Bob's utterance as true.

¹³Egan et al. (2005) take the bearers of monadic truth to be utterances. I shall here stick with MacFarlane (2007b, 2011a) and take ordinary 'true' as applying to propositions.

¹⁴Egan et al. (2005) treat the binary predicate 'true relative to context C' as applying to utterances. MacFarlane prefers the metalinguistic predicate 'true at C_U and C_A ' used in the formulation of the relativist semantics, which relates either sentences or propositions with contexts of use and contexts of assessment. Whatever the exact details, the common denominator here is the explicit relativisation of truth to a context in which the truth bearer is evaluated. I will use 'true relative to context C' as a predicate that applies to truth bearers, be they utterances, sentences(-in-context), or propositions.

However, speakers in HIGH such as Mary seem to judge explicitly relativised truth ascriptions to knowledge attributions made in Low as false. So on the relativist picture, they are systematically mistaken in these judgments. The conclusion of this objection is that relativism, too, is committed to an error theory; the best argument for relativism fails.¹⁵

2.4 Relativist Replies

There are at least two replies available to the relativist in response to this error-theoretic objection.

(1) *Discarding the data.*¹⁶ Relativists may remind us that while ‘true’ is a natural language English truth predicate that speakers frequently and competently use, ‘true relative to context *C*’, or ‘true at *C_U* and *C_A*’, is a technical, metalinguistic truth predicate used to formulate relativist semantics *for* English sentences. To ordinary speakers, Bob’s assertion of ‘That’s true relative to this context’ will sound stilted and odd, and contrary to Montminy’s Mary, they are most likely to have no clear truth value judgments.¹⁷ But even if speakers had clear judgments, we could not just assume that ‘true relative to context *C*’ in their mouths expressed the relativist’s technical concept of truth. In fact, trying to accommodate speakers’ intuitions about this expression would be bad methodology. By analogy, we would be ill-advised to take into account speakers’ reactions to ‘That’s a context-sensitive sentence’ in constructing a semantics for English. As a result, speakers’ reactions to expressions involving ‘true relative to this context’ do not constitute data against relativism.

¹⁵Montminy (2009) does not claim that this first objection goes through against the relativist, for roughly the reasons I give at the beginning of the next section. He does hold, however, that another variant of the strategy is successful, one on which the claims involving the explicitly relativised truth predicate are replaced by supposedly extensionally equivalent counterfactual conditionals. I argue in the next section that that objection does not succeed against the relativist either.

¹⁶In Kindermann (forthcoming), I mention, in footnote 17, MacFarlane’s views on norms of assertion from chapter 5 of his unpublished book manuscript. At the time the article was published *Online First*, the book manuscript available on MacFarlane’s website did not include a chapter on knowledge attributions. In a more recent draft (February 8, 2012), MacFarlane discusses Montminy’s objection in chapter 9 on ‘knows.’ Much of his response is broadly consonant with the first reply I am giving here. We differ, however, in the response to the part of Montminy’s objection that is concerned with retraction (see footnote 17 in Kindermann (forthcoming)/footnote 19 below for my response). More broadly, my goal in the article and this chapter is to show that the successful error-theoretic objection that relativists need to worry about comes from sceptical paradox and epistemic closure puzzles.

¹⁷Results from Google searches should not serve as conclusive evidence, but the fact that a Google search for ‘true relative to this context’ (in July 2011) resulted in only four hits – three papers in philosophy (including Montminy (2009)) and one in linguistics – strongly suggests that English speakers are not well-acquainted with the expression. ‘True relative to context’ offered 29 hits, all of which were philosophy or linguistics papers. A feeling of oddity was also the reaction of most philosophers and non-philosophers when I presented them with dialogue (2.2).

Relativists could further argue that there is no ordinary, non-technical English expression that correctly captures the theoretical concept of relative truth. This would forestall all attempts to run the error-theoretic objection with everyday expressions that purport to stand a better chance of expressing relative truth in the mouths of English speakers. Relativists could give this reply without compromising their claim that ordinary speakers can make sense of relative truth as a concept of *truth*. According to MacFarlane, we understand the meaning of ‘true at C_U and C_A ’, or ‘true relative to context C ’, if we grasp ‘the role [this predicate] plays in a broader theory of language use: specifically, an account of the speech act of assertion’ (MacFarlane, 2005b, 329). In brief, an assertion is a commitment to the truth of what is asserted. This commitment is honoured by providing adequate grounds for the truth of what is asserted relative to the context of assessment in which the assertion is challenged. The commitment also requires that one withdraws the assertion in any future C_A in which what is asserted is shown to be untrue relative to C_A .¹⁸ Retraction data, the relativist can conclude, gives evidence that speakers take themselves to be bound by such a relativised assertoric commitment. Thus, even if speakers lack ordinary expressions to express the technical concept of relative truth, they have an implicit grasp of it.¹⁹

¹⁸ See for instance MacFarlane (2005a, 329) and (2005b, 336-7) for further details of the commitment undertaken by assertions.

¹⁹ Montminy argues that even this connection between the notion of truth at C_U and C_A and language use requires the attribution of systematic error to speakers. As he points out, the commitment to withdrawing an assertion in any future context of assessment in which what is asserted is shown to be untrue relative to that context of assessment entails (*):

(*) In asserting ‘ S knows that p ’ (‘ S does not know that Q ’), one commits oneself to withdrawing the assertion in any future context $C_{A(High)}$ ($C_{A(Low)}$) in which what is asserted by ‘ S knows that p ’ (‘ S does not know that Q ’) is shown to be untrue relative to $C_{A(High)}$ ($C_{A(Low)}$).

Montminy grants for the sake of argument that speakers do in fact withdraw their assertions of ‘ S knows that p ’ (‘ S does not know that Q ’) made in Low (HIGH) when challenged in HIGH (Low). But he denies that speakers ‘take themselves to be bound by [(*)]’: ‘the fact that a speaker in Low would withdraw her previous knowledge denial made in HIGH does not entail that *when she is in HIGH*, the speaker takes herself to be committed to withdrawing her current knowledge denial, if this denial is challenged in some future Low. As a matter of fact, a speaker in HIGH would reject this commitment, that is, such a speaker would hold that it would be *incorrect* to withdraw her current knowledge denial in some future low-standards context’ (Montminy, 2009, 354). Montminy concludes that relativism implies that speakers are systematically mistaken about their commitments to withdraw knowledge claims.

Relativists can resist this argument in several ways. First, it is not clear that speakers would in fact reject the commitment to withdraw knowledge attributions (denials) once the standards have risen (fallen). What could the evidence be for speakers’ rejection of the commitment if it is not their linguistic behaviour (which, Montminy grants, honours this commitment)? Presumably Montminy has in mind speakers’ explicit judgments about what they take their commitments to be (what they ‘hold’). Pending empirical evidence, relativists may simply doubt that speakers would make judgments that stand in contrast to their actual behaviour – after all, they do seem to withdraw. (It has been widely noted that raising standards is easier than lowering them, which might in part explain why speakers are more reluctant to explicitly withdraw knowledge denials (because they do not accept switching to Low). But this does not threaten MacFarlane’s claim about the commitment to withdraw, which only says that when speakers *are* in Low, they will withdraw a knowledge denial that is untrue relative to Low.)

But perhaps the relativist's blank dismissal of any sort of data involving the use of vocabulary expressing relative truth is too quick. Montminy agrees that the relativist's reply successfully rebuts the objection based on an explicitly relativised truth predicate in cases like (2.2). But he maintains that there are perfectly ordinary English locutions that do express relativised truth claims. The following counterfactual conditionals are supposed to express in ordinary English the relativist's thesis that the proposition that John and Bob know that Neil Armstrong was the first man to set foot on the moon is true relative to John and Bob's context LOW:

- (2.3) If it were to be assessed in John and Bob's context, the proposition that they know that Neil Armstrong was the first man to set foot on the moon would be true.
- (2.4) If I were in John and Bob's context, then they would know that Neil Armstrong was the first man to set foot on the moon.
- (2.5) If the stakes were low and no error possibilities had been mentioned, then it would be true that John and Bob know that Neil Armstrong was the first man to set foot on the moon.

Montminy claims that ordinary speakers like Mary in HIGH would understand (2.3) – (2.5) and reject them as false. But the relativist must hold that (2.3) – (2.5) are true, since they express the relativist's thesis that the attribution of knowledge to John and Bob is true relative to John and Bob's context. The relativist, Montminy concludes, is after all committed to the claim that Mary's judgments are systematically mistaken (Montminy, 2009, 350-2).

I think relativists have good reasons to reject both that (2.3) – (2.5) express relativised truth claims and that Mary's falsity judgments concerning these conditionals (if she has any clear judgments) run against relativism's predictions. On a mainstream semantics of counterfactuals, the antecedent has an effect on the *worlds* at which the consequent is evaluated for truth and falsity. For instance, on the Lewis-Stalnaker account, a conditional $A > B$ is true iff B is true at the closest world(s) at which A is true. But

Second, relativists may even grant that speakers do make judgments to the effect it is incorrect to withdraw a knowledge attribution in HIGH, but deny that speakers' explicit judgments are relevant for an account that links the concept of relative truth to speakers' language *use*. By analogy, syntacticians would have to predict widespread error if part of their evidence was speakers' acceptance or rejection of explicitly stated grammatical rules which speakers, as a matter of fact, employ in their use of language.

Finally, it is worth pointing out that MacFarlane's (2005b) commitment account of assertion is by far not the only way of linking the concept of relative truth to language use. MacFarlane (2012) argues that many of the common accounts of assertion corroborate the concept of relative truth as the link between the compositional semantics and language use as long as a corresponding norm of retraction is added. And Egan et al. (2005), Egan (2007), and Stephenson (2007a) offer alternative accounts of assertion for relativists.

it is an open question whether, when combined with relativism, the antecedent also affects the epistemic standards coordinate in the index. Relativism has Montminy's problem only if the antecedent does affect the epistemic standards coordinate. However, there seem to be no independent theoretical reasons why this must be the case. And the alleged data from Mary's falsity judgments give us pre-theoretical reasons to think that the antecedent has no effect on the epistemic standards coordinate with respect to which the consequent is evaluated. So relativists can make the right predictions by simply adopting the unamended Lewis-Stalnaker analysis of conditionals: When Mary assesses the knowledge claim in the consequent at the closest world(s) in which she is in John's context, or in which the stakes are low and no error possibilities have been mentioned, she will still judge this knowledge claim as false relative to her actual context of assessment HIGH. As a result, (2.3) – (2.5) are not extensionally equivalent to, and do not express, relativised truth ascriptions, and Lewis-Stalnaker-cum-relativist semantics of (2.3) – (2.5) make the intuitively right predictions concerning Mary's (putative) rejection.

(2) *Override*. It is worth noting that there is a reply available to relativists even if they are inclined to take the data from (2.2) at face value and affirm that 'true relative to context C ', as used by ordinary speakers, is the relativist's metalinguistic truth predicate. On this assumption, what the data shows is that speakers take their own current contexts of assessment to override all other contexts of assessment. Adapting Moruzzi and Wright (2009), the relativist can capture the phenomenon of overriding by introducing two principles governing the metalinguistic truth predicate.

Let the relativist's metalinguistic truth predicate be represented by the dyadic predicate $T(S, C_A)$ relating sentences-at-contexts²⁰ and contexts of assessment.²¹ First-order truth ascriptions such as Bob's in (2.2) are represented by sentences of the form $T(S, C_A)$; second-order truth ascriptions such as Mary's by $T(T(S, C_{A1}), C_{A2})$. (Of course, Mary makes a second-order ascription of falsity; see below.) A context C_{A1} *overrides* a context C_{A2} when the assessment in C_{A1} of the facts dictates their proper assessment in C_{A2} . The fact that Mary judges Bob's use of 'That's true relative to this context' as false in accordance with the epistemic standards in her own context of assessment, and in spite of the standards at play at Bob's context of assessment, is captured by the ideas of (Upward Override) and (Downward Override).²² Upward Override with

²⁰Sentences-at-contexts can be understood as pairs of sentences and contexts of use. We could make contexts of use explicit and conceive of the truth predicate as a tryadic predicate relating sentences, contexts of use, and contexts of assessment. To keep the discussion as simple as possible, we will let ' S ' stand for sentences-at-contexts.

²¹In the following, I will accept harmless use/mention confusion in favour of better readability.

²²(Upward Override) and (Downward Override) are derived from Moruzzi and Wright's idea of trumping (2009, §3), which they promote for relativism about future contingents. Override contrasts with trumping in one crucial respect: 'Being trumped [...] involves that another perspective gets, not to *override* the

respect to a sentence S takes place if whenever S is true at a context of assessment C_{A1} it is thereby true at C_{A1} that S is true at *any* context of assessment. In short:

$$\text{(Upward Override)} \quad (\forall C_{A1}, C_{A2})(T(S, C_{A1}) \supset T(T(S, C_{A2}), C_{A1}))$$

This principle is called ‘Upward Override’ because the direction of dictation is from the first-order truth ascription in C_{A1} to the second-order ascription in C_{A1} of truth to the first-order truth ascription in some C_{A2} . Provided that S ranges over knowledge sentences, (Upward Override) entails that speakers will judge as true that a knowledge sentence is true relative to any context of assessment C_{A2} if they judge it true relative to their own (C_{A1}).

The converse of this principle is (Downward Override). Downward Override with respect to S takes place if whenever it is true at a context of assessment C_{A1} that S is true at a context of assessment C_{A2} , S is true at C_{A1} . In short:

$$\text{(Downward Override)} \quad (\forall C_{A1}, C_{A2})(T(T(S, C_{A2}), C_{A1}) \supset T(S, C_{A1}))$$

This principle is called ‘Downward Override’ because its direction of dictation is from the second-order ascription in C_{A1} of truth to the first-order truth ascription in any C_{A2} to the first-order truth ascription in C_{A1} . Provided that S ranges over knowledge sentences, (Downward Override) entails that speakers will judge a knowledge sentence true relative to their own context of assessment C_{A1} if they judge as true that the knowledge sentence is true relative to a context of assessment C_{A2} . (Upward Override) and (Downward Override) together entail that it is inconsistent for a speaker to accept a knowledge sentence as true relative to their own context of assessment while rejecting as false truth ascriptions to this knowledge sentence at other contexts of assessment, and vice versa.

Mary’s second-order falsity ascription to Bob’s first-order truth ascription is a contraposition instance of (Downward Override). Let S_K stand for John’s knowledge attribution, C_M for Mary’s context of assessment, and C_B for Bob’s context of assessment. Interpreting falsity as the negation of truth, we get:

$$(2.6) \quad \sim T(S_K, C_M) \supset \sim T(T(S_K, C_B), C_M)$$

Loosely speaking, (2.6) says that since Mary judges John’s knowledge attribution false from her context of assessment, she also judges that Bob’s truth ascription (in his context of assessment) to John’s knowledge attribution is false. The epistemic standards in her context of assessment override those in Bob’s context of assessment.

(Upward Override) and (Downward Override) thus allow the relativist to make the right predictions for (2.2). Acceptance of these principles is also compatible with the mandates of one’s own perspective, but to determine what they are. (Moruzzi and Wright, 2009, 314 n.10)

other data relativists have been keen on. Firstly, it does not affect the predictions about intra-contextual truth ascriptions. Secondly, it supports retraction data and makes even slightly stronger predictions. We retract knowledge attributions or denials if they turn out to be false from our current context of assessment. Given Override, we *also* judge that the knowledge claim was false in our earlier context of assessment (in fact, in any context of assessment), contrary to our previous judgment. We concede that we were wrong back then and do not tend to justify our knowledge claim by insisting that it was true relative to our previous context of assessment. Finally, (Upward Override) and (Downward Override) leave the data from ordinary disquotational ‘true’ untouched, since they exclusively target the explicitly relativised, technical truth predicate.

Override principles also lend support to an intuitive difference between, on the one hand, discourse featuring knowledge claims and discourse about the contingent future, and, on the other hand, discourse about matters of personal taste and perhaps moral and aesthetic discourse. Relativism about knowledge sentences and future contingents (see Moruzzi and Wright) promote override principles. On the contrary, the idea of override is questionable in moral and aesthetic discourse and entirely out of place in matters of personal taste. Consider (2.7).

- (2.7) John: Watching *2001: A Space Odyssey* is fun.
 Bob: No, watching *2001: A Space Odyssey* is not fun. Still, it may be true for you that it’s fun.

Accepting on the current strategy that ‘true for you’ in (2.7) expresses a relativised truth predicate, relativists can say that it is true relative to John, i.e. the standards of taste at his context of assessment C_J , and false relative to the standards of taste at Bob’s context C_B , that watching *2001: A Space Odyssey* is fun: $T(S_{2001}, C_J)$ and $\sim T(S_{2001}, C_B)$, where S_{2001} is the sentence ‘Watching *2001: A Space Odyssey* is fun’ in (2.7). Would John (or Bob) be tempted to think that because from his context of assessment, it is true (false) that watching *2001: A Space Odyssey* is fun, this must also be true (false) from any other context of assessment? If override principles held in the case of predicates of personal taste, then Bob’s assessment of ‘Watching *2001: A Space Odyssey* is fun’ as false would mandate the judgment that from John’s context of assessment, that sentence is equally false. This prediction, however, seems empirically inadequate. Typical speakers may concede that while something is not fun for them, it may well be for others. Their own standards of taste do not mandate what taste judgments others must make. This is different for knowledge attributions. Notice the bizarreness of Bob’s concession to John in the following dialogue:

- (2.8) John: I know that Neil Armstrong was the first man to set foot on the moon.
 Bob: # No, you don't know that. Still, it may be true for you that you know it.

Override principles plausibly hold for relativism about knowledge sentences and future contingents but are misplaced in the case of discourse involving predicates of personal taste. Formally, this is easily captured by restricting the range of S in (Upward Override) and (Downward Override) to sentences with which we make contingent claims about the future as well as knowledge claims.

Override principles mark the difference between discourses that typically exhibit *Tolerance* and those which do not. *Tolerance* here captures the idea that speakers can acknowledge that there might be other perspectives than their own which are (equally) appropriate; that is, speakers can acknowledge that given different relevant features in other contexts of assessment, what they judge to be true (false) may not be true (false) from other contexts of assessment.²³ Speakers are likely to be tolerant towards other perspectives on taste, but not on whether or not someone has knowledge.²⁴

2.5 Sceptical Paradox and Index-Blindness

The error-theoretic strategy from speakers' inter-contextual judgments in sections 2.3 and 2.4 appealed to data involving the use of 'true relative to context C ' and other locutions allegedly expressing relativised truth claims. The replies available to the relativist correspondingly targeted the use of these expressions to avoid the charge of systematic error attribution. However, there is a simpler error-theoretic objection that makes

²³Cf. Moruzzi and Wright's (2009) idea of 'even-handedness'

²⁴Override principles commit the relativist to second-order relativism. First-order relativism about knowledge sentences is the claim that there is at least one sentence containing 'know' which, as used in a context of use C_U , is true relative to one context of assessment and false relative to another:

$$(\text{REL}) \quad (\exists S, C_{A1}, C_{A2})(T(S, C_{A1}) \wedge \sim T(S, C_{A2}))$$

Second-order relativism about knowledge sentences is the claim that there is at least one sentence of the form $T(S, C_A)$, where S ranges over, say, knowledge sentences, and which is true relative to one context of assessment and false relative to another:

$$(\text{REL}^2) \quad (\exists S, C_{A1}, C_{A2}, C_{A3})(T(T(S, C_{A1}), C_{A2}) \wedge \sim T(T(S, C_{A1}), C_{A3}))$$

Given first-order relativism, an instance of (REL^2) can be derived from override principles. Suppose for a knowledge sentence S and two contexts of assessment c_{a1} and c_{a2} that $T(S, c_{a1})$ and $\sim T(S, c_{a2})$. Then from $T(S, c_{a1})$ it follows by (Upward Override) that $T(T(S, c_{a2}), c_{a1})$. And from $\sim T(S, c_{a2})$, taking c_{a2} for both C_{A1} and C_{A2} , it follows by contraposition on (Downward Override) that $\sim T(T(S, c_{a2}), c_{a2})$.

A further question is whether first- and second-order relativism can coherently be combined with third-order absolutism, or whether higher-order relativism is required all the way up. It is beyond the scope of the present work to determine whether higher-order relativism of any sort is a price worth paying for accommodating data from inter-contextual truth ascriptions.

no appeal to data involving relativised truth claims and is thus immune to the above replies. As I argue in this section, the solution to sceptical paradoxes and epistemic closure puzzles that naturally falls out of the relativist semantics requires the implausible attribution of a kind of semantic blindness I call index-blindness. Although relativists have not explicitly addressed sceptical paradoxes, they cannot reject this solution without significant costs, as I show in section 2.7.

Consider the following sceptical argument:

- (SA) I don't know that I'm not a BIV (i.e., a bodiless brain in a vat who has been caused to have just those sensory experiences I've had).
 If I don't know that I'm not a BIV, then I don't know that I have hands.

I don't know that I have hands.

As Schiffer (1996, 317) remarks, 'this argument presents a paradox because it tempts us to say three things that are mutually inconsistent: its first premise is true; its second premise is true; and its conclusion is false.' That is, the following three sentences are mutually inconsistent: 'I don't know that I'm not a BIV'; 'If I don't know that I'm not a BIV, then I don't know that I have hands'; and 'I know that I have hands.' Yet each of these sentences strikes us as intuitively true.

A 'fully satisfactory' solution, Schiffer continues, must accomplish two things: First, it must explain why (SA) in fact does not present a paradox. That is, it must show which one of the three sentences is false, and explain why the argument is in fact valid. And second, it must explain why (SA) *seemed* to present a paradox. That is, it must explain why the false sentence seemed true, and why we were tempted to believe that the premises are true but the conclusion false.

The relativist semantics has the resources to provide a fully satisfactory solution to sceptical paradox. Going for the first task, relativists may remind us that we assess knowledge attributions and denials from a context of assessment. Sceptical hypotheses, as introduced by the first premise of (SA), raise the epistemic standards in the context of assessment to a high level (HIGH), since one must be in quite a strong epistemic position to rule out that a sceptical hypothesis like the BIV scenario obtains. Thus, we assess 'I don't know that I'm not a BIV' as true at C_U and $C_{A(High)}$. Importantly, the epistemic standards are now raised for a context of assessment that encompasses the entire argument (SA). Hence, we must also assess the second premise and the conclusion of (SA) at C_U and $C_{A(High)}$. This allows the relativist to explain why (SA) in fact is valid, thus capturing the intuitive force of the sceptic's argument. To understand how, we need to have a look at the relativist's notion of validity.

The validity of an argument is commonly understood as necessary preservation of truth from premises to conclusion. Since truth, on a relativist semantics, is truth at a

context of use C_U and context of assessment C_A , validity is necessary preservation of truth at C_U and C_A from premises to conclusion.²⁵ Thus, (SA) is valid if, necessarily, whenever the premises are true at C_U and $C_{A(High)}$, its conclusion is also true at C_U and $C_{A(High)}$.²⁶ Since both premises and its conclusion are true at C_U and $C_{A(High)}$, (SA) is valid and sound.

This gives relativists a neat explanation of the validity and soundness of (SA). It also allows them to show that (SA) does not present a paradox. A paradox is ‘a set of mutually inconsistent propositions each of which enjoys some plausibility when considered on its own’ (Schiffer, 1996, 324). Propositions, and sentences, are mutually inconsistent just in case they cannot be true together. Obviously, relativists will relativise truth here as well. Thus, a paradox is a set of sentences, or propositions, each of which enjoys some plausibility when considered on its own but which cannot be true together at any C_U and C_A . But there are no contexts C_U and C_A at which each of the following three sentences enjoys some plausibility on its own yet they cannot be true together: ‘I don’t know that I’m not a BIV’; ‘If I don’t know that I’m not a BIV, then I don’t know that I have hands’; ‘I know that I have hands.’ From any $C_{A(High)}$, the first two are true and plausible, but the third lacks plausibility – it is false at $C_{A(High)}$. From any $C_{A(Low)}$, the second and third are true and plausible, but the first is not.²⁷ Provided that the first premise of (SA) induces $C_{A(High)}$ for the entire argument, and so for the three sentences of the apparent paradox, ‘I know that I have hands’ turns out to be the apparently true, but in fact false, sentence – it is false at $C_{A(High)}$. This completes the first part of the relativist solution.

To deliver on the second part, relativists must explain why ‘I know that I have hands’ seemed true and why it appeared that while the premises of (SA) are true, the conclusion is false. Why our resistance to accept the sceptical conclusion ‘I don’t know that I have hands’? The only answer for relativists is this: We intuitively assess the conclusion from our everyday context in which low epistemic standards prevail, while we assess the sceptical hypothesis introduced by the premises from a context with extraordinarily high epistemic standards. In (SA), we reason from the premises’ truth at C_U and $C_{A(High)}$ to the conclusion’s falsity at C_U and $C_{A(Low)}$. Yet we are ignorant of our switching contexts of assessment in moving from premises to conclusion, and its effect on our truth and falsity judgments. In the context of assessment induced by the premises of the argument (HIGH), the conclusion is in fact true. But we believe it to be false because we unknowingly fall back to assessing it from LOW. Likewise, the

²⁵Cf. MacFarlane’s definition of logical consequence (MacFarlane, 2011a, 167)

²⁶Note that to get necessary and sufficient conditions for the validity of (SA), it must be the case that, necessarily, for every C_U and C_A , if the premises are true at C_U and C_A , then the conclusion is true at C_U and C_A . It is easy to see that for any C_A with sufficiently high standards to make the premises of (SA) true, the conclusion will also be true.

²⁷For simplicity, I assume that all contexts $C_{A(High)}$ and all contexts $C_{A(Low)}$ together exhaust all possible contexts of assessment.

premises together with the unnegated conclusion strike us as paradoxical because we are unaware of the fact that we assess them from different contexts of assessment, and that this influences our judgments.²⁸

The second part of the relativist solution²⁹ involves the attribution of error to ordinary speakers – indeed to anyone who is tempted to think (SA) presents a paradox. We mistakenly think (SA) gives rise to paradox because we mistakenly switch contexts of assessment midway. Relativists must explain this error by appeal to a specific kind of semantic blindness. For relativists, knowledge attributions are assessment-index-sensitive. Ignorance of this sensitivity is what leads speaker to their mistaken judgments. They betray index-blindness:

INDEX-BLINDNESS

Speakers are blind to the fact that the truth value of contents expressed by sentences involving a particular expression ('know') can vary with a particular coordinate in the index (epistemic standards).

More precisely, we can say that relativists are committed to the thesis that speakers are epistemically assessment-index-blind. They are ignorant of the fact that the truth value of knowledge sentences can vary with the epistemic standards coordinate in the index, which is determined by the context of assessment. In assessing (SA), they are blind to the fact that the truth value of knowledge attributions is sensitive to the epistemic standards coordinate in the index, which is determined by their context of assessment. Speakers do not realise that when they assess the premises as true and the conclusion as false, they do so from different contexts, and that this affects their assessment.³⁰

²⁸I have chosen a phenomenology of sceptical arguments according to which, as we reason from premises to conclusion, we go from corresponding truth judgments to a falsity judgment. Accordingly, the relativist solution diagnoses a switch of contexts of assessment in the move from premises to conclusion. Nothing hangs on this choice. One might prefer to say that in reasoning through (SA), we come to accept the sceptical conclusion and only then remember our usual acceptance of the unnegated conclusion in everyday contexts. These two judgments then strike us as inconsistent, creating the air of paradox. If we prefer this diagnosis, what needs explanation in terms of index-blindness is the mistaken feeling of inconsistency between one's truth judgment of the conclusion in the high-standards context of assessment of (SA) and one's falsity judgments of 'I don't know that I have hands' in everyday low-standards context of assessment.

²⁹Note that the relativist 'solution' to sceptical paradox, like the contextualist's, does not amount to a refutation of scepticism. On the contrary, the relativist account gives sceptical intuitions their due. This does not imply that we can never (in no context of assessment) truly attribute knowledge. Whether or not sceptical standards are reasonable ones to adopt is a question on which the relativist semantics is neutral.

³⁰Relativists need a *local* index-blindness thesis. That is, speakers are imputed with index-blindness only locally, namely regarding sceptical arguments, situations in which they are confronted with the three knowledge sentences from (SA), and epistemic closure puzzles, as I argue below. Local index-blindness must be distinguished from *global* index-blindness, according to which speakers are blind to the index-sensitivity of an expression (knowledge sentences) in any situation in which the expression is used or evaluated. The latter is not needed to explain sceptical paradox, and its attribution would undermine

The attribution of index-blindness to speakers spoils the relativist's self-proclaimed status as a no-blindness-theory and undermines the main advantage of relativism. But one need not share MacFarlane's general worry about blindness attributions to find index-blindness an unpalatable consequence of relativism. A more serious problem with index-blindness is that it is particularly implausible in light of speakers' reliable competence in handling other kinds of index-sensitivity. Take as an example the implementation of temporalism according to which the contents expressed by sentences in contexts of use vary their truth value with a time coordinate in the index.³¹ Consider the following little argument:

- (T) John is asleep at home.
 If John is asleep at home, then he is not in his office.

 John is not in his office.

Clearly, the argument strikes us as valid. And we would easily detect a change in the value of the time coordinate in the index. Suppose Bob wants to convince Mary at midnight that John never makes it to his office before 11am. Bob gets Mary to agree to premise one and two. If he were then to say, 'You see, John is not in his office!', Mary would respond that John is not in his office at the time of their conversation (midnight), but that this did not show that he is not in his office mornings before 11am. Mary is acutely aware of the fact that she agreed to the premises' truth at midnight, and she is happy to concede the conclusion's truth at midnight. But she will immediately spot the oddness of Bob's switching the time coordinate in reasoning from the premises' being true at midnight to the conclusion's being true in the morning. Temporal (use-)index-sensitivity does not pose a challenge to ordinary speakers. But if more uncontroversial phenomena of index-sensitivity do not cause speaker error, why should we assume that sentences involving 'know' are index-sensitive despite speakers' blindness to this index-sensitivity?³²

other data in favour of a relativist semantics. Note that the distinction between local and global blindness is orthogonal to the distinction between content-blindness and index-blindness. The extent to which a semantic theory needs a local content-blindness thesis is determined by the range of data that this thesis needs to explain.

³¹A similar case could be constructed for 'orthodox' semantics, on which sentences vary in truth value only with a worlds coordinate in the index.

³²The objection in this section applies to nonindexical contextualist solutions of the paradox as well. Nonindexical contextualism (Brogaard, 2008; Kompa, 2002; MacFarlane, 2009) shares with relativism the theses of invariant content and index-sensitivity (cf. theses 1 and 2 in section 2.2) but assumes that the epistemic standards coordinate in the index is determined by the context of *use*. But the kind of use-index-blindness required for nonindexical contextualist solutions is by no means more plausible than the one relativists are committed to.

2.6 Context Confusion without Index-Blindness?

Relativists might reply that their solution does not require attribution of any form of semantic blindness. Instead, the appearance of paradox can be explained by speakers' confusion of contexts of assessment alone: Speakers are not index-blind, they are merely confused about the epistemic standards relevant to their truth value judgments. Given the nature of sceptical arguments, this confusion is neither surprising nor is its attribution a theoretical cost for relativists.

The exclusive appeal to confusion of contexts is tempting, but it falls short of explaining the paradoxical appearance of the sceptical argument. Let us be clear about the proposal. Speakers are supposed to be confused about the contexts of assessment they are occupying, yet to be fully competent with respect to the index-sensitivity of knowledge sentences. Call such speakers context-confused yet index-competent. Due to their context-confusion speakers do not notice the switch of contexts, which explains why they arrive at the puzzling combination of truth value judgments.³³

Now remember what the paradox consists in: We are inclined to accept the premises as true and reject the conclusion as false (affirm the unnegated conclusion), but we are also inclined to think that the conclusion's truth follows from the truth of the two premises (we are inclined to think that the two premises and the unnegated conclusion are inconsistent). The problem with the idea of index-competent but context-confused speakers is that it cannot explain both speakers' robust truth value judgments and their sense of inconsistency. Yet both are needed to explain why (SA) appears to present a paradox.

To see this, notice, on the one hand, that the appearance of paradox only arises from (SA) if speakers have robust truth value judgments in the first place. Speakers judge the premises true and the conclusion false. So index-competent speakers, whose truth value judgments are sensitive to the epistemic standards of their context of assessment, cannot be at a loss about which epistemic standards to use in evaluating (SA)'s premises and conclusion for truth and falsity. If they were, they would not arrive at robust judgments at all. Thus, speakers' truth judgments of the premises are competently informed by $C_{A(High)}$ and their falsity judgments of the conclusion by $C_{A(Low)}$. They switch contexts. On the other hand, to explain why index-competent speakers take the truth of the premises and the unnegated conclusion to be inconsistent, we must hold that

³³For illustration, consider a *plausible* case of a context-confused but index-competent speaker. Suppose that on a rainy St. Andrew's Day, John falls into a coma, from which he awakens three weeks later on a sunny day without any memory of having been in a coma, nor any awareness of time having passed since St. Andrew's Day. Suppose John then incorrectly judges 'It's raining' to be true. Given a temporalist semantics, his judging so can be explained by the fact that John is mistaken about the time of his context of use. Unaware of time's passing, he assumes it is still St. Andrew's Day, when it *was* raining. The explanation of John's mistake in judging need not declare him blind to the time-parameter of the index (given a temporalist semantics). All John suffers from is mislocation in time. He is ignorant of his actual context, yet he is fully index-competent.

they are judging from the same context of assessment, i.e. they do not switch contexts. For otherwise index-competent speakers would not get a sense of inconsistency. After all, the premises' truth at $C_{A(High)}$ is only inconsistent with the unnegated conclusion's truth at $C_{A(High)}$. It is not inconsistent with the unnegated conclusion's truth at $C_{A(Low)}$. So we must say that index-competent yet context-confused speakers take themselves to be at $C_{A(High)}$ and $C_{A(Low)}$ at the same time, or that they unknowingly oscillate between the two contexts. But then we lose an explanation of their unchanging, robust truth value judgments: index-competent speakers only have stable truth value judgments if they competently pick up on their context's epistemic standards, i.e. if they are not oscillating between or confused about their context.

In sum, we only get a full explanation of the appearance of paradox if we drop the assumption of index-competence and assume a certain degree of index-blindness. Speakers switch contexts of assessment and thus arrive at their truth value judgments, and they take these to be inconsistent because they are blind to the fact these judgments are informed by their contexts of assessment.

2.7 Epistemic Closure Puzzles

As a matter of fact, relativists have not addressed sceptical paradoxes in any detail.³⁴ So they might try to dodge the objection from sceptical paradox by rejecting any relativist solution. As concerns sceptical paradoxes, they might say, relativism has no advantage over traditional invariantist theories. Whatever works as a solution of the paradoxes for the latter will do for the former.

This rejection is not promising. Remember that relativists appeal to retraction data in support of their view. When John is presented with the possibility that thieves might have stolen his car, he will retract his earlier claim 'I know that my car is parked in the driveway.' Relativists explain this by pointing out that while in the earlier context Low it was right to judge the knowledge claim as true, the standards have been raised to HIGH, so judged from John's later context of assessment, the earlier claim is false. Now, consider this anti-sceptical argument corresponding to (SA):

(AS) I know that I have hands.

If I know that I have hands, then I know that I'm not a BIV.

I know that I'm not a BIV.

³⁴Perhaps with one exception: Mark Richard seems to advertise relativism as a plausible development, or fix of, epistemic contextualism and recommends the 'insight into what is going on in skeptical arguments' offered by contextualism as one of the view's attractions (2008, 167). Presumably, then, Richard would not be tempted to reject the relativist treatment of sceptical paradoxes in section 2.5 above.

The relativist story in line with the solution to sceptical paradox will say that, given the mention of a sceptical hypothesis, we will switch contexts of assessment from LOW to HIGH in moving from premises to conclusion. We will thus find the premises true at $C_{A(Low)}$ but the conclusion false at $C_{A(High)}$. But notice the similarity between the explanations of retraction data and of this anti-sceptical paradox. The crucial point is the switch of contexts of assessment. If relativists make so much of their explanation of retraction data, on what grounds could they refuse the similar explanation in the case of (anti-)sceptical paradox?

It cannot be the peculiarity of external-world scepticism that supports the relativist's potential rejection of the relativist solution. The error-theoretic objection is not limited to sceptical paradoxes concerning BIVs, evil demons, and the Matrix. It generalises to paradoxes from everyday conversation data about people, things, or activities – data that relativists are keen on explaining. Consider the following argument based on Vogel's (1990) car theft case – data that MacFarlane readily appeals to in support of relativism:³⁵

- (C) I know that my car is parked in the driveway.
 If I know that my car is parked in the driveway, then I know that it has not been stolen.
-
- I know that my car has not been stolen.

Such 'semi-sceptical' cases (to borrow Vogel's phrase) are equally puzzling, and similar formulations of them have been widely discussed in the literature on epistemic closure.³⁶ But the relativist solution to external-world sceptical paradox is an equally good explanation of why we find these cases puzzling. So relativists would have to reject the relativist solution across the board, which undermines their contention to explain data concerning the use of 'know' in everyday situations. But everyday talk about cars, lotteries, and zebras is the very data with which relativists have made their case.

³⁵Other cases include lottery-style cases and Dretske's zebra case (see Dretske (1970); Hawthorne (2004); Vogel (1990)).

³⁶A simple version of single-premise epistemic closure principle is (EC) (cf. Hawthorne (2004, 31-50) for a discussion and refinement of the principle):

- (EC) Necessarily, if S knows p and S knows that p entails q , then S knows q .

Note that we can formulate the sceptical argument by explicitly using a contraposition instance of (EC) and the plausible assumption that one knows that one's having hands entails that one is not a BIV: (1) I don't know that I'm not a BIV. (2) Necessarily, if I don't know that I'm not a BIV, then either I don't know that I have hands or I don't know that my having hands entails that I'm not a BIV. (3) I know that my having hands entails that I'm not a BIV. (4) Therefore, I don't know that I have hands.

2.8 Conclusion

My purpose in this chapter has been to show that relativism about knowledge sentences requires a semantic blindness thesis to explain the data from speakers' use of knowledge sentences. I argued that this cannot be shown by appeal to speakers' alleged use of a relativised truth predicate, or of any natural language stand-in for this predicate. Relativists have good reasons to reject the data, and even if they were to accept it, they could amend their semantics to accommodate the data. I suggested a simpler way to show the relativist's need for an error theory. An adequate relativist solution to sceptical paradoxes and epistemic closure puzzles systematically commits relativists to the attribution of a particular kind of semantic blindness to speakers: epistemic assessment-index-blindness, i.e. blindness to the fact that the truth value of contents expressed by sentences involving the expression 'know' can vary with the epistemic standards coordinate in the index, which is determined by the context of assessment. This solution falls naturally out of the relativist semantics, and as I argued, its rejection would undermine the relativist's explanation of a wide variety of data from speakers' use of knowledge sentences. But index-blindness is highly implausible in light of speakers' robust competence with other kinds of index-sensitivity. Being committed to index-blindness, relativists lose what MacFarlane advertised as their main advantage over invariantist and contextualist competitors: the ability to predict the empirical data while avoiding the 'double-edged sword' of positing speaker error.

KNOWLEDGE ATTRIBUTIONS AND SPEAKER ERROR

3.1 Introduction

The case for relativism about knowledge attributions, as championed by MacFarlane (2005a, 2011b), is based on three claims. First, all competing views of contextualist and invariantist stripes make some incorrect predictions regarding ordinary speakers' use of 'know.' Their best account of speakers' use, in the troubling cases, attributes systematic error, or blindness, to speakers. Second, all of these error theses are implausible. Worse, there is a principled worry about attributions of systematic error. In MacFarlane's own words,

a *general* problem with positing speaker error to explain away facts about use is that such explanations tend to undermine the evidential basis for the semantic theories they are intended to support. All of these semantic theories are justified indirectly on the basis of facts about speakers' use of sentences, and the more error we attribute to speakers, the less we can conclude from these facts. (2005a, 215)

The third claim is that relativism escapes this problem because it makes the correct predictions for all use facts and thus need not appeal to speaker error.

In chapter 2, I argued that this last claim is false. In order to account for our judgments about sceptical paradoxes and epistemic closure puzzles, the relativist must attribute a particular kind of semantic blindness, i.e. assessment-index-blindness. That is, speakers are blind to the fact that the truth value of knowledge sentences can vary with the epistemic standards coordinate in the index, which is determined by the context of assessment.¹ This attribution of blindness leaves the relativist's position itself threatened by the general problem.

So is the relativist back to square one, in no better position than any of the views that need error attributions? This question receives a positive answer only if all error attributions are equally implausible. But one need not accept that all error is alike. It would not come as a surprise if error-theoretic objections varied in their force. It is easy enough to see how they would do so. Here is the template for error-theoretic objections.

¹For the definition of index-blindness, see section 2.5.

ERROR-THEORETIC OBJECTION TEMPLATE

- (E₁) Theory *X* makes the wrong predictions about some set of facts *Y* regarding ordinary speakers' use of expression *Z*.
- (E₂) To account for use facts *Y*, theory *X* must appeal to systematic speaker error.
- (E₃) But this attribution of speaker error is implausible.

Where (E₁) and (E₂) are true, the force of the error-theoretic objection depends on (E₃). (E₃) ascribes implausibility. But implausibility comes in degrees. Surely, an extreme view to the effect that no error attribution in semantic theorising is ever any plausible has little credibility. Humans are fallible, and there may be good reasons why they may even persistently make systematic errors.² Moreover, this extreme view would provide no positive guidance in deciding between a number of views all of which are committed to attributing some error or other.

This, however, is exactly the dialectical situation we are facing. As I will argue in this chapter, *all* extant views need to attribute error of some kind to account for some of the data, even where they are supplemented with a pragmatic or other account for the problematic data. A reasoned decision between competing views thus needs to be based on the relative costs and merits of their respective error attributions. But while the philosophy of language has generated extensive debate concerning the adequacy conditions for semantic as well as for pragmatic theorising, it is far from clear what the adequacy conditions for error theoretic explanations are. The present chapter is an attempt to identify criteria for the evaluation and comparison of different kinds of error and their attribution. These criteria will provide the basis for a reasoned decision between invariantist, contextualist, and relativist views. They should also be useful in other areas of philosophical theorising where error attributions are made with the intention of accounting for ordinary speakers' use of the philosophically interesting expression.³

I will proceed as follows. First, I will show that relativism's competitors – versions of contextualism and invariantism about knowledge attributions – all need to appeal to speaker error in their explanation of some use facts, especially in their solution to sceptical paradoxes (§§3.2–3.5). I will identify the kinds of semantic, pragmatic, and metaphysical error the theories must attribute, and the data points for which they must attribute error. This will allow us to put all views and their error attributions side to side on the scoreboard (§3.6). I will then provide criteria for the evaluation and comparison

²See for instance the literature on heuristics and biases in Kahneman et al. (1982) that describes strategies that people use in probability and frequency judgments. These strategies are cognitively efficient, often correct, but they lead to predictable error in a significant range of cases.

³Some areas where the results of this chapter may usefully be applied are the debates on expressions of causality, vague expressions in natural language, and ascriptions of free vs forced action.

of different error attributions (§3.7). Finally, I will suggest a strategy relativists can pursue in defence of their error attributions (§3.8).

3.2 Nonindexical Contextualism and Use-Index-Blindness

We saw in chapter 2, section 2.2, that indexical contextualism, the view that the content expressed by an occurrence of ‘*S* knows [does not know] that *p*’ depends in part on the epistemic standards salient in the conversational context, has trouble accounting for speakers’ judgments of knowledge attributions made in contexts in which different epistemic standards are salient. Moreover, disagreement and retraction data pose a problem for indexical contextualism. We concluded that the contextualist’s best explanation of the data commits her to the attribution of content-blindness to ordinary speakers.

It might be thought that nonindexical contextualism provides a better alternative, one that preserves the key contextualist insights while avoiding the attribution of content-blindness. Nonindexical contextualism (Brogaard, 2008; Kompa, 2002; MacFarlane, 2009) can be seen as an intermediate view between indexical contextualism and relativism. It is the view that sentences of the form ‘*S* knows [does not know] that *p*’ are epistemically use-index-sensitive. That is, there is an epistemic standards parameter in the index which is determined by the context of use or utterance. More precisely, nonindexical contextualism can be defined by the following three theses.

1. INVARIANT CONTENT

‘*S* knows that *p*’ expresses a content that is invariant across contexts of use (bar indexical expressions in the substitution instances of ‘*S*’ and ‘*p*’).

2. INDEX SENSITIVITY

The index contains an epistemic standards coordinate, to which the truth value of ‘*S* knows that *p*’ is sensitive.

3. USE-SENSITIVITY⁴

The truth value of ‘*S* knows that *p*’ depends on the epistemic standards salient at the context of use. Put together with 2, the epistemic standards-coordinate in the index is determined by the context in which a knowledge sentence is used.

Nonindexical contextualism thus combines the contextualist’s insight that the truth of knowledge attributions depends on the epistemic standards salient in the *context of use* with the relativist’s insight that the variability is one of truth value, not of content.

⁴The term ‘use-sensitivity’ is used in MacFarlane’s sense. See for instance MacFarlane (2012, ch.3) for discussion.

Let us examine nonindexical contextualism's predictions about the core use facts that have been identified in the debate as problematic for at least some views. A quick note before we proceed. I will discuss the same set of core data for each view in order to keep the discussion focused. This is not to say that for any particular view, these data exhaust the use facts with which the view may have trouble. I urge the interested reader to keep an eye on the footnotes for additional discussion.

Basic Variability

Recall that indexical contextualists motivate the variability of content by appeal to ordinary speakers' use of 'know.' While speakers accept many knowledge attributions as true in mundane contexts of use, they tend to give in to, e.g., sceptical considerations that raise the epistemic standards and reject these attributions as false in such contexts. It is eminently natural to believe and assert in everyday situations that one knows that one has hands. But once sceptical considerations are brought forward, it is equally plausible to give that belief up and take to be true, and assert, that one does not know that one has hands. What varies here are speakers' propensities to make knowledge attributions or denials, as well as their judgments about the truth/falsity (or acceptability) of knowledge attributions or denials uttered in the very context they occupy. Call this data point *Basic Variability*.

Nonindexical contextualists explain the variability in intra-contextual truth/assertability judgments in the same way as indexical contextualists. On both views, the epistemic standards relevant for the truth evaluation of a sentence in context is determined by the context in which the sentence is uttered.

Inter-contextual Truth Ascriptions

The shared contextualist claim that the context of use determines the epistemic standard relevant for judgments of truth/assertability is also responsible for both contextualist views' failure to predict inter-contextual truth ascriptions. Speakers tend to judge knowledge claims true or false by the standards salient in their context of evaluation. Once error possibilities are taken seriously, speaker not only judge that a given subject does not have knowledge, they also judge knowledge attributions made in contexts with lower epistemic standards as false. Nonindexical contextualism's best explanation of this fact is that speakers are in error when they judge knowledge claims made in contexts that differ vastly in epistemic standards from those in their own context.

(Dis)Agreement

Nonindexical contextualists have roughly the same resources to explain disagreement data as relativists. When speaker A in Low asserts 'S knows that p ' and speaker B in HIGH asserts 'S does not know that p ,' then there is a content, *S knows that p ,*

which is expressed in A's assertion and whose negation is expressed in B's assertion. So while both speaker's claims are true in their respective contexts, their disagreement can be explained in terms of the incompatibility of their expressed contents.⁵ This is a significant advantage over indexical contextualism.

Retraction

Nonindexical contextualism does not make the right prediction about speakers' retraction of knowledge attributions. Consider John's retraction of his earlier claim 'I know that my car is parked in the driveway' after hearing from Mary that car thieves are roaming his neighbourhood (cf. section 2.2). According to nonindexical contextualism, it would be correct for John to say 'I guess I don't know that my car is parked in the driveway. But when I said before that I do know it, I spoke truly.' However, this reply is no more natural than the one that indexical contextualism licenses. Nonindexical contextualism is committed to the claim that speakers are in error when they instead say 'I was wrong,' 'I spoke falsely,' or 'My earlier assertion is false.'

Inter-contextual Assessment of Truth Ascriptions to Knowledge Claims

In chapter 2 we considered the dialogue in (2.1), taking place in Low:

- (2.1) John: We both know that Neil Armstrong was the first man to set foot on the moon.
 Bob: That's true.

Nonindexical contextualists can predict that Mary in HIGH judges as false both John's knowledge attribution and Bob's truth ascription. Nonindexical contextualist may assume that 'true' is ascribed to propositions and is governed by the equivalence schema *It is true that p iff p*. By the equivalence schema, the proposition Bob expresses is true just in case the proposition John expresses is true. Thus, Mary in HIGH can truly assert 'That's false' in talking about the proposition Bob expresses just in case relative to the high epistemic standards in her context, it is false that John and Bob both know that Neil Armstrong was the first man to set foot on the moon.⁶

Sceptical Paradox

Recall from chapter 2 the sceptical argument.

⁵For discussion of a nonindexical contextualist account of disagreement, see section 6.2.5 below.

⁶This prediction has some odd consequences for nonindexical contextualism that result from the bifurcation of the truth of utterances and ascriptions of propositional truth (cf. MacFarlane (2009, 246-9)). I will postpone discussion of these consequences until chapters 5 and 6, where we will investigate the nonindexical contextualist view in more detail.

(SA) I don't know that I'm not a BIV.

If I don't know that I'm not a BIV, then I don't know that I have hands.

I don't know that I have hands.

In chapter 2, section 2.5, I argued that relativism needs to appeal to semantic blindness if they are to give a solution to the sceptical paradox – a solution that explains both why (SA) does not in fact present a paradox and why it seemed to us to be a paradox. The semantic resources that nonindexical contextualism brings to the table mirror those of relativism. Unsurprisingly, then, the nonindexical contextualist solution to the paradox has similar virtues and vices as the relativist solution. It can explain why (SA) does not give rise to paradox by pointing out that *as long as the context of utterance is kept fixed*, the argument is valid but it is not the case that both premises are true while the conclusion is false. It can also explain the air of paradox by pointing out that speakers mistakenly switch contexts in moving from premises to conclusion and are unaware of doing so. This latter claim, however, commits nonindexical contextualism to the claim that speakers are in error in their use of knowledge attributions in (SA).⁷

In sum, nonindexical contextualists are committed to ascribing error to speakers in cases of inter-contextual truth ascriptions, retraction, and sceptical paradox. However, this error is not a form of content-blindness. On the contrary, what speakers are blind to is the use-index-sensitivity of knowledge sentences. They ascribe truth and falsity depending on the epistemic standards *in their own current context* and thereby ignore the fact that knowledge sentences are sensitive to the epistemic standards in the index as determined by the *context of utterance*. This ignorance of the semantical workings of knowledge sentences is a form of index-blindness. In chapter 2, index-blindness was defined as follows.

INDEX-BLINDNESS

Speakers are blind to the fact that the truth value of contents expressed by sentences involving a particular expression ('know') can vary with a particular co-ordinate in the index (epistemic standards).

In particular, we can say that nonindexical contextualism is committed to the claim that speakers are epistemically use-index-blind. Is this kind of blindness plausible? In chapter 2 I argued that epistemic *assessment*-index-blindness is implausible in light of speakers' robust competence regarding non-epistemic forms of index-sensitivity. The

⁷For details see section 2.5. Note that the difference between nonindexical contextualism's use-sensitivity and relativism's assessment-sensitivity is immaterial to the explanation of at least those cases of sceptical paradox in which assessor and speaker – context of use and context of assessment – are identical.

same point also holds for epistemic *use-index-blindness*. As a result, the view cannot explain all the data without recourse to an implausible semantic blindness thesis.

3.3 Moderate Insensitive Invariantism and Pragmatic Blindness

Indexical contextualists, nonindexical contextualists, and relativists are all burdened with some kind of semantic blindness thesis. Perhaps, then, some form of invariantism is the way to go.

Invariantists hold that neither the content expressed by ‘*S* knows [does not know] that *p*’ nor its truth value varies with the epistemic standards salient in the context of utterance. Traditional, insensitive invariantists moreover believe that whether or not a subject knows does not depend on the knowing subject’s attention to counterpossibilities, interests, or practical stakes. Traditional invariantists can be further distinguished. *Moderate insensitive invariantists* have it that the constant epistemic standards, i.e. the requirements on our evidence, reliability, etc. to count as knowing, are reasonably low. We know a lot of things. *Sceptical insensitive invariantists* think that the epistemic standards are fixed at an extraordinarily high level, so high in fact that we hardly ever meet them. In contrast to these traditional views, *interest-relative invariantists* (also called subject-sensitive invariantists) such as Fantl and McGrath (2002, 2009), Hawthorne (2004), Stanley (2005), and Weatherson (2011) believe that the truth of knowledge sentences depends on the epistemic standards in the subject’s context, and these are determined by the subject’s attention to counterpossibilities, interests, and practical stakes.

Before looking at the views and their predictions one by one, let me make one observation about all of them.

Basic Variability (all invariantist views)

All invariantist views have at least some trouble with the variability in truth value judgments about knowledge attributions. In everyday conversational contexts we are confident in ascribing to John the knowledge that he has hands by saying ‘John knows that he has hands,’ and in these contexts we judge this sentence as true. But when epistemic standards are raised, in the epistemology classroom or in court, we deny that John knows that he has hands, even when John’s situation has not changed in any respect. In the latter context, we think the sentence ‘John knows that he has hands’ is false. Invariantists cannot account for this variability in intra-contextual truth ascriptions to knowledge attributions because they hold that epistemic standards are fixed once and for all by facts about the subject’s situation: Either John meets the standards and knows or he does not meet them and does not know. Moderate invariantists must claim that speakers who deny knowledge in sceptical contexts are in error, or they must find an

alternative explanation. Sceptical invariantists are committed to the claim that speakers who attribute knowledge in ordinary contexts are mistaken, or they must find an alternative explanation. Interest-relative invariantists must claim that as long as John's situation does not change, one of the speakers is wrong.

Let us now turn to moderate insensitive invariantism, the view that the invariable epistemic standards that a subject has to meet in order to count as knowing are reasonably low. In response to the view's problem with basic variability data, Brown (2005, 2006) and Rysiew (2001, 2007) propose a WAM – a 'warranted assertability manoeuvre' – on behalf of moderate insensitive invariantism. A WAM is the attempt to explain speakers' judgments about the correctness or incorrectness of knowledge claims by arguing that these judgments reflect not the truth conditions of knowledge sentences in context but their warranted assertability conditions. Brown and Rysiew hold that '*S* knows [does not know] that *p*' literally expresses, simply and invariantly, the proposition *that S knows [does not know] that p*. The sentence's truth conditions do not depend on any epistemically relevant factors at the context of use. But on a given use, the sentence also pragmatically conveys, via familiar Gricean mechanisms, the content that *S* is in a very strong epistemic position with respect to the proposition that *p*.⁸ Brown uses sensitivity to spell out the notion of strength of epistemic position. On her account, it is pragmatically conveyed that *S*'s belief that *p* matches the facts across a wide range of worlds, where the context of utterance settles just how wide this range is.⁹ Rysiew adopts the framework of relevant alternatives theory. On his account, what is pragmatically implicated is that *S* can rule out the contextually salient (but irrelevant) alternatives.¹⁰ According to both Brown and Rysiew, whether a knowledge claim is warrantably assertable is a matter of the truth or falsity of its pragmatically conveyed content, which may vary from one context of utterance to another depending on the conversationally accepted demands on *S*'s strength of epistemic position (the range of worlds in which

⁸Two points are important to stress. First, Brown and Rysiew claim that their accounts exploit general and independently plausible Gricean mechanisms of implicature that are needed for a theory of communication anyway. So the accounts are theoretically parsimonious. Second, they claim that their accounts exploit mechanisms that have well-studied precedents. For instance, Brown (2006) cites the Kripke-Donnellan-Bach approach to definite descriptions, where in some cases, what is literally said is false but what is pragmatically conveyed is true. Rysiew and Brown point to cases where what is literally said differs from a richer, pragmatically conveyed content, as in 'I have eaten,' which literally expresses just that the speaker has eaten at some time in the past but conveys that the speaker has eaten recently. If implicatures triggered by knowledge sentences have suitable precedents in such implicatures, Brown's and Rysiew's accounts avoid the charge of *ad hocness*.

⁹See (Brown, 2006, 424–6). DeRose (1995, 35) defines sensitivity as follows: 'For one's belief that *p* to be sensitive one must not believe that *p* in the closest not-*p* worlds.' More intuitively, one's belief that *p* is sensitive just in case, if it were not the case that *p*, one would not believe that *p*.

¹⁰On Rysiew's (2001) moderate insensitive invariantism, the alternatives to the claim that *p* that are *relevant* for *S* to possess knowledge that *p* are constant and meetably lax, yet in sceptical contexts the *salient* alternatives may well exceed the relevant ones. The differences in detail between Brown's and Rysiew's account do not matter for the argument in this section. For simplicity, I shall stick with Rysiew's account.

S's belief must match the facts, or the range of alternatives that *S* must be in a position to rule out). Call the moderate insensitive invariantist view that is complemented by one of these pragmatic accounts *sophisticated moderate invariantism*.

It is an open question whether Brown's and Rysiew's accounts are linguistically plausible.¹¹ In order to account for the data, however, even the best version is committed to the claim that speakers confuse literal content and pragmatically implicated content and are even blind to pragmatically implicated content.

Basic Variability

Consider first the basic variability data. Moderate insensitive invariantists need to explain why in HIGH, speakers are inclined to make knowledge denials, and why they judge knowledge denials true and knowledge attributions false. Brown and Rysiew claim that speakers' judgments in HIGH track the contextually variable warranted assertability conditions of knowledge sentences, not their truth conditions. Thus, a knowledge attribution made in HIGH pragmatically implicates the false content that the subject is in a position to rule out the contextually salient error possibilities (which include possibilities the subject need not be able to rule out to count as knowing).

Note that speakers judge the knowledge claims made in their context to be *true* or *false*, not just warrantably (un)assertable or acceptable. If speakers, making a knowledge attribution in HIGH, were aware of the fact that what they literally say is true and what they pragmatically convey is false, they could naturally defend their knowledge attribution by saying 'Of course what I implied was false, but what I literally said was true' or 'What I said was true, although I wasn't warranted in asserting it.' Yet this is not what speakers say. Of course, it is unreasonable to assume that speakers' judgments track the theoretical distinction between semantic and pragmatic content. Rysiew (2007, 660, n.31) himself points this out.

But, of course, that our pretheoretical intuitions as to what we're 'saying' are generally insensitive to the semantic/pragmatic distinction is essential to the view being presented here. And, to the extent that speakers *are* reading merely pragmatically-generated information onto the relevant sentences, it's not that they will utter a sentence that they know to be false [...] only so as to communicate some further, true information (e.g., that Smith does not meet certain unusually high epistemic standards); whether they see themselves as 'saying' something true/false will depend on the perceived truth value of the information that is merely pragmatically generated.

¹¹ See Dimmock and Huvenes (2012) for relevant discussion

It is plausible that speakers' judgments are not generally sensitive to the semantics/pragmatics distinction. But Brown's and Rysiew's views require a stronger claim, namely that speakers have no means whatsoever to distinguish between what they say with their knowledge claims and what they merely implicate. This stronger claim is implausible and reveals a disanalogy between the implicatures of knowledge claims and typical cases of pragmatic implicatures.¹² In many of the standard cases of implicature, speakers *are* in fact in a position to distinguish between what they literally say and what they merely pragmatically convey. A case of pragmatic implicature that Brown (2006, 416–7) mentions is that of a speaker saying 'I have eaten,' which literally says that the speaker has eaten at some time in the past but which in normal circumstances pragmatically conveys that the speaker has eaten recently. In a case like this, it is not altogether unnatural for the speaker to respond to the challenge, 'Of course you've had food before in your life. What I want to know is whether you've eaten recently,' by saying 'I know I didn't say it literally, but what I meant was that I've eaten recently.' In contrast, it could not be further from common practice to defend a knowledge denial in HIGH by saying, 'Sure, I know that *S* in fact knows that *p*. What I meant was that he isn't in an even stronger epistemic position.'

Furthermore, pragmatic implicatures have the feature of being cancellable (Grice, 1975, 39).¹³ If 'X is meeting a woman this evening' typically has the conversational implicature that the woman in question is not X's wife, mother, sister, or platonic friend, this implicature can on occasion be cancelled by saying 'X is meeting a woman this evening. That woman is no one other than his wife.' In contrast, it is odd at best to cancel the implicature of a true knowledge attribution made in HIGH by saying, 'I know that my car is parked outside. Of course, I cannot rule out that it has been stolen by car thieves.'

The cancellability of conversational implicature and speakers' ability to defend their claims by pointing out what they meant to communicate rather than what they literally said suggest that although speakers' truth judgment may not generally be sensitive to the distinction between semantics and pragmatics, speakers do possess some resources to keep apart semantic content ('what is said') and pragmatically implicated content ('what is meant'). Yet they do not appear to be able to keep apart the semantic and

¹²Brown prefers to model the content pragmatically by knowledge claims after Bach's (1994) cases of 'implicature' (such as the 'I have eaten' case below). The differences between paradigm cases of Gricean conversational implicature and cases of Bach's implicatures should not matter to the points I wish to make. I will speak of pragmatic implicatures to cover all cases of pragmatically implied rather than literally expressed content.

¹³Cancellability is widely considered to be one of the most robust properties of what Grice calls conversational implicatures. For some alleged counterexamples, however, see Weiner (2006). In contrast to conversational implicatures, conventional implicatures are not supposed to be cancellable. Neither Rysiew nor Brown give any indication claim that the implicatures of knowledge attributions are conventional implicatures, so I will put that option to one side.

pragmatically implicated content of knowledge claims in the way that they do with typical implicatures. Recording just this fact, we may say that sophisticated moderate invariantists are committed to the claim that speakers systematically confound semantic and pragmatic content. In what follows, I will use the slightly overstated title *semantics/pragmatics confusion* to refer to just this fact. Semantics/pragmatics confusion with respect to knowledge attributions – as manifest in cancellability facts and responses to challenges – is implausible. If knowledge claims in fact trigger pragmatic implicatures in a wide range of cases, these implicatures should give rise to similar use facts as uncontroversial cases of implicatures do.

Disagreement

Disagreement data shows that speakers in different contexts, LOW and HIGH, can felicitously be reported to disagree if one says ‘*S* knows that *p*’ and the other says ‘*S* does not know that *p*.’ This would be surprising if our judgments focused on pragmatically conveyed content, as Brown and Rysiew claim. For in LOW, it is pragmatically conveyed that *S* can rule out a reasonably small range of not-*p* alternatives, whereas in HIGH it is conveyed that *S* cannot rule out an exceptionally wide range of not-*p* alternatives. These claims do not conflict, so we should not be led to think the speakers disagree. To explain disagreement, Brown and Rysiew must hold either that disagreement judgments focus on literal content and truth conditions, and truth judgments on pragmatic content and warranted assertability conditions, or that we mistakenly take the speakers’ pragmatically implicated contents to be in conflict. In the former case, we are implausibly oscillating between literal and pragmatic content without being aware of it, in the latter we are blind as to what the pragmatically implicated contents are. Thus, sophisticated moderate invariantism cannot account for disagreement data without appealing to blindness.

Inter-contextual Truth Ascriptions

Brown’s and Rysiew’s accounts make the wrong predictions about speakers’ inter-contextual truth judgments. Consider a speaker in HIGH, who judges a knowledge attribution made in LOW to be false. If that speaker properly understood what is said and implicated by the knowledge attribution in LOW, she should have no qualm with it. After all, the knowledge attribution is both true and warrantably assertable in LOW. What is literally said by it is true, and what is pragmatically implicated – that *S* is in a position to rule out a reasonably small range of not-*p* alternatives – is not in conflict with the speaker’s beliefs in HIGH. Here, however, appeal to speakers’ confusion of literal and pragmatic content is not sufficient. To explain speakers’ falsity judgments, Brown and Rysiew must hold that speakers are blind to pragmatic content. They are ignorant of the fact that the pragmatic implicatures of knowledge claims vary in content depending

on the epistemic standards at the context of utterance. They take the pragmatic content implicated by knowledge sentences in *other* contexts of utterance to be the pragmatic content these knowledge sentences would implicate were they uttered in the speakers' *own* contexts.

Retraction

Speakers also tend to retract knowledge attributions, made in LOW, in light of raised epistemic standards by saying 'I was wrong,' 'What I said was false,' or 'I take that back.' Brown's and Rysiew's accounts make it a mystery why speakers would retract. After all, they did not say anything false in LOW, nor did they pragmatically convey anything false in LOW. So they should instead insist that what they said (and implicated) is true. Here, again, appeal to speakers' confusion of literal and pragmatic content is not sufficient to explain retraction. Speakers have no reason to retract either literal or pragmatic content.¹⁴ To explain retraction, Brown and Rysiew must hold that speakers are blind to the fact that pragmatic implicatures of knowledge claims vary with the epistemic standards at play in the *context of utterance*. Although in LOW they in fact implicated that the subject is in a position to rule out the alternatives *salient in LOW*, they now take themselves to have implicated that the subject is in a position to rule out the wider range of alternatives *currently salient in HIGH* and hence retract their earlier assertion. Moreover, they are confounding literal and pragmatic content, which explains why they do not insist that what they said is (literally) true.

Inter-contextual Assessment of Truth Ascriptions to Knowledge Claims

Consider again dialogue (2.1), taking place in LOW:

- (2.1) John: We both know that Neil Armstrong was the first man to set foot on the moon.
 Bob: That's true.

On the moderate invariantist semantics, John's utterance literally expresses a true content that does not vary with epistemic standards. Bob's truth ascription is invariably true, too. Thus Mary in HIGH is wrong in judging both John's and Bob's claims false. Her error can be explained, however, by appeal to semantics/pragmatics confusion and pragmatic blindness. Mary takes John to pragmatically implicate what his utterance would implicate if it were made in her context HIGH, namely that they are both in a position to rule out all those alternatives, salient in HIGH, in which Neil Armstrong was not the first man to set foot on the moon. This implicature is false, hence Mary's judgment. Yet to arrive at her falsity judgment, Mary must not only confuse literal and

¹⁴See also Dimmock and Huvenes (2012) for relevant discussion

pragmatic content but also be blind to what the content that is in fact pragmatically implicated by John in Low is. Furthermore, Mary also takes Bob's assertion to be about what John's assertion would pragmatically implicate in HIGH. So she in fact takes him to ascribe warranted assertability to John's utterance. She is not aware of this, however, because she is ignorant of the distinction between the literal and the pragmatic content of John's utterance.¹⁵

Sceptical Paradox

Brown's and Rysiew's WAM also affords moderate insensitive invariantism a solution to sceptical paradox. According to moderate insensitive invariantism, premise one, 'I don't know that I'm not a BIV,' and premise three, 'I don't know that I have hands,' are false. Hence the sceptical argument – 'I don't know that I'm not a BIV./If I don't know that I'm not a BIV, then I don't know that I have hands./Therefore, I don't know that I have hands' – is not sound. Simply rejecting speakers' truth judgment of premise one, however, does not yield a satisfactory solution in Schiffer's sense. It does not explain why the argument *seems* to present a paradox. Adding Brown's and Rysiew's WAM, moderate insensitive invariantist can explain why the argument seems paradoxical. Premise one mentions a sceptical hypothesis and thereby introduces a context with high epistemic standards, HIGH. The pragmatically implicated content of premise one in HIGH is that the speaker cannot rule out an exceptionally wide range of (not-not-)BIV alternatives. This implicature is true, so we come to judge premise one as true. Premise two is true and warrantably assertable in HIGH (as well as in Low).¹⁶ The conclusion is false on the moderate invariantist view. Moreover, we seem

¹⁵ Dialogue (2.2) in chapter 2 involved the inter-contextual assessment of a putatively relativised truth ascription to a knowledge claim.

- (2.2) John: We both know that Neil Armstrong was the first man to set foot
on the moon.
Bob: That's true relative to this context.

Sophisticated moderate invariantists can respond to (2.2) in at least three ways. First, they could reject (2.2) as proper data (cf. relativism's first reply in section 2.4). Second, they could accept (2.2) but claim that 'relative to this context' is vacuous, since propositional truth does not vary with any epistemically significant features of contexts. Then they can explain Mary's falsity judgment just like they did for (2.1). Third, they could interpret Bob as making a true statement about Bob's knowledge attribution being warrantably assertable in their context. (Warranted assertability conditions vary with context on the moderate invariantist's pragmatic story.) Hence, Mary would be wrong in her falsity judgment. This could be explained by the combination of semantics/pragmatics confusion and pragmatic blindness that led her to believe that John's knowledge attribution is false. Being both ignorant of the distinction between truth and warranted assertability and of the indexical nature of the content that John's assertion pragmatically implicates, she judges what John's assertion would implicate in a HIGH context to be false and thus thinks Bob is wrong in his ascription.

¹⁶It is not obvious to me what the warranted assertability of a conditional amounts to. The most plausible understanding in this philosophical context seems to be that the conditional consisting of the antecedent's and consequent's pragmatically implicated contents is true. For premise two of (SA), the conditional

to have switched to a context with low epistemic standards, Low, where the content that is pragmatically implicated by the conclusion is that the speaker cannot rule out a reasonably small range of alternative possible worlds in which she has no hands. This strikes us as false. Were we still in HIGH, the pragmatically implicated content would be true, and, if pragmatic content is what our judgments are tracking, we would judge the conclusion true. In order to explain the appearance of paradox, the sophisticated moderate invariantist needs to claim that we switch context without being aware of the switch, and thereby arrive at three judgments we take to be inconsistent: Premise one is true, premise two is true, but the conclusion is false.

Importantly, this solution involves the attribution of speaker error: We are not aware of the context-switch. Moreover, the solution must also attribute semantics/pragmatics confusion to speakers. We come to think of the premises as *true* and the conclusion as *false*, and we let our intuitions concerning validity and consistency be guided by these judgments, because we confuse what our assertions literally express with what they pragmatically convey. And finally, we must be blind to pragmatic content. If we were aware of the pragmatic content of the premises, which we judge to be true, we would not find it inconsistent with the unnegated conclusion. In sum, then, adding Brown's and Rysiew's WAM allows moderate insensitive invariantists to provide a solution to sceptical paradox. But this solution commits the view to a context confusion thesis and a pragmatic blindness thesis.

To sum up, the various data show that even a sophisticated version of moderate insensitive invariantism involving a pragmatic story must charge speakers with a certain amount of blindness. Minimally, speakers exhibit what I called *semantic/pragmatics confusion*: Speakers confuse, in implausible ways, what their assertions (of knowledge sentences) literally express and what they pragmatically convey. Moreover, data from inter-contextual truth judgments and retraction show that sophisticated invariantists need a stronger form of blindness I have called *pragmatic blindness*: Speakers are blind to the fact that the content that their assertions (of knowledge sentences) pragmatically implicate can vary with the epistemic standards salient at the context of utterance.¹⁷

consisting of the antecedent's and consequent's pragmatically implicated contents in HIGH is the following: 'If I cannot rule out an exceptionally wide range of alternatives in which I'm a BIV, then I cannot rule out an exceptionally wide range of alternatives in which I have no hands.' This conditional is true if we give it a material implication reading: Both antecedent and consequent are true. However, I find fault with a natural language conditional reading, on which the consequent is supposed to follow, in some intuitive sense, from the antecedent. Not being able to rule out an exceptionally wide range of BIV-alternatives does not seem to imply not being able to rule out an exceptionally wide range of handlessness-alternatives. The latter range, it seems to me, might well not contain any BIV-alternatives.

¹⁷Williamson (2005) offers an alternative explanation of the data that troubles moderate invariantism. Instead of appealing to pragmatic implicatures, Williamson gives a psychological explanation of speakers' systematic error, appealing to the literature on heuristics and biases (e.g., in Kahneman et al. (1982)). A very similar explanation is given by Hawthorne (2004) on behalf of interest-relative invariantism. I will discuss it in section 3.5 below.

3.4 Sceptical Insensitive Invariantism and Pragmatic Blindness

Sceptical insensitive invariantism is the view that the epistemic standards that one must meet in order to know are extremely high. According to Unger (1971, 1975), knowledge implies certainty: In order to know, one must be able to eliminate, or rule out, every possibility of falsehood (where the domain of ‘every’ is not contextually variable). As a consequence, (nearly) all of our knowledge attributions are in fact false, contrary to speakers’ judgments. A variant of the view claims the bar for knowledge is slightly lower; we know little but not as little as the sceptic claims. The discussion in the present section is intended to apply to both of these positions, although the focus will be on the stronger version. Both versions entail that speakers are often in error about what they know. A quick error theory might attribute the error to semantic blindness: Speakers widely overestimate the possession of knowledge because they are blind to the semantics of ‘know.’ In particular, they are blind to the fact that knowing entails that one is able to eliminate every possibility of falsehood.

Instead of this rather crude and *ad hoc* error theory, sceptical insensitive invariantists have offered an explanation of speakers’ truth intuitions that exploits the distinction between literal content (‘what is said’) and pragmatically implicated content (‘what is meant,’ ‘what is conveyed’). I shall consider two versions, Schaffer’s (2004b) hyperbole view and Davis’ (2007) loose use view. Both take their cue from Unger (1975).

According to Schaffer, acceptable knowledge attributions are cases of hyperbole. On the sceptic’s semantics, ‘I know that I have hands’ entails that I can eliminate all possibilities of handlessness, which exaggerates the range of possibilities I can eliminate, and which is false. Compare this with uncontentious examples of hyperbole. In saying ‘I’m dying of thirst,’ or ‘That pool is a mile long,’ we are exaggerating for effect. What we literally say is false. But where we find the utterance acceptable, it implicates a truth, e.g. that I am very thirsty or that that pool is very long by the contextually salient standards. The hyperbolic effect may be understood in Gricean terms as a flouting of the maxim of quality, which enjoins one to make one’s contribution true. Since what is said is false, a hearer who takes the speaker to be cooperative will be led to infer that the speaker intended to convey something else, namely the information most similar to what is said which is true and makes the utterance cooperative. Similarly, when we attribute knowledge in Low by saying ‘I know that I have hands,’ what is literally said is false, but what is meant (conveyed, implicated) is that I can eliminate those possibilities of handlessness that are *relevant in the current context*. Since the current context is one with low epistemic standards, what is meant is true. I am in a position to eliminate the possibility that I have stumps etc. This explains why in Low, we judge the knowledge attribution true.¹⁸

¹⁸As Schaffer (2004b, 141) points out, one may also understand hyperbole in terms of Sperber and Wilson’s theory of Relevance. The details of either account do not make any difference to the points I wish to

Davis (2007) explains the basic variability of our truth judgments in terms of differences between strict and loose usage of knowledge sentences.¹⁹ Consider a case of loose use:

(3.1) THE TIME MEASUREMENT CASE

A. Wondering how hard the final exam was, I ask Mike how long he took to finish. He answers ‘Two hours.’ B. When Nora says that she took two hours and four minutes to finish the exam, Mike responds ‘You took even longer than me. It took me two hours and two minutes.’ (Davis, 2007, 407)

On Davis’ account, if what Mike said in B is true, what he said in A is false. But in A, Mike used ‘two hours’ loosely. What he meant (conveyed, implicated) is that it took him *about* two hours to finish. Once the conversational purposes require more precision in B, Mike gives more precise information. In both conversational contexts, Mike’s utterance is perfectly appropriate. While the literal content of his utterance in A is false, the implicated content is true.

Knowledge sentences, too, can be used loosely and strictly. Standards of strictness vary with contexts of utterance and allow for loose usage in Low contexts. While what is said by ‘I know that I have hands’ in Low is false, what is implicated is that I am close enough to knowing for contextually salient purposes, which may well be true. While Davis rejects the Gricean constraint that pragmatic implicatures be derivable from the cooperative principle and the conversational maxims, he holds that loose use still enables us to be cooperative and to efficiently contribute adequate information to the conversation. When the purposes of the conversation do not require us to commit to a subject’s knowing, it is sufficient to use knowledge attributions loosely and communicate that the subject is close enough to knowing. When the difference between knowing and being close enough to knowing matters to the purposes of the conversation, we use knowledge attributions strictly. This explains our truth judgments in Low and our falsity judgments in HIGH.²⁰

On both the hyperbole and the loose use account, which content is implicated by an utterance of a knowledge sentence depends in part on the context of utterance. (Notice the context-sensitive expressions ‘relevant in the current context’ and ‘for contextually indicated purposes’ in the formulation of these contents). So we can expect *sophisticated sceptical invariantism* – sceptical insensitive invariantism plus one of the

make.

¹⁹For the differences between hyperbole and loose use, see Davis (2007, 410).

²⁰As the car case (ch. 2) illustrates, loose uses of knowledge sentences may be more or less strict. Raising commonplace error possibilities raises the conversational standards of strictness, so what counted as being close enough to knowing will not do so anymore. Yet standards of strictness could be raised further. As far as I can see, it is merely a question of convention where to draw the line between loose and strict uses.

above pragmatic accounts – to run into trouble with some data that indexical contextualists struggled with. The subsequent discussion will mirror the above discussion of moderate insensitive invariantism at points. My discussion will again focus on the accounts' error-theoretic commitments and leave aside other issues concerning their linguistic plausibility.²¹

Basic Variability

Sophisticated sceptical invariantists, just like sophisticated moderate invariantists, rely on the distinction between a sentence's being true or false in a context, and its being appropriately assertable in a context. On Schaffer's and Davis' accounts, the latter is tied up with the truth or falsity of the pragmatically implicated content, the former with its literal content. While we can grant that speakers' correctness judgments often do not discriminate between truth and assertability, we should assume that they *can* and *will* if the difference is of importance to them. In standard cases of hyperbole, speakers may take themselves to have asserted something true by saying 'That pool is a mile long.' But when challenged – 'That can't be! The brochure says it's a 50m, long-course competition pool' – speakers are likely to point out that they did not mean that the pool *literally* measures a mile, just that it is very long for a community pool.²² Likewise, in THE TIME MEASUREMENT CASE above, Mike would not feel seriously challenged if his interlocutor said 'That can't both be true! You cannot have taken two hours as well as two hours and two minutes.' On the contrary, he is likely to clarify that what he meant before was just that it took him *about* two hours – two hours and two minutes, to be exact. However, when speakers are challenged after making a knowledge attribution, they do not tend to say 'I didn't mean to say that I *literally* know that I have hands,' or 'I only meant before that I am close enough to knowing given the purposes of our conversation back then.'²³ If these are nonetheless cases of hyperbole or loose use, then it must be the case that speakers are ignorant of the fact that they are engaging in a hyperbolic, or loose, use of knowledge sentences. They are blind to the contrast between what they said and what they meant, or implicated. But given that speakers do not manifest this ignorance in standard cases of hyperbole and loose use, its attribution in knowledge attributions cases is implausible.

²¹For further criticism of the hyperbole view, see Hawthorne (2004, 119–21).

²²This point has been made by MacFarlane (2005a, 206).

²³Schaffer (2004b, 140 n.3) points out that hyperbole can be highly formulaic and non-obvious. But as MacFarlane (2005a, 206) notes, the problem is not obviousness. Even non-obvious hyperbole must be deliberate: People would not be exaggerating if they believed what they literally say. And this deliberateness guarantees that speakers, at least when pressed, can differentiate between what they said and what they implicated. In Davis' words: 'In hyperbole we want hearers to notice the contrast between what was said and meant.' (Davis, 2007, 410) See also Hawthorne (2004, 120).

Retraction

Sceptical invariantists may explain speaker behaviour under challenge by pointing to their sceptical semantics. When speakers who have made a knowledge attribution are challenged with error-possibilities – ‘But you can’t be sure you’re not a BIV’ – they come to realise that in fact they do not know that they have hands. So they concede that what they said before is false, and they retract their claim. Retraction, it may seem, is strong data *in favour of* sceptical invariantism.

However, sophisticated sceptical invariantism seems to lose this straightforward explanation of retraction. Davis (2007, 411) tells us,

Using ‘*p*’ loosely involves saying that *p* while meaning only that it is close enough to being the case that *p*. In loose usage, the speaker does not intend to commit himself to what he strictly speaking said. Hence he cannot be criticised for being sloppy or ignorant.

If speakers using knowledge sentences loosely in Low contexts never intend to commit themselves to what they literally say – that they *know* that they have hands – but only to what they mean, and if they cannot be criticised for being sloppy or ignorant, then they should not give in so easily and concede that they were *wrong*. After all, they cannot be criticised for their loose use. What they committed themselves to in the context of utterance was only the weaker claim that they are close enough to knowing for the purposes of their conversation at that stage. And that is still true. Thus, once under pressure, speakers should be at pains to point out that what they meant is true, even if strictly speaking, what they said is false. But if speakers do in fact use ‘know’ loosely, it must be that they do not have recourse to such defensive replies because they are unable to distinguish between what they said and what they meant. In other words, they are blind to what they meant.

The same point can be made against the hyperbole variant. Surely, a speaker will not concede error and retract her claim ‘That pool is a mile long’ after she has been corrected. Rather, she will admit that strictly speaking, what she said is false, but she will hasten to add that she was exaggerating and merely meant that that pool is very long for a community pool. Yet in knowledge cases, speakers do not say that they were exaggerating. They retract.

Inter-contextual Truth Ascriptions

Consider once more a speaker in HIGH who has just been convinced by sceptical considerations that nobody knows that they have hands. The speaker will judge an assertion of ‘I know that I have hands,’ made by a speaker in an ordinary Low context, as false. Sceptical invariantism predicts this judgment, if they assume that speakers’ judgments

focus on what is said (which is false, in any context), not on what is meant (which is true, since the utterance was made in LOW). However, reference to speaker error is needed to explain falsity judgments, made in LOW, of knowledge denials uttered in HIGH. Since an utterance of ‘I don’t know that I have hands’ in HIGH makes strict, or non-hyperbolic, use of the negated knowledge sentence, what is said is true, and no implicature is generated. Yet speakers in LOW may come to judge the knowledge denial in HIGH false, because they think the utterer does know that she has hands. If sceptical invariantists want to exploit their pragmatic story to explain these judgments, they must claim that speakers wrongly take the knowledge denial in HIGH to have been used loosely, or hyperbolically, in the way it would have been in the LOW context. Thus, it would pragmatically implicate that the speaker in HIGH is not close enough to knowing for the purposes *indicated in the LOW context*. The latter implicature is false, hence their judgment. This explanation, however, requires that speakers be ignorant of the fact that the negated knowledge sentence in HIGH is used strictly. This ignorance can in part be attributed to context confusion. Because speakers confuse their own context with the context of utterance, they take the knowledge denial to be used loosely, or hyperbolically. It also requires that speakers are blind to what is meant. They hear a false implicated content being communicated where there is none.

Disagreement

Sceptical invariantism seems to offer a straightforward explanation of disagreement judgments. As long as two assertions of ‘*S* knows that *p*’ and ‘*S* does not know that *p*’ are about the same subject, time, and proposition entertained by the subject at the time, they literally express incompatible contents. But this account is in tension with the sophisticated sceptical invariantist’s claim that in LOW, we are using ‘know’ loosely, or hyperbolically, and our truth judgments track the pragmatically conveyed content. For this would require that a speaker in LOW makes truth judgments about what knowledge attributions in her context pragmatically convey, but makes disagreement judgments based on the (in)compatibility of the literal content of knowledge claims. But that is implausible. A speaker who intends to engage in hyperbole by saying ‘That pool is a mile long’ would not take herself to disagree with someone who says that the pool is not literally one mile long. So if this is the sophisticated sceptical invariantist’s account of disagreement, she is committed to the thesis that speakers confound literal and pragmatic content.

Perhaps sophisticated sceptical invariantists can offer a better account of disagreement. They may hold that a speaker in HIGH, uttering ‘*S* does not know that *p*,’ takes herself to disagree with a speaker in LOW, who utters ‘*S* knows that *p*,’ because she takes both utterances to be strict (or non-hyperbolic) uses, with contradictory literal contents. A speaker in LOW takes herself to disagree with a speaker in HIGH because she

takes both utterances to be loose (or hyperbolic) uses, with contradictory pragmatically implied contents. However, this explanation requires that speakers are confused about each others' contexts. They take each others' utterances to be used in the way that it would be used in their own context. This ignorance of context might not seem utterly implausible. Whenever we pedantically correct someone who has used an expression loosely, we seem to be mistaken about standards of strictness in their conversational context (cf. the above challenges, 'That can't be! The brochure says it's a 50m, long-course competition pool,' 'That can't both be true! You cannot have taken two hours as well as two hours and 2 minutes.'). But there is an important difference between these cases and knowledge cases. The perceived disagreement in the pool and time measurement case is ephemeral. It is easily, and often, dissolved by the clarification that one of their utterances was a loose use, as in 'I meant that the pool is very long for a community pool,' or 'I meant that it took me *about* two hours.' No such linguistic moves seem available in the knowledge case. It would be utterly bizarre for the speaker in LOW to point out that she does not mean to contradict the knowledge denial by the speaker in HIGH, since she only meant that she is close enough to knowing for her purposes. Context confusion in the knowledge case is not ephemeral.

Inter-contextual Assessment of Truth Ascriptions to Knowledge Claims

According to sceptical invariantists, both John's and Bob's claims in (2.1) are false, so Mary is right in judging them false. What sceptical invariantists need to explain is why Bob's utterance seems correct in LOW. This is just a case of intra-contextual truth ascriptions, which I discussed above under the heading 'Basic Variability.' There I argued that sophisticated sceptical invariantists need to appeal to semantics/pragmatics confusion in their explanation.²⁴

Sceptical Paradox

On sceptical invariantist semantics, sceptical arguments like (SA) are valid and sound. They are not paradoxical. Adding the pragmatic accounts of Schaffer or Davis, sceptical invariantists can also explain why arguments like (SA) *seem* to present paradoxes. What they need to explain is how we come to believe that the conclusion is false. Suppose that

²⁴Sophisticated sceptical invariantists have a range of options to account for (2.2). Reject it, accept it but claim that 'relative to this context' is vacuous, or accept it but claim that Bob is making a true statement about the warranted assertability of John's knowledge attribution. The first response is straightforward and motivated by the same reasons as the relativist's dismissal above. The second response is just what sophisticated sceptical invariantism has to say about (2.1). On the third response, Bob is understood to judge what John implicated, namely that they are both close enough to knowing, for the purposes of their context, that Neil Armstrong was the first man to set foot on the moon. This claim is true. Mary, however, is ignorant of the distinction between truth conditions and warranted assertability conditions. Since she focuses on the literal content of John's assertion, she takes Bob to make a claim about the *truth* of John's assertion. That claim is false, hence her judgment.

premise one induces a sceptical context. This leads speakers to focus on the premises' literal contents, which are true. Once they come to the conclusion, however, they focus on pragmatically implicated content and switch back into a Low context, where 'I don't know that I have hands' is used loosely to mean 'I am not close enough to knowing for the indicated (lax) purposes.' And this is false. Yet they are unaware of their switch of context and focus and thus take (SA) to generate a paradox.

Unsurprisingly, this explanation involves context confusion. Speakers do not realise that they switch contexts in moving from premises to conclusion. They also do not realise that their focus switches from literal content to pragmatically implicated content. If they did, speakers would see that the literal contents of the premises and conclusion are true, and thus the argument sound, while they could clarify that they took the conclusion to mean that the speaker is not close enough to knowing for the purposes of the context. But speakers are puzzled by (SA). This clarification is not available to them. They confuse literal and pragmatic content. What is more, they are blind to the pragmatically implicated content. If the content speakers take the conclusion to pragmatically implicate was transparent to them, they would see that its negation is not inconsistent with the premises.

3.5 Interest-Relative Invariantism and Supervenience-Blindness

Interest-relative invariantists, or subject-sensitive invariantists (Fantl and McGrath, 2002, 2009; Hawthorne, 2004; Stanley, 2005; Weatherson, 2011)²⁵ hold that neither the truth value nor the content of knowledge sentences depends on any epistemologically significant features of the context of utterance or the context of assessment. They also do not supplement the invariantist semantics with an account of pragmatic implicatures. However, whether or not a subject *S* has knowledge of the proposition that *p* at a given time depends on more than just the traditional factors – whether *S* believes that *p*, whether the proposition that *p* is true, whether *S* has good evidence, whether *S* is using a reliable method, and so forth. The truth value of instances of 'S knows [does not know] that *p*' also depends on the subject's practical situation, including such factors as *S*'s attention to counterpossibilities, interests, and stakes. Metaphysically speaking, the attention to counterpossibilities, interests, and practical stakes of the subject at a given time are among the things on which the knowledge relation supervenes.²⁶ Thus, interest-relative invariantists acknowledge that the very features to which contextualists

²⁵Not all of these authors wholeheartedly embrace subject-sensitive invariantism. Hawthorne, for instance, seems to only give it his tentative endorsement, and Fantl and McGrath remain explicitly open towards a contextualist implementation of their view.

²⁶Different versions of interest-relative invariantism do not agree on all of the details here. Where the details matter, I will focus on Stanley's and Hawthorne's versions of interest-relative invariantism.

have drawn attention matter for knowledge, but they urge that it is the features of the subject's situation, not the attributor's, that matter.

Basic Variability & Inter-contextual Truth Ascriptions

Interest-relative invariantism (IRI) can account for the variability in our knowledge attributions in cases which differ with regard to the knowing subject's practical situation. But it is quite obvious that IRI's invariantist semantics does not predict that our intra-contextual truth ascriptions to knowledge sentences vary from context of utterance to context of utterance, while the subject's situation does not change. Nor does IRI make all and only the right predictions about inter-contextual truth ascriptions. As long as a subject *S*'s situation at a time *t* remains unchanged, including *S*'s attention to counterpossibilities, *S*'s interests and stakes, all utterances of knowledge sentences concerning the subject *S* at *t* and the proposition that *p* will have the same content and truth value. Stakes or standards in the context of utterance are irrelevant. Suppose that *S*'s stakes at *t* are low and *S* is in a sufficiently strong epistemic position to count as knowing. Then in all contexts of utterance, '*S* knows (at *t*) that *p*' is true, and its negation is false. Speakers in a context *Low* will be correct in attributing knowledge to *S*, and they are correct in ascribing truth to '*S* knows (at *t*) that *p*,' no matter if it is uttered in their own or some other context. But speakers in a context *HIGH* will, *pace* IRI, deny that *S* knows (at *t*) that *p*, and they will judge '*S* knows (at *t*) that *p*' to be false, both as uttered in their own context and as uttered in other contexts. So speakers whose stakes are significantly different from those of the subject are wrong in their knowledge claims and truth ascriptions to knowledge claims. (I shall use 'stakes' to include all the factors of attention to counterpossibilities, interests, and practical stakes, and will use 'Low'/'HIGH' to contexts with low/high stakes in this wide sense.) If no pragmatic account is available, an error theory is needed.²⁷ Attributors in *HIGH* do not take into account the stakes in the subject's situation when these stakes are low. That is, speakers prove to be ignorant of the fact that knowledge supervenes on the subject's stakes. Call this kind of error *supervenience-blindness*.

Hawthorne (2004, 162-6) provides an ingenious explanation of our tendency to deny knowledge when we are in *HIGH*, even if the knowing subject is in a low stakes situation and, according to IRI, does have knowledge. Psychological research on heur-

²⁷Hawthorne (2004, 160) provides a pragmatic explanation of the fact that attributors in *HIGH* do not and cannot truly utter '*S* knows that *p*' even when *S*'s stakes are low and *S* in fact knows that *p*. He draws attention to the knowledge norm of assertion, according to which one should assert that *p* only if one knows that *p*. So if an attributor in *HIGH* felicitously asserts '*S* knows that *p*,' she has to know that *S* knows that *p*. Because knowledge is factive, she will then have to know that *p*. But since the attributor is in *HIGH*, she does not know that *p*, and is aware of that fact, so she cannot, and does not, felicitously assert that *p*. But as MacFarlane (2005a, 208-9) and Stanley (2005, 99) point out, this only explains why attributors in *HIGH* cannot felicitously assert '*S* knows that *p*.' It cannot explain why attributors in *HIGH* can felicitously assert '*S* does not know that *p*.'

istics and biases shows that in our probability and frequency judgments we use a small number of heuristics that are cognitively effective, often lead to correct judgments but equally lead to biases that skew our judgments.²⁸ One such heuristic is the psychological availability of relevant events. For instance, ‘recently experienced floods appear to set an upward bound to the size of loss with which managers perceive they ought to be concerned.’²⁹ For attributors in HIGH certain not- p counterpossibilities are salient, or ‘available,’ with their raised stakes. We then tend to pessimistically overestimate the knowledge-destroying danger of these counterpossibilities and as a result come to deny knowledge. Moreover, Hawthorne claims, we tend to project our overestimations of these counterpossibilities onto the subject’s position and come to also deny knowledge of the subject, even when the subject’s stakes are low. Call this the *Psychological Pessimism and Projection* explanation of error (PPP).³⁰

Whatever its credits as a non-*ad hoc* explanation drawing on psychological features for which there is independent evidence, PPP does not give a plausible explanation of the supervenience-blindness IRI is committed to attributing to speakers. For one thing, when subjects with a basic understanding of probability theory are shown that and how their judgment is mistaken, they tend to accept the correction.³¹ Similarly, if attributors in HIGH tend to pessimistically overproject their own stakes, yet were otherwise sensitive to subjects’ stakes in their knowledge claims, they should find it acceptable to be corrected by explicit mentioning of the subject’s stakes. Suppose we are in HIGH and have just judged that S (who is in Low) does not know that p . Then we are told: ‘Look, there isn’t much at stake for S , so of course S knows that p .’ We would presumably find this correction utterly bizarre and would stand our ground. When made explicit, we do not think the subject’s stakes matter to whether or not she knows. This disanalogy between knowledge attribution cases and probability judgments shows that PPP, appealing to the distorting heuristics agents use in probability judgments, does not provide a plausible account of speakers’ blindness to the stakes-features in the supervenience base.

For another thing, Hawthorne’s PPP only explains why attributors in HIGH judge that subjects in low stakes situations do not know. However, we can also imagine attributors in Low who judge that a subject in a high stakes situation knows. For

²⁸See Kahneman and Tversky (1982) for an introduction and the articles in Kahneman et al. (1982) and Kahneman and Tversky (2000) for further detail.

²⁹Slovic et al. (1982, 465). This passage is cited by Hawthorne (2004, 164).

³⁰Nagel (2008, 2010) develops alternative psychological explanations of the error non-sceptical invariantists must ascribe. I learned of her work too late to discuss it in this dissertation.

³¹They may, under pressure, still make the same mistakes in similar cases. But this point is independent of the fact that they accept their judgments to be corrected by proper probabilistic calculation. It is an interesting fact, noted by Kahneman and Tversky (1982, 18) that ‘the reliance on heuristics and the prevalence of biases are not restricted to laymen. Experienced researchers are also prone to the same biases – when they think intuitively.’

instance, we may discuss in the pub whether John, currently a witness in a criminal trial, knows that he saw his friend Bob, the main suspect, near the crime site around the estimated time the crime was committed. We know that John said, under oath, that he knows that he saw Bob. But when asked by the judge whether he could rule out that he saw Bob's younger brother Hob that night, John withdraws his claim and says 'I don't know that it was Bob I saw that night.' In the pub, it is likely that we are unimpressed by the judges' overly demanding standards, and, relying on former conversations with John about what he saw, may insist that of course John knows that it was Bob whom he saw. So even while there is a lot at stake for John – he is under oath, and he cares about what happens to his friend – we in Low attribute knowledge. But PPP cannot explain our projection of low stakes onto a high stakes situation because it involves that we are overly impressed with 'available' error-possibilities.³²

Before we discuss IRI's handling of other data, it is worth noting that Stanley (2005, 101-4) has offered an alternative explanation of the data from inter-contextual truth ascriptions. According to Stanley, attributors in HIGH have a practical interest of their own in finding out whether the information state of the subject *S* in Low is sufficient for *them* to know that *p*. So they are in fact trying to find out if the counterfactual 'If *S* were in our practical situation, *S* would know that *p*' is true. Their falsity ascription to '*S* knows that *p*' reflects their judgment of this counterfactual, in whose truth or falsity they are interested given the practical stakes of their own situation.

Stanley (2005, 116) claims that his error theory for IRI attributes metaphysical error, not semantic error, and is therefore better than contextualism's semantic blindness theory:

[T]he contextualist explanation in terms of general semantic blindness is considerably *more disruptive to our conceptual scheme* than the explanation suggested by IRI. Semantic Blindness involves the dramatic claim that we are blind to the semantic workings of our language, and we make errors about the truth value of knowledge ascriptions as a result of that blindness. In contrast, the explanation I have provided of High Attributor-Low Subject Stakes is not similarly disruptive; I have not attributed to people a *new kind of hitherto unforeseen error*. People ignore some of the metaphysical determinants of knowledge, rather than being ignorant of

³²Another worry, pointed out by Cohen (2004b, 489), MacFarlane (2005a, 214), Stanley (2005, 101), and Schaffer (2006, 92-3) is that Hawthorne's Projection and Pessimism explanation overgenerates and undermines the data that is supposed to support IRI. For if we are in the grip of pessimism and often overestimate counterpossibilities, then we might be *overestimating* the stakes of a subject and deny knowledge to the subject even in high stakes cases that motivate IRI in the first place. A principled account is needed that explains why we are overly pessimistic and project in high attributor-low subject stakes cases but not in high subject stakes cases. See also section 3.7 below. For further criticisms of PPP, see Cohen (2004b, 489-90), DeRose (2005) and Schaffer (2006, 92-3).

features of their language. [...] This [Stanley's explanation of our error in High Attributor-Low Stakes cases] is certainly an error theory, but it does not involve imputation to agents of a *hitherto unknown* form of linguistic ignorance. (Emphasis D.K.)

If Stanley is right and his point can be generalised to the claim that metaphysical error attributions are less costly than semantic error attributions, this would give IRI a significant edge over all views that have to attribute semantic error. But unfortunately, Stanley gives no reasons why in general, semantic blindness attributions are worse – ‘more disruptive to our conceptual scheme’ – than metaphysical blindness attributions. Moreover, contrary to Stanley's contention, his explanation does involve the imputation of a form of *linguistic* blindness. When speakers in HIGH evaluate a knowledge attribution ‘S knows that *p*’ to a subject *S* in a situation with low practical stakes, they systematically do not evaluate the proposition expressed by ‘S knows that *p*’ but instead evaluate the counterfactual proposition *If S were in our practical situation, S would know that p*. Yet they seem to be unaware of this. If they were aware of the fact that ‘S knows that *p*’ does not say, nor implicate, that if *S* were in the attributors' practical situation, *S* would know that *p*, speakers would not make the judgments they in fact make. So Stanley needs to attribute a form of linguistic (semantic or pragmatic) blindness, too (cf. Schaffer (2006, 93)). In consequence, it seems false that Stanley's explanation does not impute to speakers a ‘hitherto unknown’ form of linguistic ignorance. At the very least, the burden is on Stanley to show that there are other cases in which speakers systematically treat sentences in indicative mood as if they expressed, or implicated, counterfactual contents.

There are further problems with Stanley's explanation. One is that it overgenerates. If we are likely to confuse the counterfactual with the proposition in fact expressed, we are likely to do so in other cases, too. Stanley's account also explains why attributors in HIGH deny knowledge of subjects in HIGH. But this undermines the original data in favour of IRI. And another problem with Stanley's explanation is that it extends to Low Attributor–High Subject cases, to the effect that attributors should be expected to judge in accordance with their own stakes. But they do not. Stanley does not give an account of why in these cases, it is the subject's high stakes that guides attributors' judgments (cf. Schaffer (2006, 93)).

Stanley provides an interesting alternative explanation of the troublesome data to Hawthorne's PPP. Given its numerous problems, however, it does not seem to offer any significant advantages over Hawthorne's explanation. I will therefore continue to discuss the prospects of IRI when supplemented with Hawthorne's explanation.

Retraction

Hawthorne's PPP also explains why speakers retract first-person knowledge ascrip-

tions for the form ‘I know that p ’ once they are confronted with undermining error-possibilities. Assume that the speaker’s stakes at t_1 are such that ‘I know that p ’ is true at t_1 . At t_2 , an error-possibility is raised, which changes the speaker’s stakes. So from t_2 onwards the speaker does not any more know that p . According to IRI, however, it is still true that the speaker knew that p at t_1 . So even if speakers are correct in saying, ‘Okay, so I don’t know that p ,’ they have no reason to retract their earlier claim and say ‘Okay, I was wrong. I did not know that p .’³³ PPP explains why people retract and deny that they knew at t_1 . They project their current situation, in which the stakes are raised, onto their earlier situation, and because they take their current situation to be more enlightened, they let their current judgment override their earlier judgment.³⁴

Again, however, PPP commits the interest-relative invariantist to the attribution of supervenience-blindness. From t_1 to t_2 , the supervenience base changes. If speakers were sensitive to the stakes feature in the supervenience base, and merely as a matter of psychological automatism projecting their current stakes at t_2 back onto t_1 , then they would be receptive to correction. But it would be awkward for them to respond to correction by saying ‘Yes, before you mentioned these error-possibilities, I did know, but now I don’t.’ To explain the awkwardness of explicitly mentioning the effect stakes have on our knowledge, IRI requires a supervenience-blindness thesis.³⁵

Ignorant Attributor

Another case whose explanation requires the attribution of supervenience-blindness is the situation in which the attributor is completely ignorant of the subject’s stakes. Here, IRI predicts that the attributor should have little confidence in her knowledge claims – after all, she does not know what the subject’s stakes are. (Similarly, an attributor would not confidently attribute or deny knowledge when she does not know what the subject’s evidence/justification/strength of epistemic position is.) But suppose I am asked whether my office mate John, who is present, knows that his car is parked in the university parking lot. Notice how bizarre it would be for me to answer ‘I can’t say. Hey

³³Williamson (2005, 220-1) notes that it is equally natural to say ‘So I was wrong. I don’t know that p after all,’ and adds that ‘after all’ has the subtle effect implying an admission of error. ‘A bomb disposal expert might report ‘The bomb is armed’ and then, after disarming it, ‘The bomb is not armed,’ but to say ‘The bomb is not armed after all’ in the latter circumstances would be to invite the interpretation that one was admitting an error in the original claim.’

³⁴IRI has a straightforward explanation of retraction in first-person cases where the stakes were in fact high from the beginning, and the knowledge attribution was false even when it was made (perhaps because the subject failed to see just how much was at stake for her, or she unjustifiedly ignored counterpossibilities). In this case, the subject has a genuine insight into what is, and has been from the beginning, at stake and so correctly withdraws her earlier knowledge attributions as false.

³⁵Hawthorne (2004, 166) and Stanley (2005, 106-7) point out that modal embeddings just like temporal embeddings of knowledge sentences present a problem to IRI. We do not say things like ‘If I had been less anxious back then, more committed to p , less worried about alternatives, I would have known p back then,’ or ‘John knows that he won’t be able to afford health insurance, but if he were discussing the possibility that he might win the lottery, he would not know this.’

John, how much does it matter for you right now if your car is parked in the university parking lot?’ or ‘I can’t say. John, which counterpossibilities are you considering right now?’ PPP might explain our confidence: We project our own stakes onto the subject’s situation. But again, our projection should not make us frown when the relevance of the subject’s stakes is pointed out to us. To explain the oddity of the reply ‘That depends on what’s at stake for John,’ IRI needs to appeal to supervenience-blindness.

Disagreement

IRI has a simple invariantist semantics, which is untainted by any variable pragmatic story. So they have an easy explanation of why we take two speakers in different contexts to disagree when one says ‘S knows that *p*,’ and the other says ‘S does not know that *p*’: The proposition that one speaker expresses is in conflict with the proposition the other speaker expresses.

Inter-contextual Assessment of Truth Ascriptions to Knowledge Claims

IRI predicts that Mary is wrong in HIGH in judging both John’s knowledge attribution in (2.1) and Bob’s claim ‘That’s true’ in their low stakes context as false. Interest-relative invariantists can explain Mary’s judgments by appeal to their PPP explanation of High Attributor-Low Subject cases. Mary projects her own stakes onto John and Bob’s situation and thus judges that they do not have knowledge. Accordingly, she judges that Bob’s truth ascription is false. Once more, however, PPP does not fully explain the supervenience-blindness, as Mary would not be receptive to correction in the way she would be with her judgments of probability and frequency.³⁶

Sceptical Paradox

Interest-relative invariantist do not, to my knowledge, explicitly address (SA). But it is not difficult to see what they can and should say about (SA). In considering IRI’s options for providing a fully satisfactory solution to sceptical paradox, we must distinguish between first-person versions of the sceptical argument (SA) and third-person versions (see section 2.5, or section 3.2 above). Consider first the first-person version, (SA), where subject and attributor are identical (we as readers are invited to let ‘I’ refer to ourselves). To explain (SA), interest-relative invariantists can mimic the contextualist-

³⁶Regarding dialogue (2.2), interest-relative invariantists have two options. They can reject (2.2) as improper data, or they can accept it. If they accept it, they can explain it in the same way as they explained (2.1), namely by appeal to PPP and supervenience-blindness. For IRI, ‘in this context’ will not do any significant work here. It might, on the one hand, refer to John and Bob’s context of utterance; then it is vacuous since knowledge attributions, and truth ascriptions to knowledge attributions, are not context-sensitive on IRI. ‘In this context’ might, on the other hand, refer to John and Bob’s practical situation; then it states a feature on which John and Bob’s knowledge supervenes, just like ‘That’s true given our evidence’ does. In either case, Bob’s claim in (2.2) is true and, as the latter interpretation makes obvious, an explanation of Mary’s falsity judgment will need to attribute supervenience-blindness to Mary.

relativist strategy and say that the first premise raises the stakes. Hence the speaker (reader) does not know that she is not a BIV. The stakes are now raised for the speaker's (reader's) situation that encompasses the entire reasoning in (SA), which guarantees that the argument is valid and sound. (Remember that by 'stakes', I am referring to factors that include attention to counterpossibilities such as the BIV hypothesis.) However, as speakers (readers) we find (SA) paradoxical because in moving to the conclusion, we somehow take the stakes to have dropped. Accordingly, we judge the conclusion false. Yet we are unaware of our changed, mistaken conception of the stakes. That is, we are supervenience-blind. We get puzzled because we are ignorant of the fact the supervenience base for knowledge includes stakes-features.

The trouble with this explanation is that it seems *ac hoc*. How and why would we, mistakenly and unconsciously, take the stakes in our context to suddenly be lowered? While it is easy to raise the stakes, it is usually very hard to deflect the force of sceptical considerations and lower the stakes. Of course, this is an explanatory burden that IRI shares with all of the above accounts. But IRI already appeals to PPP, which cites a heuristic that can skew our judgments where risk-*raising* considerations are psychologically available. Adding to this a psychological explanation that declares our judgments skewed where stakes-*lowering* considerations are psychologically available seems to give IRI an explanatory tool that is too powerful and moreover lacks independent evidence.

Consider next a third-person version of the sceptical argument, where the subject, John, is absent from the attributor's (the reader's) speech context:

(SA³) John doesn't know that he's not a handless BIV.

If John doesn't know that he's not a BIV, then John doesn't know that he has hands.

John doesn't know that he has hands.

In (SA³), the supervenience base remains unchanged, assuming that John is not in the context in which we reason through (SA³) and the claims concern John's epistemic state at the same time. John is either in HIGH or in LOW. If he is in HIGH, then the argument is sound and IRI needs to explain why it seems paradoxical to us. If John is in LOW, then the first premise and the conclusion are false, and IRI needs to explain why we take the first premise to be true. In both cases, a projection story is needed.³⁷ We project our stakes onto John's situation. As we reason through (SA³), we judge the premises true (given that our supervenience base now contains high stakes).

³⁷Again, interest-relative invariantists need a different projection story than PPP here. PPP claims that we are overly impressed with sceptical considerations. While this explains why we judge the premises true, it makes implausible that we would unreflectively let our stakes change to low stakes as we come to the conclusion, or that we would be impressed with the judgments we frequently pass in low stakes situations.

Our stakes change to Low as we reach the conclusion, and we thus judge it false. IRI's solution of third-person sceptical paradoxes, where subject and attributor are in distinct practical situations, suffers from the same *ad hocness* charge as its solution to first-person versions.³⁸

3.6 The Views on the Scoreboard

It is time to summarise the findings of this chapter and put them side to side with the results of chapter 2 on the scoreboard (see Table 3.1 below). The use facts on the vertical axis include basic variability data, inter-contextual truth ascriptions, retraction, disagreement, inter-contextual assessment of truth ascriptions to knowledge claims, and sceptical paradox. Note that the latter category also includes epistemic closure puzzles, as I argue in section 2.7, although in the present chapter we focused on the sceptical argument itself. Let me emphasise again that this list of data points is not exhaustive. But it represents the core data theories of knowledge attributions are expected to explain.³⁹

Two observations from Table 3.1 bear emphasising. First, no view straightforwardly predicts all of the data without recourse to a semantic, pragmatic, or supervenience-blindness thesis. Most of the views even need some form of blindness for the majority of the data. Surprisingly, perhaps, attributions of speaker error are ubiquitous. Second, relativism is no exception, contrary to MacFarlane's advertisement. While relativists predict most data correctly, their account of sceptical paradoxes and epistemic closure puzzles commits them to the attribution of assessment-index-blindness. Hence, none of the views can afford the argumentative strategy that their view is to be preferred on the grounds that it altogether avoids the attribution of error.

³⁸ Another problematic set of data for IRI is the behaviour of 'know' in temporal and modal embeddings. For instance, in a given context speakers tend to apply the same epistemic standards in their judgments of present-tense knowledge attributions and past-tense knowledge attributions. It is somewhat bizarre to say 'I don't know that my car is parked outside, but a moment ago before I considered the possibility that car thieves are roaming the neighbourhood, I did know that my car is parked outside.' Similarly, it is odd to say 'I know that the bank is open on Saturdays, but if it was of great practical importance whether the bank is open on Saturdays or not, I wouldn't know it.' Yet IRI predicts that these assertions are fine, since in the subject's situation in the past (counterfactual situation) the stakes were low enough (high enough) for her (not) to know. See for instance Blome-Tillmann (2009), DeRose (2004), Hawthorne (2004, 166), MacFarlane (2005a, 201-2), and Stanley (2005, 106-7) for discussion.

³⁹ Note also that I have not discussed the indexical contextualist solution of sceptical paradox on which the relativist solution in chapter 2 is modelled. The reader may extrapolate from the objection to relativism that the indexical relativist solution commits the view to the attribution of content-blindness. See DeRose (1995) and Schiffer (1996).

Table 3.1: Error Attributions

	Sophisticated Moderate Invariantism	Sophisticated Sceptical Invariantism	Interest-relative Invariantism	Indexical Contextualism	Nonindexical Contextualism	Relativism
Basic Variability	Semantics/Pragmatics Confusion	Pragmatics/Semantics Confusion	Supervenience-Blindness	✓	✓	✓
Inter-contextual truth ascriptions	Sem./Prag. Confusion & Prag. Blindness	Context Confusion & Prag. Blindness	Supervenience-Blindness	Content-Blindness	Use-Index-Blindness	✓
Retraction	Sem./Prag. Confusion & Prag. Blindness	Pragmatic Blindness	Supervenience-Blindness	Content-Blindness	Use-Index-Blindness	✓
(Dis)Agreement	Sem./Prag. Confusion or Prag. Blindness	Context Confusion	✓	Content-Blindness	✓	✓
Inter-contextual assessment of truth ascr. to knowledge claims	Semantics/Pragmatics Confusion	Semantics/Pragmatics Confusion	Supervenience-Blindness	Content-Blindness	✓	✓
Sceptical Paradox	Context Confusion & Pragmatic Blindness	Context Confusion & Pragmatic Blindness	Supervenience-Blindness	Content-Blindness	Use-Index-Blindness	Assessment-Index-Blindness

3.7 On the Evaluation and Comparison of Blindness Attributions

The discussion in the present and previous chapter contained error-theoretic objections against indexical and nonindexical contextualism, relativism, and sophisticated versions of invariantism. Together these objections show that all views are stuck with the attribution of some error to account for some data. If we want to make a reasoned decision between the views, we need to distinguish the objections by their force. We need to get clear about how to evaluate and compare the attribution of different kinds of error, or blindness. In this section, I offer five criteria for the evaluation of attributions of different kinds of error.

Before we look at the criteria, it will be instructive to register the responses to error-theoretic objections that are available to all theories of knowledge attributions.

- I. Argue that error attributions are a theoretical cost that is favourable to the costs incurred by other views

2. Make error plausible by finding precedents
3. Make error plausible by providing an independently motivated psychological explanation of the error
4. Argue that every theory must attribute error

The first strategy is one that I will deliberately bracket here. Clearly, the attribution of error may be an acceptable cost in light of a view's other theoretical virtues, especially if the view outranks all competitors with regard to these virtues. The importance of empirical adequacy as one theoretical virtue among others is an issue to which theorists in the debate take different attitudes.⁴⁰ The goal of my discussion here is to provide grounds for the adjudication between views purely in terms of how they fare with respect to their empirical predictions about use.

The second strategy of finding precedents of the attributed error in the use of expressions that allegedly exhibit the same semantic or pragmatic features is has been popular with indexical contextualists.⁴¹ Precedence is one of the evaluative criteria below, so I postpone the discussion of this strategy.

Hawthorne's PPP is an instance of the third strategy. We will come back to psychological explanations below.

The fourth strategy is employed for instance by DeRose (2006, 333-5) in defence of indexical contextualism. DeRose argues that while the attribution of semantic blindness is undesirable, it is a commitment that invariantist theories incur as well. 'Either way, many speakers are 'bamboozled by their own words.'" (DeRose, 2006, 335) The fourth strategy is also the one I have been pursuing on the relativist's behalf in this chapter. But pointing out that every theory is subject to the same problem is not sufficient as an argument in favour of a view over the competition. The view's error attribution must also be shown to be less implausible than the competition's error theories. DeRose, for one, seems pessimistic about the prospects of arguing in favour of the attribution of one kind of blindness over another:

⁴⁰For instance, Stanley (2005, 97-8) rejects slavish adherence to empirical adequacy: 'The fact that IRI [Stanley's version of IRI] gives a charitable explanation for all intuitions except High Attributor-Low Subject Stakes is not a *prima facie* concern for the interest-relative invariantist. Recall that these intuitions are not intended simply to be data for an epistemological theory, as the grammaticality of various sentences may be taken to be data for a syntactic theory. Rather, the role of my appeal to our intuitions about these particular cases is to make vivid our commitment to the conceptual connections between knowledge and practical reasoning.'

Among the other theoretical virtues, theoretical simplicity is surely crucial. One metatheoretical principle along these lines can be found in Grice, dubbed the Modified Occam's Razor: 'Senses are not to be multiplied beyond necessity.' (Grice, 1989, 47) In a more general fashion, Schaffer (2004b, 146-7) cites as a principle of 'standard linguistic methodology': 'Linguistic machinery should not be complicated without necessity.'

For specific desiderata for theories of knowledge attributions beyond the explanation of our usage, see for instance Hawthorne (2004, 111-2).

⁴¹See for instance Blome-Tillmann (2008), Cohen (1999b, 2004a), and Schaffer and Szabó (forthcoming).

There may be some reason for thinking it's more problematic to suppose that many speakers are blind to the context-sensitivity of their own words than to suppose that many are blind to the context-*insensitivity* of their own words. But it's not easy to see how to give any credible argument for such an asymmetry, and it's perhaps best not to stretch to anticipate how such an argument might go. (DeRose, 2006, 335)

I am less pessimistic. Error attributions, I would like to suggest, can be evaluated with respect to the following five criteria.

CRITERIA FOR THE EVALUATION AND COMPARISON OF KINDS OF BLINDNESS

1. PRECEDENCE: Can the kind of error be found in other other areas of thought and talk?
2. QUALITY OF COGNITIVE MALFUNCTIONING: How cognitively serious is the error?
3. PERVASION OF ERROR: How much data is explained only by appeal to error?
4. SCOPE: Does the error theory explain those and only those kinds of cases for which it is needed?
5. EXPLANATION: If a psychological/cognitive explanation of the kind of blindness is offered, how good an explanation is it?

Let me illustrate these criteria by way of examples.

PRECEDENCE. Schiffer (1996) argued against indexical contextualism that the attribution of what I call content-blindness to speakers is implausible because speakers are not prone to the same kind of error in their use of other indexical expressions. The diagnosed error is implausible because it lacks precedents of the right kind.⁴² Indeed, several authors have attempted to defend contextualism's attribution of content-blindness by arguing that the same error *is* found with expressions whose semantic workings are similar to 'know.' Thus Cohen (1999b, 2004a) and Blome-Tillmann (2008), underlining the similarity between 'know' and gradable adjectives, point out that similar use patterns occur with ascriptions of flatness or emptiness and that contextualist theories of gradable adjectives must attribute the same kind of error as contextualism about 'know.' Schaffer and Szabó (forthcoming) argue in defence of contextualism that while speakers' use of core indexicals and gradable adjectives fails to exhibit patterns of blindness, the similarity in use and blindness patterns between adverbial quantifiers, epistemic modals and 'know' supports their claim that 'know' shares the A-quantificational features of these expressions.

⁴²Similarly, Schaffer (2006, 93) objects to Stanley's semantic error explanation of High Attributor-Low Subject Cases (cf. section 3.5) that this error has no precedents. 'At least, some *independent evidence* is needed that we actually are prone to such confusions.' (Emphasis D.K.)

One has to be careful with PRECEDENCE as a criterion for evaluating the plausibility of a theory's error attributions. Whether or not the existence of precedent error makes its attribution more plausible depends on the independent evidence we have for the presence of the semantic or pragmatic mechanism in the type of expression after which 'know' is modelled. For instance, if it turns out that 'know' and adverbial quantifiers share certain patterns of use that contextualist theories do not correctly predict, then the error attributions the contextualist needs for the use of 'know' are only made more plausible if there is independent strong evidence in favour of contextualism about gradable adjectives or adverbial quantifiers. There might very well be. But if not: Two bad theories do not make one good theory.

QUALITY OF COGNITIVE MALFUNCTIONING. Indexical and nonindexical contextualists are each committed to attributing speaker error in cases where speakers retract earlier knowledge attributions in light of error-possibilities. Thus, these views predict that John is systematically mistaken in saying 'I guess I don't know that my car is parked in the driveway. But what I said before is true./But when I said before that I do know it, I spoke truly.' To explain retraction, nonindexical contextualists ascribe use-index-blindness, according to which speakers are blind to the fact that the truth of the contents of knowledge sentences depends on the epistemic standards salient in the context of utterance. Indexical contextualists ascribe content-blindness, according to which speakers are blind to what they said – the content of their knowledge attribution. Are both error attributions equally (im)plausible? It might be argued that content-blindness involves a graver cognitive error to speakers. Not only are they wrong about the truth or falsity of their earlier knowledge claim, they are also mistaken about what they were saying.

It should be noted that QUALITY OF COGNITIVE MALFUNCTIONING, like any other criteria, cannot be used in isolation to deliver a verdict. *Prima facie* the attribution of a more serious cognitive malfunctioning is less plausible than the attribution of a less serious cognitive malfunctioning. But this need not be the case if there is compelling independent evidence for the presence of the more serious malfunctioning.

PERVASION OF ERROR. Here we can simply take a look at the scoreboard. For how many data points is the error theory needed? We can see that, for instance, nonindexical contextualism attributes a less pervasive blindness than, say, sophisticated sceptical invariantism. What does this tell us about the plausibility of an error attribution? On the one hand, there is MacFarlane's 'general problem' that 'the more error we attribute to speakers, the less we can conclude from these facts [about speakers' use of 'know']'. (MacFarlane, 2005a, 215) Clearly, an all-pervasive error theory leaves little evidence in favour of a view. On the other hand, a simple *the less, the better* verdict might not always be right. It might lend plausibility to an error theory that similar cases are accounted for by appeal to the same kind of blindness. In the same vein, if a kind of blindness

is needed to account for only one or two data points, one might wonder why speakers do not seem susceptible to the same kind of error in very similar kinds of cases. Like the other criteria, pervasion of error cannot adjudicate between kinds of blindness in isolation from the other criteria.

SCOPE. Where a view provides a substantial error theory, the latter should explain speakers' use in all and only those cases where the view otherwise makes incorrect predictions. There are two ways in which the attribution of error can fail to have the right scope.

(i) Overgeneration: The attribution of a kind of blindness explains more than those kinds of cases for which it is designed. Schaffer (2006, 92-3), for instance, points out that Hawthorne's PPP for High Attributor-Low Subject cases equally provides explanations for the original data that motivated IRI (the 'High' cases in which subject and attributor are both in high-stakes contexts). This undermines the support for IRI: '[T]he sort of projection error Hawthorne has in mind would equally explain away any residual intuition we have to deny knowledge to the subject in the original *High*. Perhaps the high-stakes subject does know [*pace* IRI's prediction], but we intuit otherwise because we overestimate the possibility that he is at risk. This would thus undermine the original case for IRI.'⁴³

(ii) Undergeneration: The attribution of a kind of blindness does not explain all of the cases for which it is needed. For instance, it might explain the original cases it is designed for but does not cover slightly modified cases for which the view also makes incorrect predictions. Thus, Schaffer argues against Hawthorne's PPP that it cannot explain a High Attributor-Low Subject case in which it is part of the case's description that the subject, Sam, 'has little at stake.' Schaffer writes:

In the case as described it is quite clear that only you are at risk of bankruptcy. Sam is safe. We can test for this by asking *both* whether Sam knows *and* whether Sam is at risk. If you answer that Sam does not know *and* is not at risk, then your basis for denying him knowledge cannot be the sort of projection error Hawthorne posits. (Schaffer, 2006, 93)

EXPLANATION. This last criterion is admittedly vague. What makes for a good psychological explanation is a complicated question I cannot hope to answer here. General epistemic virtues such as empirical correctness, simplicity, testability, fertility, and conservativeness are certainly relevant. For our purposes, it may suffice to record that the plausibility of an error attribution that is based on a particular psychological explanation of the error will depend on the plausibility of that story as a psychological explanation.

⁴³Schaffer attributes this objection to Williamson (2005). See also Cohen (2004b, 489), MacFarlane (2005a, 213-4), and Stanley (2005, 101). Schaffer also presents a similar overgeneration objection to Stanley's semantic blindness explanation (Schaffer, 2006, 93).

Unfortunately, there is no simple algorithm that takes us from some views' performance with respect to these criteria to a clear ranking of the views. There are two obstacles. The first is that it is far from clear how to weigh the criteria against each other. To illustrate, suppose that a view does well on all criteria except SCOPE, because it overgenerates significantly. Is that view to be preferred to a view that does well on all criteria except PRECEDENCE, because the attributed error is hitherto unforeseen? Or is a high score on PRECEDENCE to be considered more important than a high score on SCOPE? Second, decision-making is hampered by the fact that whether or not a view fares well with respect to a criterion can sometimes not be answered in isolation from other criteria. For instance, we saw above that some theory T_1 that needs to posit error for only a few data points (PERVASION) may seem preferable to a theory T_2 that posits error for a significantly larger range of data points because T_1 simply leaves more empirical evidence that supports it. But if T_2 can be combined with a plausible psychological explanation (EXPLANATION) that supports all and only the data points it is needed for (SCOPE), whereas it would not explain why the error occurs only in the cases in which theory T_1 posits error, T_2 's higher pervasion of error does not make the theory less plausible. It will thus be a delicate question, and one on which theorists may disagree, how to interpret and weigh the criteria in any given case.

3.8 Relativism Rescued?

At last, let me use these criteria to sketch a strategy of defence for relativism. Before I do so, a caveat is in order. Even though the argument for relativism is one for the plausibility of its error attribution *relative to* the error attributions of the competing views, I cannot hope to give a fair evaluation of each competing view that makes relativism's superiority obvious. I can only give comparative hints and suggest how the case for relativism will go. I will go through the criteria in their opposite order.

On EXPLANATION. Relativists have not given a psychological explanation of index-blindness. I will not be able to present a full error *theory* on behalf of relativism here. This would require an explanation of *why* otherwise competent speakers are index-blind in sceptical paradox and epistemic closure cases. In the meantime, then, a verdict on EXPLANATION will have to wait.

On SCOPE. Relativists have recourse to index-blindness to explain all and only sceptical paradoxes and epistemic closure cases. Does their explanation under- or overgenerate? Again, this question cannot be answered before relativists provide a specific explanation of why index-blindness occurs in these cases. The prospects of such an explanation vis-à-vis SCOPE will depend on whether or not it explains the error by appeal to facts about our cognitive processes that can be found in cases like (SA) but not in other cases where relativism predicts no error.

On Pervasiveness of Error. Relativists need index-blindness *only* in their account of sceptical paradoxes and epistemic closure puzzles. In comparison, every competitor theory needs a blindness thesis for a variety of data. If less (pervasive error) is better, then relativists have by far the best record on Pervasiveness of Error. There is a lot of data that supports the relativist view, for which competitors need error theories.

As pointed out, there is some reason not to evaluate Pervasiveness of Error in isolation from other criteria such as Scope. Overall, an error theory that provides a convincing explanation of why errors would be committed in just the relatively large range of cases may be more plausible than an error theory that makes it a mystery why speakers would systematically fail to get it right only in particular kinds of cases. Since relativists have not provided a full error theory, the final judgment on Pervasiveness of Error must wait. In the meantime, relativists are at no disadvantage. Most views have not given full error theories either, and where they did (PPP), it significantly over- and undergenerated and so did not support the pervasiveness of error IRI is committed to attributing.

On Quality of Cognitive Malfunctioning. To impute index-blindness to speakers is to specify the kind of error involved, but it is not to explain *why* this error occurs, nor is it to give an account of the potential cognitive malfunctioning involved. As long as such an explanation is wanting, it is difficult to reach a verdict regarding Quality of Cognitive Malfunctioning.

In absence of a full error theory, however, one point can be made. Relativists need to explain why error occurs with sceptical paradoxes and epistemic closure puzzles. So it is likely that in the explanation, specific facts about whatever is going on in our reasoning of (SA) and similar cases are important. But all theories are united in their need to explain some error or other being performed by competent speakers in cases like (SA). So it is likely that all theories will refer to these facts in their fully worked-out error theories. Moreover, if it is correct that content-blindness is indicative of a greater cognitive shortcoming than index-blindness, the relativist has a *prima facie* plausible case that the cognitive shortcoming they have to posit is less serious than, for instance, the one indexical contextualist must posit.

On Precedence. In section 2.5, I argued that the attribution of index-blindness is implausible because there seem to be no precedents. Less controversial forms of index-sensitivity, time- and world-index-sensitivity, pose no similar challenge to speakers. Relativism's score on Precedence looks bad. It is not clear, however, whether its score is worse than other views' score on Precedence. Let us consider indexical contextualism and let us grant for the moment that Blome-Tillmann and Cohen (or Schaffer and Szabó) succeed in showing that a lot of the error patterns with 'know' are also manifest in our use of gradable adjectives (or A-quantificational expressions such as epistemic 'might' or adverbs like 'always'). Remember that the only error the relativist has to attribute is error arising from sceptical paradoxes and epistemic closure

puzzles. Do gradable adjectives (A-quantifiers) cause the same appearance of paradox with competent speakers as ‘know’ does? It seems doubtful. Consider (SA^F).

(SA^F) This lab desk isn’t flat. It’s got some bumps.

If this lab desk isn’t flat, then that meadow isn’t flat.

That meadow isn’t flat.

(SA^F) resembles the original (SA) in relevant respects. Premise one raises the standards for flatness, premise two is an indicative conditional with negated flatness ascriptions in antecedent and consequent, and the conclusion follows by modus ponens. Contextualists may want to argue that (SA^F) causes puzzlement just as (SA) does. That is, competent speakers who feel the pull of the argument may worry that accepting the conclusion is in conflict with a lot of their everyday flatness ascriptions and may as a result want to deny the conclusion while still accepting the premises. I actually doubt that (SA^F) causes this puzzlement in competent speakers. But even if it does, there is still a difference between the two arguments. Speakers have simple resources to resolve the conflict they may feel between accepting the conclusion in the context of (SA^F) and denying it in ordinary contexts. It is natural to clarify one’s acceptance of the conclusion in (SA^F) by modifying it: ‘Well, that meadow isn’t flat by the standards needed in the lab, but it’s perfectly flat by normal standards,’ or ‘That meadow isn’t flat for a lab desk, but it is flat for a meadow.’ In contrast, similar modifiers for ‘know’ are not naturally available to competent speakers. It takes thorough exposure to epistemology to not balk at an assertion of ‘Well, I don’t know that I have hands by strict standards, but I perfectly know it by ordinary standards.’ But if this is so, indexical contextualists still owe proper precedents for the systematic error arising from cases of sceptical paradox. These are the only the cases for which the relativist lacks precedents of error. It is hence not clear whether the relativists is worse off than the indexical contextualist.

In sum, relativism has *prima facie* advantages over the competition regarding PERVASION OF ERROR, seems to stand a good chance of scoring at least as high as, say, indexical contextualism on QUALITY OF COGNITIVE MALFUNCTIONING, and does no worse than indexical contextualism with respect to PRECEDENCE. A full defence of relativism with respect to the error attributions it needs to make will require an explanation of why speakers systematically manifest index-blindness in sceptical paradox cases. An explanation of this kind is likely to appeal to specific facts about sceptical paradox cases. Since all views need to explain some error or other in these cases, the need for such an explanation puts relativism at no disadvantage.

3.9 Conclusion

Chapters 2 and 3 together provide an argument that all theories of knowledge attributions – relativist, contextualist, or invariantist – make some incorrect empirical predictions. None of the views escaped attributing error to speakers in their account of the problematic data. I identified the kinds of error each view is committed to attributing to competent speakers. In particular, I argued that even sophisticated versions of invariantism that provide pragmatic explanations of the data are stuck with the attribution of what I called pragmatic blindness.

The fact that all views attribute error left us with the need to evaluate and compare the views' error theories and the kinds of error they involve. Despite the ubiquity and importance of error-theoretic objections in the debate on knowledge attributions, clarity on such evaluative criteria has been wanting. I suggested five criteria, which I labelled PRECEDENCE, QUALITY OF COGNITIVE MALFUNCTIONING, PERVASION OR ERROR, SCOPE, and EXPLANATION. While these criteria do not quite provide an algorithm for our decision process between views inflicted with error theories, they help us get clearer on the force of error-theoretic objections and provide us with firmer grounds for a reasoned decision. I concluded on a positive note for relativism, suggesting that pending a full error theory, relativism does not fare worse than competing views on many criteria.

Part II

***De Se* and Subjective Attitudes**

PERSPECTIVE IN CONTEXT: SEQUENCED WORLDS

4.1 Introduction

There is, or appears to be, a conflict between, on the one hand, the perspectival nature of many of our attitudes and, on the other hand, the received picture of linguistic communication. According to this picture, there is a single content which the speaker believes, expresses in speech, and which hearers come to believe if they understand and trust the speaker. But it is difficult to see how this picture fits with a natural account of two kinds of perspectival attitudes: so-called *de se* attitudes, i.e. attitudes about oneself, and attitudes about ‘subjective’ matters such as personal taste, like the belief that liquorice is tasty. According to this account, mental attitudes have so-called *centered content* – roughly, content whose truth depends on an individual at a world and time. I will argue that the conflict between the received picture of communication and the centered content view of perspectival attitudes can be resolved without giving up either of these attractive views. The solution I will propose is a unified theory of mental attitudes and linguistic communication on which content is modelled not in terms of centered worlds but as a set of *sequenced worlds* – roughly, possible worlds that are ‘centered’ on a sequence of individuals at a time.

It has long been argued that *de se* attitudes – attitudes about oneself – are a distinctive category of thought, irreducible to either *de dicto* or *de re* thought.¹ On Lewis’s elegant and influential proposal, the content of a thought one would express by using the words ‘I am hungry’ is a set of centered worlds, where a centered world is a possible world ‘centered’ on an individual at a time. To believe *that I am hungry* is to locate oneself in the set of centered worlds whose center is hungry. A number of philosophers and linguists have recently argued that the content of thoughts about ‘subjective’ matters such as personal taste must also be understood as centered content.² To believe that liquorice is tasty is to locate oneself in the set of centered worlds to whose center liquorice tastes good.

How do we communicate these self-locating beliefs? The standard picture of communication says, very roughly, that we exchange information by simply passing it on, from speaker’s head to hearer’s head. But this widely endorsed picture is in conflict

¹ See, for instance, Castañeda (1966, 1967), Perry (1977, 1979), and Lewis (1979a)

² See, for instance, Egan et al. (2005), Egan (2007, 2010a), Lasersohn (2005) and Stephenson (2007a)

with self-locating belief. If I believe that I am hungry, and I say to you, ‘I am hungry,’ you do *not* come to believe the content I believe. For that would be for you to believe that you are hungry. Instead you come to believe another content, namely that I am hungry.

The conflict at hand suggests that we reject either the natural standard picture of communication or the elegant Lewisian account of self-locating belief. In this chapter, I will argue that neither is necessary. I will begin by stating the self-locating account of belief and other attitudes (§4.2), the received picture of linguistic communication (§4.3), and the conflict between these two views (§4.4). Then I will develop the sequenced worlds framework in a broadly Stalnakerian picture of communication, and I will show that it affords an attractive reconciliation of the two views (§§4.5–4.7). I will give a semantics for predicates of personal taste and for personal pronouns that yields the desired sequenced worlds content in communication (§§4.8, 4.9). I will then turn to a brief discussion of competing solutions to the conflict (§4.10) and will close by answering objections to the sequenced worlds account (§4.11). In chapters 5 and 6, I will introduce relativist and contextualist versions of the model and the empirical data on which a decision between these versions turns, and I will explore some of the pragmatic and semantic details of the sequenced worlds model.

4.2 Centered Content

4.2.1 Attitudes *De Se*

De se attitudes are thoughts about oneself when one thinks about oneself in the first-person way. They are thoughts one would typically express with a sentence containing a 1st-personal pronoun (‘I’, ‘me’, ‘my’).³ David Lewis famously argued that the objects of *de se* attitudes are (or determine) not possible worlds propositions but sets of *centered worlds*. At least three kinds of motivation have been given for the distinctiveness of *de se* thoughts – that is, for the irreducibility of *de se* attitudes to either *de dicto* or *de re* attitudes. Let me quickly rehearse what I consider the strongest motivation: similarity arguments (to borrow Egan’s (2010b) term).⁴

³*De se* attitudes are often understood in a wider sense to include *de nunc* attitudes – thoughts about one’s location in time such as the thought *that the meeting starts now* – and thoughts about one’s location in space such as the thought *that this is the Mt. Tallac trail*. For simplicity, I will here focus on the narrow class of *de se* attitudes that are characteristically expressed by using 1st-personal pronouns.

⁴The second type of considerations in favour of a distinctive category of *de se* thought come from arguments that purport to show that while one may know all relevant standard (possible worlds) propositions, one may still lack knowledge about oneself (see for instance Lewis’s (1979a) two gods case and Perry’s (1979) case of Lingens in the Stanford Library). The third kind of motivation comes from the semantics of attitude verbs like ‘expect,’ ‘want,’ and ‘imagine’ that take infinitival complement clauses as in ‘Joe wants to meet Obama.’ Their truth conditions have been shown to be sensitive to the ascription of *de se* attitudes (see, e.g., Morgan (1970) and Chierchia (1989)).

Mad Heimson believes that he is Hume, a belief he would express by saying ‘I am Hume.’ Hume, of course, also believed of himself that he is Hume. Hume and Heimson share a belief, they are doxastically similar, which explains similarities in their actions (given that their desires and background beliefs are similar). They introduce themselves as ‘David Hume,’ get angry when they hear Hume being badmouthed, sign with ‘David Hume,’ and so on. But there is no relevant possible worlds proposition that both Heimson and Hume believe that would explain this doxastic similarity. Why not? Heimson and Hume are worldmates. So any candidate possible worlds proposition is either true at their world or false at their world. If it is true, then both Heimson and Hume have a true belief. If it is false, both Heimson and Hume have a false belief. But Hume is right in believing that he is Hume because he *is* Hume. Heimson is wrong in so believing. So the shared object of their beliefs, which explains their similarities in action, cannot be a possible worlds proposition.

Lewis concludes that the shared object of Heimson’s and Hume’s beliefs is the property *being Hume*, which each of them self-ascribes.⁵ Equivalently, the object of their beliefs is the centered content that is the set of centered worlds whose center is Hume. A *centered world* is a possible world ‘centered’ on an individual inhabiting the world at some time.⁶ Just like possible worlds can be understood as ways the world might be, centered worlds can be understood as ways one might be in the world, as possible locations in logical space, or as perspectives one might have. Formally, a *centered world* can be represented by an ordered triple $\langle w, t, x \rangle$ consisting of a world w , a time t , and an individual x inhabiting w at t .⁷ The triple determines the individual’s spatiotemporal location in w and every other fact concerning the individual in w at t , including her attitudes towards matters of taste. I will also sometimes speak of an individual x ’s location determined by the triple $\langle w, t, x \rangle$ as x ’s *perspective* (in w at t). Intuitively, a perspective is a subject’s viewpoint, which includes where she is, what the distance is to things around her, what she can see from her position, as well as what she thinks and feels about things.

On a standard possible worlds account, one’s overall belief state determines a set of possible worlds, the possible worlds compatible with what one believes. Analogously, on the centered worlds account, one’s overall belief state determines a set of centered worlds, the set of centered worlds compatible with what one believes. A centered world

⁵Perry (1977, 1979) draws the lesson that Hume and Heimson believed different propositions and that their doxastic similarity was to be accounted for by their shared ‘belief state’ – roughly, the first-personal mode of presentation of the propositions. For reasons to prefer Lewis’s analysis over Perry’s account, see Lewis (1979a, 524–5, 535–8).

⁶The term *centered world* was originally coined by Quine (1969).

⁷Alternatively, a centered world can be represented by an ordered quintuple $\langle w, x, y, z, t \rangle$ of a possible world w and the space-time coordinates x, y, z, t . For some metaphysical differences between the two conceptions of centered worlds underlying these ways of picking them out as well as for some problems of each conception, see Liao (2012).

$\langle w, t, x \rangle$ is compatible with what one believes iff one's beliefs do not rule out the possibility that one is x in w at t . One believes a centered worlds content p iff every centered world compatible with what one believes is contained in p . Hume believes that he is Hume iff every centered world compatible with what he believes is a member of

(4.1) *HUME*: $\{\langle w, t, x \rangle: x \text{ is Hume in } w \text{ at } t\}$.

His belief in *HUME* at some time t_1 is correct iff in addition, his actual location at t_1 , $\langle @, t_1, \text{Hume} \rangle$,⁸ is a member of *HUME*.

Lewis speaks of belief as self-ascription of properties. On centered worlds talk, belief is self-location in a set of centered worlds. Since properties correspond to sets of centered worlds, we will switch back and forth between these two ways of talking.⁹

On the centered worlds account, *all* attitudes have centered content. However, not all centered contents are *de se* contents. Following Egan's (2006, 107) terminology, we can call a centered content p *boring* if it does not distinguish between locations in a world. More precisely, p is boring iff for every world w and pairs $\langle t_1, x \rangle, \langle t_2, y \rangle$ such that individual x inhabits w at time t_1 and y inhabits w at t_2 , p contains $\langle w, t_1, x \rangle$ iff it contains $\langle w, t_2, y \rangle$. Because boring centered contents distinguish between worlds but not between locations in a world, they are equivalent to possible worlds contents.¹⁰ *De se* contents do distinguish between locations in a world. They are *interesting*. A centered content p is interesting iff there is a world w and pairs $\langle t_1, x \rangle, \langle t_2, y \rangle$ of individuals inhabiting w at times t_1 and t_2 , respectively, such that p contains $\langle w, t_1, x \rangle$ and does not contain $\langle w, t_2, y \rangle$.

A quick note on belief *de re*. For Lewis (1979a), belief *de re* is 'other-ascription' of properties and receives a descriptivist analysis in terms of self-ascription. Very roughly, to believe *de re* of Goethe that he is the author of *Faust* is (i) to be uniquely acquainted with Goethe in a particular way and (ii) to self-ascribe the property of being uniquely thusly acquainted with someone who is the author of *Faust*. For simplicity, I will try to avoid talk of attitudes *de re*. Where needed, I will depart from Lewis and adapt a standard non-Lewisian analysis on which the content of the belief that Goethe is the author of *Faust* is the boring centered worlds content $\{\langle w, t, x \rangle: \text{Goethe is the author of } Faust \text{ in } w \text{ at } t\}$.

⁸Here, '@' denotes the actual world.

⁹A property determines its extension at possible worlds and times. Given a world and time, it determines the set of individuals that instantiate the property at that world and time. A property can thus be thought of as a function from worlds and times to sets of individuals. This function is equivalent to a function from world-time-individual triples to truth values. The set characterised by the latter function is a set of world-time-individual triples $\langle w, t, x \rangle$ such that x instantiates the property at t in w (cf. Lewis (1979a, 532)).

¹⁰See Nolan (2006), however, for objections to the claim that one can do with centered worlds whatever one can do with possible worlds.

4.2.2 Attitudes Towards ‘Subjective’ Matters

Egan et al. (2005), Egan (2007, 2010a), Lasersohn (2005) and Stephenson (2007a) have recently argued that attitudes towards ‘subjective’ matters like personal taste are best understood as having interesting centered contents, too. Intuitively, the truth of claims about what is tasty, fun, or entertaining depends not just on what the objects concerned are like, but on some subject not made explicit. The degree to which we differ in our beliefs about matters of taste seems to support this subjective dimension of taste judgments and stands in stark contrast to the convergence in the judgments about objective matters which we pass with comparable confidence. On the centered worlds account, to believe that some cookie is tasty is to locate oneself in the set of centered worlds to whose center the cookie tastes good:

(4.2) *COOKIE*: $\{\langle w, t, x \rangle$: the (contextually salient) cookie tastes good to x in w at $t\}$.

COOKIE is the set of centered worlds to whose center the cookie in question tastes good; it is the set of perspectives on which the cookie in question tastes good. *COOKIE* can be true of one person but false of another. If the cookie happens to taste good to Ben, he is right in believing *COOKIE*, while Anna is wrong in believing *COOKIE*, if the cookie does not taste good to her. *COOKIE* is an interesting centered content.

In self-ascription talk, to believe *COOKIE* is to self-ascribe the property of being an x such that the cookie tastes good to x . Just what the property is that I am self-ascribing when I believe of something that it is tasty is a substantial question. The options Egan (2010a) discusses include *being disposed to enjoy this cookie*, *having gustatory capacities and sentiments that are robustly disposed to interact with this cookie in a way that produces gustatory experiences of an intrinsically desirable sort*, *being someone whose suitably idealised self would be disposed to enjoy this cookie*, and *being a member of a group typical members of which are disposed to enjoy this cookie*. I will here remain neutral on the metaphysics of tastiness – on the question of what experienter and experienced object have to be like for the object to be tasty to the experienter.

Egan et al., Egan, Lasersohn and Stephenson have coupled their pledge for centered worlds content with a relativistic semantics for predicates of personal taste such as ‘tasty’ and ‘fun.’ I will turn to semantic questions in section 4.8 below.

4.2.3 Egocentricity and Other Similarities

There are important differences between *de se* beliefs and beliefs about matters of personal taste. For the moment, however, notice their similarities. Crucially, both kinds of belief are beliefs in *interesting* centered contents and follow an egocentric belief norm:

EGOCENTRICITY Believe p only if *you yourself* are correctly located by p .

Importantly, EGOCENTRICITY requires only the believer's correct location to be contained in p . An agent's *de se* belief that she is hungry is correct as long as that agent is hungry, even if she were to be the only person in logical space ever to be hungry. Similarly, an agent appropriately believes that some cookie is tasty as long as she herself is such that the cookie tastes good to her, even if she were the only person ever to enjoy its taste.

EGOCENTRICITY also applies to belief in boring centered contents, even though intuitively, the appropriateness of one's belief that, say, snow is white depends on more than just one's own location. However, an agent is correctly located by a boring centered content if and only if she and all of her worldmates are correctly located by it. Thus it is only the appropriateness of believing interesting centered contents that depends *merely* on the location of the attitude holder.

The first-person orientation of beliefs about taste explains the general availability of the inference from a belief *that X is tasty* to a belief *that X is tasty to me*. The inference is available because on the view, my belief that some food is tasty entails my belief that the food is tasty to me. In fact, on the centered worlds framework, the beliefs have the same content, viz. $\{\langle w, t, x \rangle: \text{the food is tasty to } x \text{ in } w \text{ at } t\}$.¹¹

In addition to the interestingness of content and the egocentricity of belief, *de se* attitudes and attitudes about taste also have in common an account of believing alike in terms of shared content. When Ben and Anna each believe, *I am hungry*, their similar disposition to act is explained by the shared interesting centered content of their beliefs. Likewise, when they each believe, *This wine is tasty*, they are disposed to similar behaviour (given similar background beliefs and desires); they are reaching for their glass frequently, will not refuse a refill, etc. Their similar disposition can be explained by their believing alike, which is accounted for by the shared content: $\{\langle w, t, x \rangle: \text{the (contextually salient) wine is tasty to } x \text{ in } w \text{ at } t\}$.

The most important difference between the two kinds of interesting centered contents concerns their role in communication. We will turn to this difference in section 4.4. At the level of mental attitudes, one difference is worth pointing out. There is a second notion of believing alike, captured in agreement reports regarding people's beliefs, which sets *de se* attitudes apart from taste attitudes. When Ben and Anna both believe, *This wine is tasty*, we can report that they agree that the wine is tasty. But when they both believe, *I am hungry*, there is nothing they agree about, despite the

¹¹On the sequenced worlds analysis of belief I propose in section 4.6, my belief *that X is tasty* entails the belief *that X is tasty to me* but the two beliefs need not have the same content. The analysis distinguishes between belief in a solitary context and belief in a conversational context. In a solitary context, my beliefs *that X is tasty* and *that X is tasty to me* have the same content. But in a conversational context, the beliefs will not have identical contents (on one natural way of understanding identity of contents in the framework). However, even in a conversational context, the content of my belief *that X is tasty* entails the content of my belief *that X is tasty to me*, thereby explaining the availability of the inference from the former to the latter.

shared content of their beliefs. We will need an explanation of why these two interesting centered contents come apart with respect to our intuitions of believing alike qua agreement. This issue will be addressed in section 6.4 below.

4.3 The Lockean Picture of Communication

Received wisdom paints a simple and attractive picture of linguistic communication as the transfer of information. This picture is famously expressed by John Locke:

They suppose their words to be marks of the ideas in the minds also of other men, with whom they communicate: for else they should talk in vain, and could not be understood, if the sounds they applied to one idea were such as by the hearer were applied to another, which is to speak two languages. But in this men [...] think it enough that they use the word, as they imagine, in the common acceptation of that language; in which they suppose that the idea they make it a sign of is precisely the same to which the understanding men of that country apply that name.

(An Essay Concerning Human Understanding, book III, ch. 2, §4)

The picture attributed to Locke and arguably endorsed by Frege and much of 20th century philosophy is this: a speaker succeeds in communicating when she has an idea in her head and uses the words that express this idea in language and arouse in the hearer the very same idea. The speaker's mental content is, as it were, transported from her head to the hearer's head, who comes to share this content.¹²

On the Lockean picture, one kind of content plays the following three roles:

1. **SPEAKER'S MENTAL CONTENT:** what the speaker believes and intends to communicate
2. **SPEECH ACT CONTENT:** what the speaker's (assertoric) speech act literally expresses
3. **HEARER'S MENTAL CONTENT:** what the hearer comes to believe, if she understands and trusts the speaker

¹²For the attribution of the 'Lockean' picture of communication to Locke, as well as to Frege, and for a defence of the picture against criticism, see Pagin (2008). Interestingly, Frege has been interpreted by some to be committed to the view that contents ('thoughts') expressed by sentences involving indexicals like the first-person pronoun 'I' are private, non-shareable and incommunicable, contrary to the shareable nature of thoughts that Frege stresses in his writings (see for instance Kripke (2011) for this interpretation). First-personal thoughts being of a particular concern in this chapter, I will have to leave it open whether Frege can correctly be claimed to endorse the Lockean picture of communication.

So if I believe that snow is white and intend to tell you, and I assert the right words, ‘Snow is white,’ I express my very belief. If you understand what I assert and trust me, you come to have a belief with the same content as I.

Every theory of communication needs to say what plays these three roles. It also needs to say how SPEAKER’S MENTAL CONTENT, SPEECH ACT CONTENT, and HEARER’S MENTAL CONTENT are related:

BELIEF-SPEECH COORDINATION: the connection between the speaker’s belief content and the content of the speech act

SPEECH-BELIEF COORDINATION: the connection between the speech act’s content and the hearer’s belief content

The Lockean picture again offers a straightforward account: These connections are simply the identity relation. Note that SPEECH ACT CONTENT is what the expressions used in a speech act semantically express. It coincides, roughly, with what Grice calls *what the speaker says* rather than with Grice’s *what the speaker means*, which may involve pragmatically implicated content.¹³

4.4 The Conflict

There is *prima facie* evidence that we do communicate our self-locating beliefs. If mad Heimson, not without a certain reputation among his contemporaries, were to ask, ‘Who am I?’ and received the answer, ‘You are Heimson,’ it seems that he would be told just the piece of *de se* information that could set him straight.¹⁴ But this datum is in conflict with the Lockean picture. Suppose Ben has the *de se* belief that he is hungry and says to Anna, ‘I am hungry.’ The Lockean picture predicts that he is expressing the interesting centered content of his belief, which Anna will come to believe if she understands and accepts his assertion. That is, Anna will come to locate *herself* in the content and will thus believe *de se* that *she* is hungry. But what Ben communicates is obviously some other information – information Anna grasps if she comes to have a belief to the effect that the speaker, Ben, is hungry. Call this problem the *de se* problem.¹⁵

The problem for self-locating belief on the Lockean picture gets worse. It may seem that the right conclusion to draw from the *de se* problem is that hearers systematically *infer* an appropriate self-locating belief centered on themselves from the fact that

¹³A complete account of linguistic communication will have to account for pragmatically conveyed information, disambiguation, indirect speech acts and other pragmatic phenomena as well. Here, I am interested only in that part which accounts for the connection between mental content, linguistically expressed content, and compositional semantic theory. By a ‘theory of communication,’ I shall mean, following Lewis (1975, 1980), a systematic restatement of speakers’ common knowledge of their practice of linguistic communication.

¹⁴See Torre (2010) for reasons to think that communication of *de se* information is possible.

¹⁵This problem was first raised by Stalnaker (1981, 146-7).

the speaker asserted a content centered on himself. For instance, Anna may infer the centered content such that the center is being addressed by someone hungry from the fact that Ben asserted the centered content such that the center is hungry. But this is not what happens in the communication of self-locating beliefs about matters of taste. Suppose Ben believes that some cookie is tasty and says to Anna, ‘This cookie is tasty,’ thereby expressing the centered content *COOKIE*.

(4.2) *COOKIE*: $\{\langle w, t, x \rangle$: the (contextually salient) cookie tastes good to x in w at t \}.

What kind of thing will Anna come to believe if she understands and trusts him? She will *not* just come to locate herself in a content such that the center is addressed by someone to whom the cookie tastes good. On the contrary, if Anna accepts the claim, she will come to have the very self-locating belief that Ben has, viz. a belief with the centered content such that the cookie tastes good to the center. For Anna to accept an assertion of ‘This cookie is tasty’ is for her to locate *herself* in a cookie-liking location.

Centered worlds *de se* content and centered worlds content about matters of taste play incompatible roles in communication. With successful assertions about taste, the hearer comes to believe the same centered content as the speaker. With successful assertions about oneself, the hearer does not come to believe the same centered content. Call this problem the *incompatibility problem*.

It may seem that we have to give up either the centered content belief model or the Lockean picture of communication. But this would be hasty. In the next sections, I will propose an account that preserves the simplicity of the Lockean picture and the self-locating nature of belief by modifying the notion of centered content.

4.5 Sequenced Worlds Content

A centered world represents a possible way one individual may be. That is enough for belief as *self*-location, but not for communication. In communication, we are not just trying to locate ourselves individually. We are trying to locate ourselves as a group. We are trying to arrive at a common view about our collective location and everyone’s position in it. And for that, the possible ways different individuals may be need to be represented. If I tell you ‘It’s my turn,’ I am talking about myself in terms of my own possibilities. If I tell you, ‘It’s your turn,’ I am talking about you in terms of your possibilities. The fundamental problem with centered worlds content on the Lockean picture of communication is that the single center needs to represent sometimes the speaker, sometimes the addressee, and sometimes both.

The problem can be solved by introducing a *sequence* of centers. A *sequenced world* is a world centered on a number of individuals. It represents a possible way that a plural-

ity of individuals might be without conflating their individual possibilities. Formally, a sequenced world is a triple consisting of a world, a time, and a sequence of individuals inhabiting the world at the time. A sequenced worlds content p is the set of $\langle w, t, \langle x_1, \dots, x_n \rangle \rangle$ -triples such that p is true at $\langle w, t, \langle x_1, \dots, x_n \rangle \rangle$. Lewis thought of centered worlds contents as properties of individuals. Similarly, sequenced worlds contents can be thought of as properties of ordered n -tuples of individuals.¹⁶

I will here present a sequenced worlds model of communication that is a development of Ninan's (2010b) and Torre's (2010) accounts, which use sequenced worlds content to solve the first of the above problems – the *de se* problem.¹⁷ What I will show is that a suitably developed sequenced worlds model provides a solution to the whole problem. It yields a unified account of belief and communication for *de se* contents as well as contents about matters of taste.

SPEAKER'S MENTAL CONTENT, SPEECH ACT CONTENT, and HEARER'S MENTAL CONTENT are now sets of sequenced worlds, with one slot in the sequence for each conversational participant. Whose possibilities each slot represents must be stable in communication. Otherwise the problem would persist. If the first center, say, were to represent Ben's possibilities when he believes the content, but were to represent Anna's possibilities when she comes to believe the content, their individual possibilities would again be conflated. In order to stabilise what the content of speech acts and beliefs held during conversation represents, we relativise it to a *conversational context* c , a triple of a world w_c , time t_c , and an ordered list of conversational participants in w_c at t_c . Call the ordered list of conversational participants the *conversational sequence*. It is determined by the conversational facts. Who the participants to a conversation are depends on the mutually recognised intentions of speaker and audience. The order of the list of conversational participants does not matter, as long as we keep it stable for the entire conversation.¹⁸

Let us see how this helps with *de se* communication. We stipulate that for the conversation between Ben and Anna in w_c at t_c , $\langle w_c, t_c, \langle \text{Ben}, \text{Anna} \rangle \rangle$ is the conversational context. Then the content of Ben's assertion of 'I am hungry' is the sequenced

¹⁶Lewis (1983a, 28) himself provided the idea of worlds with multiple centers but did not use them to account for centered communication: "Besides possible individuals, world-sized and smaller, there are still other possibilities: joint possibilities for two or more individuals. These are ordered pairs, triples, etc. ... or even infinite sequences of possible individuals, all from the same world. An ordered pair of compossible individuals, for instance, is a way that a pair of individuals might possibly be." (See also Lewis (1986, 232-3))

¹⁷The model I develop here is close to Ninan's in key respects. Any shortcomings of the model are, of course, my responsibility. The model differs from Torre's in technicalities and one important philosophical respect: Torre gives up the Lockean picture of communication.

¹⁸The set of conversational contexts is a proper subset of the set of sequenced worlds – those sequenced worlds in which the individuals of the sequence are in a conversation with each other. For a given moment in a conversation, there are as many formal objects I call conversational contexts as there are ways of combining the participants into an ordered sequence.

worlds content $HUNGRY_t$:

$$(4.3) \quad HUNGRY_t: \{\langle w, t, \langle x_1, x_2 \rangle \rangle: x_1 \text{ is hungry in } w \text{ at } t\}$$

$HUNGRY_t$ says, roughly, that the first center x_1 is hungry. Given the conversational sequence $\langle \text{Ben}, \text{Anna} \rangle$, Ben's possibilities are represented by the first center, and Anna's possibilities by the second center. So for Ben to believe $HUNGRY_t$ is for him to believe *de se* that he is hungry. For Anna to believe $HUNGRY_t$ is for her to believe *de te* – 'of you' – that Ben is hungry. It is not for her to believe *de se* that she herself is hungry. So if Ben believes what he says and if Anna understands and accepts Ben's assertion, he and Anna will come to believe the same sequenced worlds content $HUNGRY_t$. However, their doxastic states are not exactly the same, as they dispose them to different actions. (We will come back to this difference in section 4.6.) This solves the *de se* problem.

Talk about taste need not distinguish between centers in the same way that *de se* communication must. If Ben successfully communicates 'This cookie is tasty' to Anna, they will each come to locate themselves in a cookie-liking location. On the sequenced world picture, we get this result if taste contents place conditions on every center. Let us again take $\langle w_c, t_c, \langle \text{Ben}, \text{Anna} \rangle \rangle$ as the conversational context. If Ben believes and asserts 'This cookie is tasty' in c , he expresses the sequenced worlds content $COOKIE_{1\&2}$:

$$(4.4) \quad COOKIE_{1\&2}: \{\langle w, t, \langle x_1, x_2 \rangle \rangle: \text{the cookie tastes good to } \langle x_1, x_2 \rangle \text{ in } w \text{ at } t\}$$

If communication is successful and Anna accepts Ben's assertion relative to the conversational sequence $\langle \text{Ben}, \text{Anna} \rangle$, she comes to locate herself in the set of sequenced worlds such that the cookie tastes good to *all* centers. And that seems right. If Ben wishes to establish that the cookie is tasty by asserting 'This cookie is tasty,' he has succeeded if they both locate themselves among the cookie-likers. This solves the incompatibility problem: *de se* and subjective sequenced worlds contents do not play incompatible roles in communication. The communication of either is successful in case the hearer comes to believe the same content the speaker believed and expressed in speech.

Belief in $COOKIE_{1\&2}$ is different from purely egocentric belief whose correctness depends only on one's own correct location. Anna should believe $COOKIE_{1\&2}$ only if she believes that the cookie tastes good to the speaker and that it tastes good to herself. The latter belief egocentrically concerns her taste, the former is safe as long as she takes Ben's assertion to be sincere. Sequenced worlds content on the Lockean picture captures the fact that success in the communication of subjective, evaluative claims involves acceptance of a common perspective on the matter.

4.6 Belief in Context

Sequenced worlds content, on the Lockean picture, is what is expressed and believed by speaker and audience. To believe a sequenced worlds content in a conversational context is to have a belief with a content whose sequenced worlds have sequences with as many individuals as there are parties to the conversation. It is to locate oneself as well as everyone else in the conversation; it is to locate the group of which one is a member, in a way that allows for the respective perspectives of the members to differ. In Lewis's terms, believing a centered worlds content is ascribing a property to oneself. Believing a sequenced worlds content in conversation is ascribing a property to the group of which one is a member.

The notion of believing a sequenced worlds content must be relativised to a conversational context and a believer. Here is why. Suppose Lingens says to his cousin Ortcutt, 'I am tired of reading.' If he is communicating successfully, then relative to the conversational context $\langle w_c, t_c, \langle \text{Lingens, Ortcutt} \rangle \rangle$ they will both end up believing $TIREDD_I$:

$$(4.5) \quad TIREDD_I: \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : x_1 \text{ is tired of reading in } w \text{ at } t \}$$

But even when Lingens and Ortcutt believe the same sequenced worlds content $TIREDD_I$ relative to the conversational context, there is still an important difference between their belief states. Their beliefs will dispose them to different actions – perhaps a disposition to stop reading for Lingens, and perhaps a disposition to say, 'Why don't you take a break?' for Ortcutt. This difference in belief states is accounted for by relativising belief to agents in conversational contexts. The complete account of sequenced worlds belief is as follows:

N-BELIEF IN A CONVERSATIONAL CONTEXT

An agent A n -believes a sequenced worlds content $\{ \langle w, t, \langle x_1, \dots, x_u \rangle \rangle : p(w, t, \langle x_1, \dots, x_u \rangle) \}$ in the conversational context $\langle w_c, t_c, \langle y_1, \dots, y_u \rangle \rangle$ iff

- (i) $A = y_n \in \{ y_1, \dots, y_u \}$
- (ii) there are relations $R_1 \dots R_u$ such that in w_c at t_c , y_n is uniquely R_1 -related to y_1, \dots , and y_n is uniquely R_u -related to y_u (where R_n is the identity relation) and y_n 's standing in $R_1 \dots R_u$ to $y_1 \dots y_u$ establishes a conversation between $y_1 \dots y_u$
- (iii) every sequenced world $\langle w', t', \langle x'_1, \dots, x'_u \rangle \rangle$ compatible with what A believes in w_c at t_c is such that $p(w', t', \langle x'_1, \dots, x'_u \rangle) = \mathbf{I}$.

Informally, N-BELIEF IN A CONVERSATIONAL CONTEXT says that an agent n -believes a sequenced worlds content p in a conversational context just in case (i) the agent is the

n th member of the conversational sequence, (ii) the agent stands in relations to every member of the conversational sequence which establish a conversation between them, and (iii) the agent believes that she might be the n th member of a group of which p is true.

N-BELIEF IN A CONVERSATIONAL CONTEXT entails that Lingens I -believes $TIRED_t$ in the conversational context $\langle w_c, t_c, \langle \text{Lingens, Ortcutt} \rangle \rangle$ just in case (i) $\text{Lingens} \in \{\text{Lingens, Ortcutt}\}$, (ii) there are conversation-establishing relations R_1 and R_2 such that in w_c at t_c R_1 uniquely relates Lingens to himself (identity) and R_2 relates Lingens to Ortcutt (e.g., the addressing relation), and (iii) every sequenced world $\langle w', t', \langle x'_1, x'_2 \rangle \rangle$ compatible with what Lingens believes in w_c at t_c is such that x'_1 is tired of reading in w' at t' . A sequenced world $\langle w', t', \langle x'_1, x'_2 \rangle \rangle$ is compatible with what Lingens believes in w_c at t_c if Lingens thinks in w_c at t_c that he might be the member x'_1 of a group $\langle x'_1, x'_2 \rangle$ in w' at t' whose members are related by R_1, R_2 .

There is nothing mysterious about an agent's believing a sequenced worlds content in a conversational context, once we accept centered worlds content and individual self-location. We can, if we want, translate sequenced worlds belief into centered worlds belief. Put simply, the idea is that to ascribe a property to the group of which one is a member is equivalent to self-ascribing the property of being a member of a group that has this property. For Ortcutt to believe the sequenced worlds content p expressed by Lingens' assertion of 'I am tired of reading' in the conversational context $\langle w_c, t_c, \langle \text{Lingens, Ortcutt} \rangle \rangle$ is for him to be addressed by Lingens and to believe the centered worlds content $p' = \{\langle w, t, x \rangle: \text{there is a } y, x \text{ is } y\text{'s addressee in } w \text{ at } t, \text{ and } p(w, t, \langle y, x \rangle)\}$. More generally, the following equivalence holds:

SEQUENCED WORLDS BELIEF AND CENTERED WORLDS BELIEF

An agent A n -believes a sequenced worlds content $\{\langle w, t, \langle x_1, \dots, x_u \rangle \rangle: p(w, t, \langle x_1, \dots, x_u \rangle)\}$ in the conversational context $\langle w_c, t_c, \langle y_1, \dots, y_u \rangle \rangle$ iff

- (i) $A = y_n \in \{y_1, \dots, y_u\}$,
- (ii) there are relations $R_1 \dots R_u$ such that in w_c at t_c , y_n is uniquely R_1 -related to y_1, \dots , and y_n is uniquely R_u -related to y_u (where R_n is the identity relation) and y_n 's standing in $R_1 \dots R_u$ to $y_1 \dots y_u$ establishes a conversation between $y_1 \dots y_u$
- (iii) A believes the centered worlds content $\{\langle w, t, x \rangle: \text{there are individuals } x_1, \dots, x_u \text{ such that } x \text{ is uniquely } R_1\text{-related to } x_1, x \text{ is uniquely } R_2\text{-related to } x_2, \dots, \text{ and } x \text{ is uniquely } R_u\text{-related to } x_u \text{ in } w \text{ at } t, \text{ and } p(w, t, \langle x_1, \dots, x_u \rangle)\}$.

For $n \neq m$, n -believing and m -believing in a conversational context are two different doxastic states, with different potential effects on action. If Lingens and Ortcutt

communicate successfully, both come to have beliefs with the same sequenced worlds content p . However, in the conversational context $\langle w_c, t_c, \langle \text{Lingens, Ortcutt} \rangle \rangle$, Lingens will come to 1-believe p , which probably disposes him to stop reading, while Ortcutt will come to 2-believe p , which will not dispose him to such action. SEQUENCED WORLDS BELIEF AND CENTERED WORLDS BELIEF makes this evident: When Lingens 1-believes and Ortcutt 2-believes p , they believe the same content p , but for each the centered worlds belief equivalent to his belief in p is different. We will come back to the connection between doxastic states and action dispositions in section 6.4.

A word of clarification on SEQUENCED WORLDS BELIEF AND CENTERED WORLDS BELIEF. The equivalence between sequenced worlds belief and centered worlds belief can be read in three ways. First, as stating an equivalence between two equally basic and theoretically useful notions of belief. Second, as stating a reductive explanation of believing a sequenced worlds content in terms of believing a centered worlds content. The fundamental notion of belief then is individual self-location. Third, the equivalence can be read as stating a reductive explanation of belief in centered worlds content in terms of belief in sequenced worlds content. The fundamental notion of belief then is collective self-location, or group-location. On the last option, individual self-location is a special case of collective self-location where the group consists of only one member. I am inclined to think the fundamental notion is group-location, but I will not offer independent arguments here. Suffice it to say that taking sequenced worlds belief and sequenced worlds content as fundamental supports the Lockean assumption that communication essentially involves the sharing of a single content. If centered worlds belief is fundamental and speaker and hearer in successful communication have beliefs with different centered worlds contents, introducing a shared content at the non-fundamental level of sequenced worlds belief hardly goes a long way towards saving the Lockean picture.

4.7 Assertion and the Common Ground

4.7.1 Common Ground

On the sequenced worlds model, SPEAKER'S MENTAL CONTENT, SPEECH ACT CONTENT, and HEARER'S MENTAL CONTENT are one and the same sequenced worlds content. Thus BELIEF-SPEECH COORDINATION and SPEECH-BELIEF COORDINATION are given by the identity of these contents. The basic Lockean idea that one piece of information travels from speaker's head to hearer's head is preserved.

The sequenced worlds model also fits naturally with a Stalnakerian implementation of the Lockean picture. I will first sketch Stalnaker's original account and then make the changes needed to accommodate sequenced worlds content.

According to Stalnaker, linguistic communication is primarily a matter of updating and establishing a body of shared information – the common ground.¹⁹ Speech acts serve to influence this body of information in various ways. In particular, the essential effect of assertion is to add the asserted content to the common ground. The attitude that speakers strike towards the common ground is the attitude of *presupposition*:

... the presuppositions of a speaker are the propositions he takes for granted as part of the background of the conversation. A proposition is presupposed if the speaker is disposed to act as if he assumes or believes that the proposition is true, and as if he assumes or believes that his audience assumes or believes that it is true as well. (Stalnaker, 1978, 84)

Presupposition, in this sense, is a public attitude: one presupposes a proposition *p* only if one presupposes that everyone else in the conversation also presupposes *p*. A speaker's presuppositions are represented by the speaker's *context set*: the set of possible worlds compatible with what the speaker presupposes. (Propositions, for Stalnaker, are also sets of possible worlds; a speaker's context set is the intersection of the propositions she presupposes.) There is a context set for each participant in a conversation, but when things go as they should, all participants make the same presuppositions and the speakers' context sets coincide with the *conversation's context set*. The common ground is represented by the conversation's context set, which is the intersection of the propositions in the common ground. A conversation is *defective* when the conversation's participants do not all make the same presuppositions.²⁰

Assertions are proposals to add information to the common ground. When an assertion of *p* is understood and accepted by all participants in a conversation, its content *p* becomes presupposed in the conversation, and its effect is to eliminate all the non-*p* worlds from the conversation's context set. "To engage in conversation is, essentially, to distinguish among alternative possible ways that things may be." (Stalnaker, 1978, 85) An assertion's primary contribution is to narrow down what the participants commonly take to be the possible relevant ways the world might be.

¹⁹Stalnaker's views on the pragmatics of communication are developed in his 1970, 1974, 1978, and 2002, among others.

²⁰In the default case, the common ground will be common knowledge or common belief. But it need not be. Interlocutors may take different non-public attitudes towards what is presupposed, depending on the purpose of the conversation. When the purpose is to establish truth, the attitude is, plausibly, knowledge or belief; when speakers are interested in exploring a hypothetical situation, the mutually recognised non-public attitude is supposition; asf. What the right non-public attitude to take is, may itself be a matter of negotiation between interlocutors. I will for the most part focus on the default case where the common ground is common belief. Even when presupposition entails belief, however, the converse is not true. Given the public nature of presupposition, a speaker typically believes a variety of things she does not believe her audience to believe, or she may believe that her audience believes them but not that they believe that she believes them, etc.

On the sequenced worlds model, the conversation's context set is a set of sequenced worlds whose sequences have as many centers as the conversation has participants. To engage in conversation is to distinguish between alternative ways that the conversational participants might be, where this does not require that they all share the ways they *individually* might be. Intuitively, the purpose of conversation is the coordination of individual perspectives, sometimes with the result of sharing a perspective, sometimes with the result of having one's individual perspective noticed.

Assertions serve this purpose, if successful, by adding the sequenced worlds content they express to the common ground. When in the common ground, that content is presupposed by all conversational participants. We can define the notion of speaker presupposition for a context set containing sequenced worlds on the basis of n -belief in a conversational context:

SPEAKER PRESUPPOSITION_{SW}

A speaker S n -presupposes a sequenced worlds content p in a conversational context $\langle w_c, t_c, \langle x_1, \dots, x_u \rangle \rangle$ iff $S = x_n \in \{x_1, \dots, x_u\}$ and S is disposed to act as if she n -assumes or n -believes p , and as if she n -assumes or n -believes that for all $x_i \in \{x_1, \dots, x_u\}$, x_i i -assumes or i -believes p as well.

If a sequenced worlds content p is part of the common ground in the default case where the common ground is common belief, every participant in the conversation x_n n -believes p .

The Stalnakerian model with sequenced worlds content vindicates the Lockean idea that one content is what is expressed by the speaker and believed by all participants if communication is successful. At the same time, belief on the model still involves self-location, although belief in a conversational context involves locating not just oneself but the conversational group of which one is a member. The model solves the *de se* problem and the incompatibility problem by keeping centers and individual possibilities separated where necessary and allowing for joint possibilities to be established where this is, intuitively, the effect of assertion. Thus, the conflict between the Lockean picture of communication and the self-locating belief model can be resolved without giving up either of them.

4.7.2 Assertion and Acceptance

We saw that belief in context is location of the conversational group of which one is a member. The correct belief-norm is thus not EGOCENTRICITY but a group-centric norm that requires the group to be correctly located in the believed sequenced worlds content. What about assertion and acceptance of sequenced worlds content in conversation? Our model of communication should tell us under which conditions it is felicitous for

speakers to assert a sequenced worlds content p , and when it is a good idea for hearers to accept p into the common ground.

Clearly, egocentric norms will not make the right predictions. To see this, let us consider one mainstream approach to norms of assertion, which states the crucial necessary condition for felicitous assertion in terms of truth of the asserted content.²¹ First, it will be helpful to distinguish between two kinds of perspective at which a content may be true.

INDIVIDUAL PERSPECTIVE $P^I = \langle w, t, x \rangle$

An individual perspective P^I represents the perspective of a single individual (her and only her location and world view) in the world w at the time t .

THE CONVERSATION'S PERSPECTIVE $P^N = \langle w_c, t_c, \langle x_1, \dots, x_n \rangle \rangle$

The perspective of a conversation is $\langle w_c, t_c, \langle x_1, \dots, x_n \rangle \rangle$, where w_c and t_c are the world and time at which the conversation takes place, and the sequence of individuals $\langle x_1, \dots, x_n \rangle$ is determined by the conversational sequence for the conversation. It represents the individual perspectives of all conversational participants x_1, \dots, x_n in w_c at t_c .

Let us say that a sequenced worlds content is true from an individual perspective just in case it correctly represents the location of that individual – no matter the location of the other individuals in the sequence. We can then state the following egocentric truth norm of assertion:

²¹ Stephenson (2007a) endorses an egocentric belief norm of assertion. MacFarlane (2012) and Kölbel (2002) endorse an egocentric truth norm equivalent to the one stated here (they do not account for interesting *de se* content, however), and so do Lasersohn (2005) and Egan et al. (2005), with some qualifications. For Lasersohn only 'autocentric' uses of predicates of personal taste, those on which the asserter takes herself to be the relevant 'judge,' are governed by an egocentric norm. Utterances with 'exocentric' uses of predicates of personal taste, those on which the asserter takes someone else as the relevant 'judge,' are assertable only if the sentence used is true in the context and relative to the perspective of the contextually determined experiencer. In Lasersohn's own words: "we typically *assert* from an autocentric perspective. That is, in making an assertion, we regard it as in some sense *justified* iff it is true relative to that context corresponding to the concrete situation of utterance in which we ourselves serve as judge. . . Under certain circumstances, we may also adopt an *exocentric* perspective, assessing sentences for truth relative to contexts in which someone other than ourselves is specified as the judge, or regarding our assertions as justified by virtue of their truth relative to such contexts." (Lasersohn, 2005, 670)

Egan et al. seem to take an egocentric norm of assertion as the default: "The correct norm [of assertion] is that one should only say something that's true when evaluated in the context you are in." (Egan et al., 2005, 153) (They provide ample evidence following this statement that 'you' here picks out the asserter.) However, Egan et al. add that the speaker's context need not *always* correspond to what I call her individual perspective: "We assume here that contexts can include more than just the speaker. If Vinny the Vulture is speaking to a group of humans he arguably cannot say '*Rotting flesh tastes great*'. The reason is that rotting flesh does not taste great to the group of speakers in the conversation, most of whom are humans." (153) The proposal of this chapter is to make the 'context' that includes not just the speaker but all parties to the conversation the default.

EGOCENTRIC TRUTH NORM

ASSERT_E A speech act content p is appropriately assertable in context c only if p is true from the speaker's individual perspective P_c^I in c .

ASSERT_E may seem to make the right predictions for *de se* assertions. Intuitively, Heimson may assert 'I am Heimson' only if the center-slot representing Heimson correctly locates him. But ASSERT_E fails to make the correct predictions for *de te* assertions. It does not prohibit speakers to tell anyone except Heimson, 'You are Heimson,' as it should. The speech act content expressed by 'You are Heimson' places a constraint only on a center different from the center representing the speaker. So as long as someone in the world and at the time of the conversation is Heimson, the speech act content is true from the speaker's individual perspective.

The right norms of assertion and acceptance, which go hand in hand with belief in context as group-location, are group-centric norms. They are here stated as truth norms:²²

GROUP-CENTRIC TRUTH NORMS

ASSERT_G A speech act content p is appropriately assertable in context c only if p is true from the conversation's perspective in c .

ACCEPT_G A speech act content p is acceptable in context c if p is true from the conversation's perspective in c .

Suppose Ben tells Heimson in a context of utterance c , 'You are mad.' Let the conversational context be $\langle w_c, t_c, \langle \text{Ben, Heimson} \rangle \rangle$ so that the speech act content of Ben's assertion is $\{ \langle w, t, \langle x_1, x_2 \rangle \rangle : x_2 \text{ is mad in } w \text{ at } t \}$. The conversation's perspective P_2^c is $\langle w_c, t_c, \langle \text{Ben, Heimson} \rangle \rangle$. According to ASSERT_G, Ben's assertion is appropriate only if $\{ \langle w, t, \langle x_1, x_2 \rangle \rangle : x_2 \text{ is mad in } w \text{ at } t \}$ is true from $\langle w_c, t_c, \langle \text{Ben, Heimson} \rangle \rangle$. This is as it should be. Ben should make the assertion only if – and Heimson should accept the assertion if – Heimson is mad at the time and world of speaking.²³

²²Truth norms may be understood as default norms. The purpose of the conversation may determine stronger norms where the purpose is to establish common knowledge, or weaker norms where the purpose is to establish, e.g., common supposition or common pretence (cf. fn. 20 in section 4.7.1). For instance, when the mutually recognised attitude towards the common ground is common pretence, the norms of assertion and acceptance are as follows:

ASSERT_{PRETENCE} A speech act content p is appropriately assertable for a speaker S in context c where pretence is the purpose of the conversation only if p is true at all the worlds compatible with what S pretends to be the case.

ACCEPT_{PRETENCE} A speech act content p is acceptable for an addressee A in context c where pretence is the purpose of the conversation if p is true at all the worlds compatible with what A pretends to be the case.

²³Note that to assess whether $\{ \langle w, t, \langle x_1, x_2 \rangle \rangle : x_2 \text{ is mad in } w \text{ at } t \}$ is true from $\langle w_c, t_c, \langle \text{Ben, Heimson} \rangle \rangle$, one need not be able to identify x_2 *de re* as Heimson. We use the names 'Ben' and 'Heimson' to state

For talk about taste, the group-centric norms entail that, for instance, ‘This cookie is tasty’ is assertable only if the cookie tastes good to speaker *and audience*. This is because all interlocutors have to be correctly located in the content, which says of each one of them that the cookie is tasty to them.

But this prediction might seem too strong. Why should a speaker have to make sure that she and her audience have a common outlook on taste in order to guarantee that her assertion about the cookie’s tastiness is appropriate? Is not the subjectivity of taste claims better captured by an egocentric norm like *ASSERT_E*?

I will postpone full discussion of the appropriateness of bare taste assertions to chapter 5, where we will see that the sequenced worlds view can account for the subjectivity of taste claims. For now, notice a counterintuitive consequence of purely egocentric truth norms. As MacFarlane (2012, 179) observes, an egocentric truth norm licenses asserting bare taste claims such as ‘Schnitzel is tasty’ just when it licenses making assertions about what is tasty to the speaker herself, such as ‘Schnitzel is tasty to me.’ This is because the former is true from the speaker’s own perspective whenever the latter is. However, there are many conversational situations in which asserting the explicitly relativised taste sentence is natural but asserting the bare taste claim is odd, as in (4.6c) and (4.7c).

- (4.6) a. Ben: Schnitzel is tasty.
 b. Anna: No, it’s not tasty! It is bland.
 c. Ben: Well, it’s tasty to *me*, at least.
- (4.7) a. Ben: Schnitzel is tasty.
 b. Anna: No, it’s not tasty! It is bland.
 c. Ben: ? Well, it is tasty.

Ben’s assertion in (4.6c) amounts to a ‘partial retraction’ of his first assertion.²⁴ The group-centric norm predicts that (4.6c) is felicitous but (4.7c) is not, since from the conversational perspective at the time of Ben’s second assertion, it is true that Schnitzel tastes good to Ben while it is not true that Schnitzel tastes good to Ben and Anna. The egocentric norm, on the contrary, predicts that both (4.6c) and (4.7c) are felicitous, since from Ben’s perspective at the time of his second utterance, it is both true that Schnitzel is tasty and that Schnitzel is tasty to Ben himself. Of course, an explanation of the conversational impropriety of (4.7c) may appeal to other norms than the norm of

what the conversational sequence is merely for convenience. All that is needed to determine whether some content is true from the conversational perspective is the ability to keep participants apart and consistently track them in derivations of content from context. For participants in one-to-one conversations, the ability to distinguish between oneself and not-onself and to track them, respectively, suffices.

²⁴Cf. Pearson (forthcoming, §4.6) on ‘partial retraction’ and what Schaffer (2011) calls ‘entrenchment’ for a related phenomenon. The example in (4.6) is adapted from Pearson (forthcoming, §4.6).

assertion. But the example shows at the very least that egocentric truth norms are blind to conversational circumstances that transcend the speaker's individual perspective.

4.8 A Semantic Proposal

Understanding *SPEECH ACT CONTENT* in terms of sequenced worlds was the key to solving the *de se* and incompatibility problems in a way that reconciles the self-locating account of belief with the Lockean picture of communication. But so far I have merely claimed that *SPEECH ACT CONTENT* is sequenced worlds content. I have not yet shown what the relation is between sentences – such as ‘This cookie is tasty’ and ‘I am hungry’ – and this kind of speech act content. For the sequenced worlds model to be plausible, it needs to be completed with a semantics of predicates of personal taste and of personal pronouns and an account of how the semantics determines sequenced worlds speech act content.

The sequenced worlds model does not require a radical departure from existing semantic proposals. For instance, the standard Kaplanian treatment of personal pronouns can be combined with a modified version of Stephenson's (2007a) semantics for predicates of personal taste to make room for sequenced worlds content. Other options are available, but for concreteness I will introduce a combination of the above in this section.²⁵ In the next section, I will show how this semantics delivers sequenced worlds speech act content.

Our starting point is a Kaplan-style intensional semantic theory on which extensions are assigned relative to a context c and an index i .²⁶ An expression's semantic value is a function from a context and an index to an extension; we will also say that an expression's semantic value *at a context and index* is an extension. A context c is a possible occasion of use of an expression, which determines at least a world, a time, a speaker and addressee(s), and a location. Formally, we will model a context as a sequenced world $\langle w, t, \langle x_1, \dots, x_n \rangle \rangle$, where x_1 represents the speaker and x_2 through x_n represent the addressees. An index i is a sequence of independently shiftable features of context, called coordinates. In the semantics we need, the index is a triple $\langle w, t, \langle x_1, \dots, x_n \rangle \rangle$.²⁷ The index is the first modification of Stephenson's system, which has $\langle w, t, x \rangle$ -triples as indices. The double brackets ‘[[]]’ denote the interpreta-

²⁵For instance, a semantics with sequenced worlds indices can also be formulated by modifying Lasersohn's (2005; 2008) semantics for predicates of personal taste (see footnote 33 in section 4.8 below for some details).

²⁶For a clear exposition of an ‘orthodox’ version of the system with possible worlds as the only coordinates of the index, see Heim and Kratzer (1998, ch.12) and von Stechow and Heim (2011).

²⁷The index may or may not be a sequenced world depending on whether a possible situation corresponds to the combination of world, time, and individuals. For instance, $\langle @, 16 \text{ June } 1902, \langle \text{Frege, Russell} \rangle \rangle$ is a sequenced world but $\langle @, 16 \text{ June } 2011, \langle \text{Frege, Russell} \rangle \rangle$ is not, since Frege and Russell inhabited @ in 1902 but not in 2011. Both triples, however, can be values of the index.

tion function, a three-place function that maps an expression, a context and an index to an extension.

The extensions of standard one-place predicates like ‘hungry’ depend on the world- and time-coordinate of the index, but not on any individual in the sequence.

$$(4.8) \quad \llbracket \text{hungry} \rrbracket^{c,i} = [\lambda y_e. y \text{ is hungry in } w \text{ at } t],$$

for $i = \langle w, t, \langle x_1, \dots, x_n \rangle \rangle$ ²⁸

Predicates of personal taste (PPTs) such as ‘tasty’ and ‘fun’,²⁹ on Stephenson’s view, are two-place predicates. They are functions from an experiencer and experienced object or individual, and a context and an index, to truth values.³⁰

$$(4.9) \quad \llbracket \text{tasty} \rrbracket^{c,i} = \llbracket \text{taste good} \rrbracket^{c,i} = [\lambda y_e. [\lambda z_e. z \text{ tastes good to } y \text{ in } w \text{ at } t]]$$

$$\llbracket \text{taste terrible} \rrbracket^{c,i} = [\lambda y_e. [\lambda z_e. z \text{ tastes terrible to } y \text{ in } w \text{ at } t]]$$

$$\llbracket \text{fun} \rrbracket^{c,i} = [\lambda y_e. [\lambda z_e. z \text{ is fun for } y \text{ in } w \text{ at } t]]$$

²⁸Read ‘ $[\lambda y_e. y \text{ is hungry in } w \text{ at } t]$ ’ as ‘the function which maps every y_e to 1 (truth) if y is hungry in w at t , and to 0 (falsity) otherwise.’ The subscript ‘ e ’ indicates y ’s semantic type. e is the semantic type of individuals, s stands for worlds, i for times, and t for truth values; combinations such as et , or $\langle e, t \rangle$, stand for functions from the first type (e) to the second (t).

²⁹I shall for the most part restrict my attention to paradigm examples of predicates of personal taste such as ‘fun’ and ‘tasty’. The demarcation of a class of predicates of personal taste from aesthetic and other evaluative predicates is difficult, as Lasersohn (2005) observes. MacFarlane (2012) gives his semantics only for ‘tasty’. Lasersohn (2005, 2008), Schaffer (2011), and Stephenson (2007a) limit their analyses to ‘tasty’, ‘taste good/terrible’, and ‘fun’. (See Lasersohn (2008, §1.2) for a non-definitional demarcation of predicates of personal taste based on the interaction of perspective with scalarity.) Cappelen and Hawthorne (2009) extend this list to include examples like ‘spicy’, ‘funny’, ‘delicious’, ‘disgusting’, ‘nauseating’ and also ‘filling’. Richard (2008, 10) gives an indication of what he considers predicates concerning matters of taste to be when he writes: ‘A likely place to look for relative truth, it seems, is in matters of taste. Our detailed conceptions of who is handsome, who sexy, what is hip, what L7, what is boring, what perverted—these typically differ.’ Richard subsequently develops his relativist semantics for standard examples of gradable adjectives like ‘rich’ and ‘tall’, and endorses truth relativism for all vague expressions (Richard, 2008, 101). For reasons related to faultless disagreement, Richard believes that at least some claims about taste require an expressivist rather than relativist treatment. Moltmann (2010) includes ‘delicious’, ‘pleasant’, ‘nice’ and intends her semantics to cover moral predicates, aesthetic predicates, and gradable adjectives such as ‘tall’, or ‘rich’. Egan (2010a) gives an account of disputes of aesthetic taste in general, but makes no claims concerning the semantic details of his key examples such as ‘good’, ‘beautiful’, ‘tasty’, ‘fun’, ‘elegant’, ‘ugly’, ‘disgusting’. Egan et al. (2005) are mainly concerned with epistemic modals but claim that relativist theories are also plausible for expressions such as ‘huge’, ‘being the same colour’, and ‘taste [good/bad]’ (2005, 152).

³⁰Stephenson does not discuss the scalar aspect of many predicates of personal taste like ‘tasty’ and ‘fun’, and I will also be neutral on the question of how to account for the gradable nature of PPTs. The analysis in this section is compatible with Richard’s (2008) relativist proposal that gradable adjectives have their comparison class determined by the context of use and thus express contents that are specific with respect to comparison class. Equally, elements of the degree analysis of gradable adjectives (Kennedy, 2000, 2007) are available to the relativist, although the details will have to be left for future work. See Lasersohn (2008, §1.2) for a brief discussion of the scalarity of PPTs. See also footnote 3 in section 5.2 below for discussion.

The entries for PPTs do not directly make their extensions dependent on the individuals in the index. However, in first-personal uses of PPTs, i.e. in bare taste claims such as ‘This cookie is tasty’ which are based on the asserter’s taste preferences but which do not have an overt experiencer argument in the sentence’s surface structure, there is a covert, phonologically null nominal item ‘PRO_C’ at the appropriate level of logical form. Simplifying the logical form, and ignoring tense and the contribution of the copula, (4.10) has the following structure.

(4.10) This cookie is tasty.
 [This cookie] [is tasty PRO_C]

PRO_C takes as its reference the sequence of centers in the index:

(4.11) $[[\text{PRO}_C]]^{c,i} = \langle x_1, \dots, x_n \rangle$,
 where the index $i = \langle w, t, \langle x_1, \dots, x_n \rangle \rangle$.

PRO_C is the second modification of Stephenson’s system, in which the nominal item PRO_J refers to the single judge given by the index. Like PRO_J, PRO_C is ‘not a pronoun in the sense of being able to be bound or controlled, nor is it an indexical since it takes its reference from the index rather than the context of use.’ (Stephenson, 2007a, 500) PRO_C thus introduces dependence of first-personal uses of PPTs on the individuals in the index into the system. The meaning of (4.10) is computed in (4.12).

(4.12) $[[\text{This cookie is tasty}]]^{c,i} =$
 $[[\text{tasty}]]^{c,i} ([[\text{PRO}_C]])^{c,i} ([[\text{this cookie}]])^{c,i} = 1$ iff
 the cookie tastes good to $\langle x_1, \dots, x_n \rangle$ in w at t

Treating PPTs as two-place predicates provides a straightforward handling of uses of PPTs with an overt prepositional phrases such as ‘for Ben’ in (4.13).

(4.13) The roller coaster is fun for Ben.

Stephenson assumes that in expressions like ‘fun for Ben’, or ‘tasty to Anna’, the prepositions ‘for’ and ‘to’ make no semantic contributions. That is, their semantic value is the identity function on individuals.

(4.14) $[[\text{for}]]^{c,i} = [\lambda y_e. y]$

The meaning of (4.13) is computed in (4.15).

(4.15) $[[\text{The roller coaster is fun for Ben}]]^{c,i} =$
 $[[\text{fun}]]^{c,i} ([[\text{for Ben}]])^{c,i} ([[\text{the roller coaster}]])^{c,i} = 1$ iff
 the roller coaster is fun for Ben in w at t

Finally, there are uses of PPTs, sometimes called ‘exocentric’, which have no overt experiencer but in which the context of use makes a particular individual salient, as in (4.16).³¹

- (4.16) Anna: How does Bill like the rides?
 Ben: Well, the merry-go-round is fun, but the water slide is a little too scary.

In uses like (4.16), the logical form contains a silent referential pronoun ‘pro_x’ referring to the individual *x*, which takes its reference from the context of use *c*.

$$(4.17) \llbracket \text{pro}_x \rrbracket^{c,i} = x$$

The logical form, at the appropriate level, of ‘the merry-go-round is fun’ in (4.16) is (4.18) with the meaning of (4.19).

$$(4.18) [\text{The merry-go-round}] [\text{is fun pro}_{\text{Bill}}]]$$

$$(4.19) \llbracket \text{The merry-go-round is fun} \rrbracket^{c,i} = \\ \llbracket \text{fun} \rrbracket^{c,i} (\llbracket \text{pro}_{\text{Bill}} \rrbracket^{c,i}) (\llbracket \text{the merry-go-round} \rrbracket^{c,i}) = 1 \text{ iff} \\ \text{the merry-go-round is fun for Bill in } w \text{ at } t$$

Like uses of PPTs with overt prepositional phrases, these ‘exocentric’ uses are non-first-personal; their semantic value at a context and index is not dependent on the individuals in the index.³²

Personal pronouns like ‘I’/‘me’ and second person singular ‘you’ receive a standard Kaplanian treatment.

$$(4.20) \llbracket \text{I} \rrbracket^{c,i} = \text{the speaker/agent of } c \\ \llbracket \text{you} \rrbracket^{c,i} = \text{the addressee of } c$$

So Ben’s utterance in (4.21) has the meaning given in (4.22).

³¹The example is adapted from Lasersohn (2005, 672).

³²In section 4.2.3, it was claimed that the inference from a belief *that X is tasty* to a belief *that X is tasty to me* is generally available. ‘Exocentric’ uses of PPTs may appear to be a counterexample to this claim. If Anna comes to believe what Ben says in (4.16) – ‘The merry-go-round is fun’ – she can *not* infer that the merry-go-round is fun for herself. But note that while it is true that Anna cannot make this inference, the ‘exocentric’ use of PPTs in (4.16) and other examples does not constitute a counterexample to the general availability of the inference from a belief *that X is tasty/fun* to a belief *that X is tasty to/fun for me*. In section 4.9, we will see that the content expressed by ‘exocentric’ uses has an experiencer argument different from any center. If made fully explicit, the content of taste beliefs corresponding to ‘exocentric’ uses is not *that X is tasty/fun*, but *that X is tasty to/fun for Y*, where *Y* may be non-identical to the believer. In (4.16), Anna will come to believe *that the merry-go-round is fun for Bill*. Thus beliefs corresponding to ‘exocentric’ uses of PPTs do not constitute counterexamples to the availability of the inference from a belief *that X is tasty* to a belief *that X is tasty to me*.

(4.21) Ben: I am hungry.

(4.22) $\llbracket \text{I am hungry} \rrbracket^{c,i} = \llbracket \text{hungry} \rrbracket^{c,i}$ ($\llbracket \text{I} \rrbracket^{c,i} = \mathbf{1}$ iff Ben is hungry in w at t)

We now have all we need for a compositional semantic theory for simple sentences expressing claims about taste and *de se* attitudes.³³

4.9 Speech Act Content

How do we get the sequenced worlds speech act content we need from the semantic values given above? In short, by taking the diagonal of a sentence's Kaplanian character. In what follows, I will give a definition of SPEECH ACT CONTENT that is a generalisation of Ninan's definition of *utterance content* (2010b, 18). Let us start with sentences expressing *de se* attitudes.

Kaplan took 'what is said' – the speech act content expressed – by an assertoric utterance of a sentence in context to be the function from index to truth values. Let us call this content, in our system, the Kaplan horizontal:

$$\text{Kaplan horizontal of } \Phi \text{ at } c: \quad \lambda i. \llbracket \Phi \rrbracket^{c,i} = \{ \langle w, t, \langle x_1, \dots, x_n \rangle \rangle : \llbracket \Phi \rrbracket^{c, \langle w, t, \langle x_1, \dots, x_n \rangle \rangle} = \mathbf{1} \}^{34}$$

³³As mentioned above, Stephenson's semantics for PPTs is not the only option. A relativist semantics with sequenced worlds indices can also be formulated by reworking Lasersohn's semantics for predicates of personal taste. Lasersohn (2005, 2009) treats PPTs as 1-place predicates whose denotation at a context and index may vary with the 'judge'-coordinate j of the index. With suitable notational adjustments, his entry for 'tasty' is as follows:

$$\llbracket \text{tasty} \rrbracket^{c, \langle w, t, j \rangle} = \lambda x_e. x \text{ tastes good to } j \text{ in } w \text{ at } t$$

This entry can straightforwardly be extended to make room for a sequence of judges:

$$\llbracket \text{tasty} \rrbracket^{c, \langle w, t, \langle j_1, \dots, j_n \rangle \rangle} = \lambda x_e. x \text{ tastes good to } \langle j_1, \dots, j_n \rangle \text{ in } w \text{ at } t$$

For Lasersohn, the preposition 'to' in 'tasty to Ben' is an operator on the judge-coordinate that shifts it to Ben. Replacing j again with a sequence of judges, we get the following entry for 'tasty to Ben':

$$\llbracket \text{tasty to Ben} \rrbracket^{c, \langle w, t, \langle j_1, \dots, j_n \rangle \rangle} = \llbracket \text{tasty} \rrbracket^{c, \langle w, t, \text{Ben} \rangle} = \lambda x_e. x \text{ tastes good to Ben in } w \text{ at } t$$

Lasersohn does not give 'exocentric' uses of PPTs a different semantic treatment but accounts for them in the pragmatics. Speakers assume a contextually salient perspective that is different from their own and evaluate the taste claim relative to the world-time-judge triple formed by the world and time of utterance and the individual whose perspective is salient.

This latter feature of an extended Lasersohn semantics and pragmatics stands in tension with the Stalnakerian model of communication I presented in section 4.7. Stephenson's treatment of 'exocentric' uses of PPTs requires no changes to that model, which makes her semantics the better fit.

³⁴Sequenced worlds contents can be characterised as sets of sequenced worlds and as the characteristic functions of such sets. For indices that take sequenced worlds with sequences of n individuals, $\lambda i. \llbracket \Phi \rrbracket^{c,i}$ is the characteristic function of the set $\{ \langle w, t, \langle x_1, \dots, x_n \rangle \rangle : \llbracket \Phi \rrbracket^{c, \langle w, t, \langle x_1, \dots, x_n \rangle \rangle} = \mathbf{1} \}$. The two formulations are essentially equivalent, and I will use both.

Given the standard Kaplanian semantic values of the pronouns ‘I’/‘me’ and ‘you,’ their reference is resolved in the derivation of the Kaplan horizontal from context. Thus, the Kaplan horizontal of (4.21) is (4.23).

$$(4.23) \lambda i. \llbracket \text{I am hungry} \rrbracket^{c,i} = \{ \langle w, t, \langle x_1, \dots, x_n \rangle \rangle : \text{Ben is hungry in } w \text{ at } t \}$$

But (4.23) is not the interesting sequenced worlds content which, as we saw above, is communicated by assertions involving first-personal pronouns. Fortunately, the Kaplan horizontal is not the only speech act content definable from the semantics. As Lewis (1980) showed, the intensional semantic theory does not determine one unique candidate for the role of speech act content. To be sure, it is convenient if the content that is the input to intensional operators – i.e., here the Kaplan horizontal – is also the content expressed in communication. But speech act content *need not* be identical to the content that combines with intensional operators to yield a sentence’s semantic value in context.³⁵ It is this freedom that allows us to define the right interesting sequenced worlds content from the semantic value of sentences containing first-personal pronouns.

Suppose Ben utters ‘I am hungry’ in a conversation with Anna, where the conversational sequence is $\langle \text{Ben}, \text{Anna} \rangle$. As we saw above, the interesting sequenced worlds content he communicates is $HUNGRY_I$, repeated in (4.24).

$$(4.24) HUNGRY_I: \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : x_1 \text{ is hungry in } w \text{ at } t \}.$$

In this particular case, the speech act content $HUNGRY_I$ is the Kaplan diagonal of the sentence ‘I am hungry’ as asserted by Ben. The Kaplan diagonal of a sentence Φ is the set of contexts at which Φ is true.

$$\begin{aligned} \text{Kaplan diagonal of } \Phi: \lambda c. \llbracket \Phi \rrbracket^{c,c} = \\ \{ c: \llbracket \Phi \rrbracket^{c,c} = \mathbf{I} \} = \\ \{ \langle w, t, \langle x_1, \dots, x_n \rangle \rangle : \llbracket \Phi \rrbracket^{\langle w, t, \langle x_1, \dots, x_n \rangle \rangle, \langle w, t, \langle x_1, \dots, x_n \rangle \rangle} = \mathbf{I} \} \end{aligned}$$

Recall that a context c is modelled by a sequenced world. So the Kaplan diagonal is a sequenced worlds content. Recall also that when a sequenced world represents a context, the first individual of the sequence, x_1 , represents the speaker of the context. ‘I am hungry’ is true at all contexts in which the speaker is hungry, which is just the set of contexts $\langle w, t, \langle x_1, \dots, x_n \rangle \rangle$ such that x_1 is hungry in w at t . So the Kaplan diagonal of Ben’s assertion in a conversation with one addressee is $HUNGRY_I$, as required.

The Kaplan diagonal is the right speech act content of first-personal assertions only in the case where the speaker happens to be the first individual in the conversational

³⁵Ninan (2010a) offers a clear and illuminating exposition of Lewis’ point, as well as the various options for defining speech act content in a Kaplan-style intensional semantics.

sequence. The Kaplan diagonal is a set of contexts, and the sequenced worlds representing contexts have by convention always the same individual-slot represent the speaker of the context – we have picked the first slot, x_1 . So given the conversational sequence $\langle \text{Ben}, \text{Anna} \rangle$, Ben’s assertion of ‘I am hungry’ will be accepted by Anna in case she thinks Ben, the first member of the conversational sequence represented in the context by x_1 , is hungry. But given that ‘I’ always picks out the speaker of the context, which we fixed as x_1 , the Kaplan diagonal of Anna’s assertion of ‘I am hungry’ is also the set of sequenced worlds such that x_1 is hungry in w at t , i.e. $HUNGRY_t$. Yet for Ben and Anna to presuppose $HUNGRY_t$, relative to the conversational sequence $\langle \text{Ben}, \text{Anna} \rangle$, is to accept that Ben is hungry. And that is clearly not the effect of Anna’s assertion of ‘I am hungry.’

What Anna in fact says is that the center representing her, relative to the conversational sequence, is hungry: $\{ \langle w, t, \langle x_1, x_2 \rangle \rangle : x_2 \text{ is hungry in } w \text{ at } t \}$. We can get this content from the Kaplan diagonal by simply swapping around the centers. So in conversations with just two interlocutors, speech act content is the Kaplan diagonal in case the speaker is the person that comes first in the conversational sequence, and it is the Kaplan diagonal with the centers swapped around in case the speaker comes second in the conversational sequence.

Generalising this rule for conversations with more than two interlocutors, it is helpful to introduce the notion of an m -inverse:³⁶

For any sequenced worlds contents p and q , q is the m -inverse of p iff for all sequenced worlds $\langle w, t, \langle x_1, \dots, x_n \rangle \rangle$ and $m \leq n$: $\langle w, t, \langle x_1, \dots, x_m, \dots, x_n \rangle \rangle \in q$ iff $\langle w, t, \langle x_m, \dots, x_1, \dots, x_n \rangle \rangle \in p$.

Thus, if the person that comes second in the conversational sequence makes an assertion of a sentence Φ , the speech act content is the 2-inverse of the Kaplan diagonal of Φ , as above. In general, an assertion of Φ by the m^{th} member of the conversational sequence expresses the m -inverse of the Kaplan diagonal of Φ – that is, the Kaplan diagonal with x_1 and x_m swapped around.³⁷ We can thus give the following definition of speech act content:

SPEECH ACT CONTENT_{SW}

In a conversational context $\langle w_c, t_c, \langle x_1, \dots, x_m, \dots, x_n \rangle \rangle$ the content of an utterance of Φ by x_m is the m -inverse of the Kaplan diagonal of Φ .

³⁶The notion of an m -inverse is a generalisation of Ninan’s notion of inverse, which he defines for sequenced worlds content with only two centers in the sequence.

³⁷In the special case where the speaker comes first in the conversational sequence – as Ben above – the speech act content is the 1-inverse of the Kaplan diagonal. Since swapping x_1 around with itself is vacuous, the 1-inverse of the Kaplan diagonal just is the Kaplan diagonal, as desired.

SPEECH ACT CONTENT_{SW} also yields the desired content for first-personal taste claims. Since such taste claims do not put any condition on one center of the sequence that they do not put on every other, swapping centers around is vacuous. And since PPTs in first-personal uses do not introduce dependence on context of the intension of the whole sentence, SPEECH ACT CONTENT_{SW} is in fact equivalent to the (*m*-inverse of the) Kaplan horizontal of the sentence. The speech act content_{SW} of Ben's assertion of 'Liquorice is tasty' in a conversation with Anna is (4.25):

$$\begin{aligned}
 (4.25) \quad & \lambda \langle w, t, \langle x_1, x_2 \rangle \rangle. \llbracket \text{Liquorice is tasty PRO}_C \rrbracket^{\langle w, t, \langle x_1, x_2 \rangle \rangle, \langle w, t, \langle x_1, x_2 \rangle \rangle} \\
 & = \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : \llbracket \text{Liquorice is tasty PRO}_C \rrbracket^{\langle w, t, \langle x_1, x_2 \rangle \rangle, \langle w, t, \langle x_1, x_2 \rangle \rangle} = \mathbf{1} \} \\
 & = \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : \text{liquorice tastes good to } \langle x_1, x_2 \rangle \text{ in } w \text{ at } t \}
 \end{aligned}$$

Two comments on SPEECH ACT CONTENT_{SW} are in order. First, it is important to realise that SPEECH ACT CONTENT_{SW}, as we have defined it, diagonalises over entire contexts. It thereby affects the interpretation of all context-sensitive expressions. So whenever we use, say, a demonstrative expression, we communicate contents in which the reference of this expression is not determinately resolved. For instance, in a conversation between two people, an assertion of 'This is Steve' expresses the speech act content_{SW} in (4.26):

$$(4.26) \quad \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : \text{the object demonstrated in } \langle w, t, \langle x_1, x_2 \rangle \rangle \text{ is Steve in } w \text{ at } t \}$$

One might object that it is implausible that an assertion of 'This is Steve' expresses the content in (4.26). For standard context-sensitive expressions other than 'I/you', one might want to stay true to one's philosophical commitments to Kaplan's direct reference view of 'this' and 'that.' It is worth noting that these commitments are not incompatible with the sequenced worlds model. Diagonalising in a way that affects all pronouns, demonstratives and other context-sensitive expressions alike is not the only option. Appendix 4.A at the end of this chapter outlines two ways of diagonalising that leave the reference of demonstratives and other context-sensitive expressions unchanged. Diagonalising over entire contexts is not forced on us.

There is, however, reason to think that standard demonstratives and other context-sensitive expressions make contributions to speech act content as in (4.26). The class of cases in which we need to diagonalise, in which horizontal content cannot be what is communicated, is wider than so far discussed. We often use demonstrative and other context-sensitive expressions when there is no common knowledge of all the features relevant to resolving the reference of these expressions in the context of use. Consider two cases from Perry (1979).

A hiker who is lost in the wilderness is looking at a lake and a mountain and does not know whether he is facing Gilmore Lake and Mt. Tallac or Clyde Lake and Jack's

peak. He is ignorant of the spatial location of his context. Suppose he is then told, ‘*This* is Gilmore Lake and *that* is Mt. Tallac’, and thus comes to know where he is. Prior to this assertion, the context set – the set of contexts that he and his interlocutor mutually accept as possibilities for their actual context – contains contexts that are located within sight of Gilmore Lake and Mt. Tallac and contexts located near Clyde Lake and Jack’s peak. So when the hiker is told, ‘*This* is Gilmore Lake and *that* is Mt. Tallac’, he is unable to derive the Kaplan horizontal, since for some contexts compatible with what he presupposes, ‘this’ and ‘that’ refer to Gilmore Lake and Mt. Tallac, and for some contexts, ‘this’ and ‘that’ refer to Clyde Lake and Jack’s peak. However, he can derive the diagonal content such that what the speaker demonstrates first is Gilmore Lake and what the speaker demonstrates second is Mt. Tallac. The diagonal will allow him to eliminate all the contexts from the context set that are located near Clyde Lake and Jack’s peak, and he will come to know that what he is looking at is Gilmore Lake and Mt. Tallac.

Similarly, consider Perry’s professor who knows that the departmental meeting starts at noon but does not know that, as he is thinking about whether or not it is time to go to the meeting, it is noon. So when he is told, ‘The meeting starts *now*,’ he won’t be able to retrieve the Kaplan horizontal, since due to his ignorance the context set contains contexts in which it is noon and contexts in which it is some time other than noon. However, he is in a position to derive the diagonal such that the meeting starts at the time of speaking, which allows him to eliminate from the context set all contexts other than those in which it is noon.

There are thus cases of *de hic* and *de nunc* communication in which diagonal content is the only plausible communicated information.

But is it plausible to assume that diagonal content is *always* the content expressed in speech, even in cases of perfectly smooth communication, where there is no relevant contextual ignorance? This concerns the second important comment: $\text{SPEECH ACT CONTENT}_{\text{SW}}$ entails that the content expressed in communication is *always* (the *m*-inverse of) diagonal content. Note that while diagonal content is the only plausible candidate for speech act content in cases in which there is ignorance among interlocutors as to the features of context relevant to resolving the reference of demonstratives and other context-sensitive expressions, it is harmless to posit diagonal content in cases of smooth communication, in which there is no such ignorance among interlocutors. In these cases, a theory on which horizontal content is communicated and our theory on which diagonal content is communicated make exactly the same predictions. This is best seen by considering a simple example.³⁸ Suppose James says to Liz, ‘Sam isn’t here’ at a point in the conversation where it is mutual knowledge that the location of the conversation is the British Library, London. Then the context set will contain only

³⁸Heim (2004) provides a proof of this result.

contexts $\langle w, t, \langle x_1, x_2 \rangle \rangle$ in which x_1 and x_2 are located in the British Library in w at t . The Kaplan horizontal of James's assertion is

$$(4.27) \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : \text{Sam is not in the British Library in } w \text{ at } t \}.$$

The (\mathbf{I} -inverse of the) Kaplan diagonal of James's assertion is

$$(4.28) \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : \text{Sam is not in the location of } \langle w, t, \langle x_1, x_2 \rangle \rangle \text{ in } w \text{ at } t \}.$$

If Liz understands James's assertion and trusts him to be truthful, the effect of James's assertion will be an update of the context set by intersecting it with the assertion's content. Now, the update with (4.27) and the update with (4.28) have exactly the same effect, for the context set prior to the assertion contains only contexts located in the British Library. So intersection with (4.27) will eliminate all of the British-Library contexts in which Sam is in the British Library, and intersection with (4.28) will eliminate all of the British-Library contexts in which Sam is in the British-Library context. In both cases, the post-assertion context set will only contain British-Library contexts in which Sam is not in the British Library.

In conclusion, diagonalisation is needed for various cases of *de se*, *de nunc*, and *de hic* ignorance, and theories that posit horizontal content make exactly the same empirical predictions as our theory on which communicated content is always diagonal content for conversations in which there is no such ignorance. Given this equivalence, we should prefer the simpler theory that makes diagonalisation the general rule: $\text{SPEECH ACT CONTENT}_{\text{SW}}$.

4.10 Competing Solutions

The sequenced worlds model shows that the conflict between the Lockean picture of communication and the self-locating belief model can be resolved without giving up either. But there are other accounts that solve the conflict by giving up either the Lockean picture or the self-locating belief model. I will here discuss some of these accounts and point out some of their respective disadvantages. With the notable exception of Egan (2007), none of these views provide a unified account of the communication of both *de se* and subjective attitudes. I will suggest ways to extend the existing accounts of the communication of *de se* attitudes to that of subjective attitudes. Some of my criticisms will concern the prospects of such extensions. Before I begin, let me repeat that both the self-location account and the Lockean picture are very attractive accounts of mental attitudes and linguistic communication, respectively, so saving the core of both of them should be seen as a major advantage of the sequenced worlds over all of these competitors.

4.10.1 Centered Belief Content without Lockean Communication

4.10.1.1 Boring Speech Act Content

Egan (2007, 2009, 2010a) solves the problem of the conflict by denying that in *de se* communication, interesting centered content is transmitted along Lockean lines. At the level of thought, both *de se* and taste attitudes have interesting centered content. But only the contents of taste beliefs get communicated directly, *de se* contents do not. Assertions about oneself express *boring* centered contents – equivalently, possible worlds contents.³⁹ First- and second-personal singular pronouns have their well-known Kaplanian semantics, on which they contribute the speaker, or the addressee, to the content expressed. So when Ben at noon thinks *de se* that he is hungry, and says, ‘I am hungry,’ he expresses the boring centered content $\{\langle w, t, x \rangle: \text{Ben is hungry in } w \text{ at noon}\}$. Egan endorses a common ground model of conversation on which the context set is a set of centered worlds. Successful taste assertions update with interesting centered contents, *de se* claims with boring centered contents.⁴⁰

To account for the fact that *de se* communication can result in the acquisition of *de te* or *de se* beliefs on the hearer’s part, Egan must complement this account. As it stands, an addressee of Ben’s assertion of ‘I am hungry’ only acquires the boring centered content that Ben is hungry. If she is to gain a piece of *de se* information, she must *infer* it from the perfectly ordinary information about the world she receives. On the completed account, then, speech is not the direct expression of *de se* thought, and what is said is not what hearers come to believe when they acquire a *de se* belief. SPEAKER’S MENTAL CONTENT, SPEECH ACT CONTENT, and HEARER’S MENTAL CONTENT are not the same, and explanations of BELIEF-SPEECH COORDINATION and SPEECH-BELIEF COORDINATION will turn out to be rather complex. All core tenets of the Lockean picture are abandoned.⁴¹

The view needs to be further complicated to account for many cases of *de se* communication. What does Anna learn when Ben introduces himself by saying ‘I am Ben?’ Not the necessary but uninformative proposition *that Ben is Ben*, which is the boring centered content the Kaplanian semantics delivers. Egan (2010a, 280) suggests that hearers come to acquire the diagonal content of the sentence by exploiting the pragmatic mechanisms described in Stalnaker (1978). On Egan’s account, the diagonal of Ben’s assertion is something like the interesting content $\{\langle w, t, x \rangle: \text{the speaker in}$

³⁹A boring centered worlds content p does not distinguish between locations in a world: p is boring iff for every world w and pairs $\langle t_1, x \rangle, \langle t_2, y \rangle$ such that individual x inhabits w at time t_1 and y inhabits w at t_2 , p contains $\langle w, t_1, x \rangle$ iff it contains $\langle w, t_2, y \rangle$. See section 4.2.1.

⁴⁰Moss (forthcoming) endorses a very similar account of *de se* communication, without taking a stance on discourse about matters of taste. Stephenson (2007a) endorses a very similar account of discourse about taste, without taking a stance on *de se* communication.

⁴¹Stalnaker (1981, 146-7) was the first to sketch this account of *de se* communication for the Lewisian and criticised it on the grounds that it must deny that speech is the straightforward expression of thought.

$\langle w, t, x \rangle$ is Ben in w at t , from which Anna infers the *de se* content such that the center is being addressed by Ben.

There are also some worries about the conception of common ground in terms of centered worlds. If conversation is the joint project of carving out possibilities, whose possibilities does the context set represent? Ways the speaker might be, or ways the hearer might be? One possibility, suggested by Egan (p.c.), is that the context set represents the speaker's possibilities relative to the speaker and the hearer's possibilities relative to the hearer. But this leaves it unclear what we are agreeing on in disputes about taste, if the speaker proposes an update of *her* possibilities and the hearer accepts an update of *his* possibilities.

Alternatively, the context set could represent the possibilities of the conversational group, possibilities all members of the group share. This is Stephenson's proposal (2007a, 509). But on this view, centered worlds play two distinct roles in private attitudes like belief and the public attitude of presupposition towards the common ground: representing possibilities for the individual attitude holder, and representing possibilities for the conversational group. While private belief is still *self*-location, the base attitude in presupposition (belief or assumption according to Stalnaker) must be *group*-location. The view needs to explain how speakers come to make the step from self- to group-location if they are to propose or accept changes to the common ground purely on the basis of their private beliefs.

4.10.1.2 Speaker-Centric Speech Act Content and Recentering

Another way to solve the *de se* problem, endorsed by Heim (2004) and Weber (2011), is to hold on to the expression of *de se* thought in speech but deny that hearers come to believe what is said. That is, SPEECH ACT CONTENT is identical to SPEAKER'S MENTAL CONTENT but NOT to HEARER'S MENTAL CONTENT. The picture is this. When Ben says 'I am hungry,' the asserted content is the interesting centered content of his belief: $\{\langle w, t, x \rangle: x \text{ is hungry in } w \text{ at } t\}$. The hearer arrives at the content she ends up believing by *recentering*: From the speech act content and her beliefs about how she is related to the speaker, she *infers* the *de se* content $\{\langle w, t, x \rangle: x \text{ is addressed by someone hungry in } w \text{ at } t\}$.⁴² In general, recentering is an operation that takes a centered content p and outputs a centered content $q = \{\langle w, t, x \rangle: x \text{ bears relation } R \text{ uniquely to } y, \text{ and } p(w, t, y)\}$, where R is a relation which is determined by the context of utterance and by which the addressee identifies the speaker.⁴³

⁴²This account, too, was first mentioned by Stalnaker (1981, 146-7), who criticised it for the indirectness it imputes on communication.

⁴³The recentering operation, as a result of which hearers comes to *self*-ascribe a property q where the speaker self-ascribed p , is essentially what for Lewis amounts to the ascription of a property to others (Lewis, 1979a, 539).

Recentering makes evident that the account also gives up the Lockean picture: The speaker believes and asserts one *de se* content, hearers infer another. And because hearers may be related to the speaker in different ways – e.g., only one of them can be the addressee of the singular pronoun ‘you’ – each hearer may infer a different *de se* content.

One drawback of the view is that the Stalnakerian model of common ground must be given up because there is no single content that speaker and audience come to presuppose as a result of successful communication. For the same reason, the account of disagreement in conversation has to be more complicated than the simple story available to other options, on which two speakers disagree in case there is a single content that one accepts and the other rejects.

More importantly, the view, designed to explain *de se* communication, cannot solve the incompatibility problem. Applied to taste communication, it makes the wrong predictions. If the speaker, in saying ‘This cookie is tasty,’ expresses her belief with the content

(4.2) *COOKIE*: $\{\langle w, t, x \rangle: \text{the (salient) cookie tastes good to } x \text{ in } w \text{ at } t\}$,

the hearer will, by recentering, infer something like the content $\{\langle w, t, x \rangle: x \text{ is addressed by someone to whom the cookie tastes good in } w \text{ at } t\}$. Surely, that is not what the hearer comes to believe. On the contrary, the hearer comes to believe, simply, *COOKIE*.

One may think that the right account of taste communication can be combined with the recentering view by declaring a systematic difference between taste and *de se* communication: In recentering, hearers use different relations – something like *being addressed by* for *de se* assertions and the identity relation for taste assertions. This proposal, however, cannot account for mixed assertions, such as ‘My cookies are tasty’ and ‘My pants are on fire and that’s not fun,’ which on the proposal require hearers to recenter in inconsistent ways.

4.10.2 Lockean Communication without Centered Belief Content

Perry (1977, 1979) argued that *de se* attitudes are attitudes to regular propositions and that their distinctive cognitive significance is due to the first-personal mode of presentation (so-called ‘belief states’) by which propositions are entertained. Simplifying somewhat, when Ben believes *de se* that he is hungry, he believes, in a first-personal way, the proposition that Ben is hungry. Ben communicates his *de se* belief by asserting ‘I am hungry,’ which given a Kaplanian semantics expresses the proposition that Ben is hungry. If communication is successful, the audience comes to believe that very proposition. The Lockean picture is thus preserved. Speaker and audience believe the same proposition, albeit under different modes.

Perry's account of *de se* attitudes may be combined with elements of Moltmann's (2010) account of thought and talk about taste. On her account, assertions about matters of taste express contents with absolute truth conditions but to which our cognitive access is essentially first-personal. Very roughly, 'This cookie is tasty' has the logical form $Gn x \text{ this cookie is tasty to } x$, where Gn is a generic quantifier whose domain is contextually restricted; an assertion of this sentence expresses, roughly, that the cookie tastes good to the individuals in the domain. We grasp this content in a first-personal way by taking the domain to contain only individuals with which we identify. This identification marks the special first-personal mode of presentation with which taste contents are entertained.

Perry's account may seem simpler than the sequenced worlds model and therefore preferable. As it stands, however, it avoids the conflict by ignoring its cause. The fact that what is a *de se* belief for the speaker becomes a *de te* belief for the hearer, and vice versa. For an utterance of 'I am hungry,' Perry's account needs to explain the systematic coordination of speaker's first-personal mode of presentation, the words used to express the thought, and the hearer's second-personal mode of presentation. On one way of understanding modes of presentation, they are associated with functions from contexts (a believer at a time and world) to propositions (such as *that Ben is hungry*); they are the mental equivalents of the Kaplanian *characters* of the sentences used to express the beliefs. On another way of understanding modes of presentation, they are associated with functions from contexts to contents true at the contexts; if contexts are understood as centered worlds, the set of contexts that get mapped to true contents, for a given mode of presentation, is the set of centered worlds the Lewisian takes to be the object of the attitude. Whatever the exact understanding of modes of presentation, a complete mode of presentation account is not simpler than the sequenced worlds model, it just locates *de se* aspects in a different place.

There are also some drawbacks to locating the subjective dimension of taste judgments in the modes of presentation. Moltmann's account, for instance, needs to explain why we would systematically entertain contents about taste in a way that leads many to erroneous judgments. For Moltmann's contents have absolute truth conditions – their truth depends on the tastes of the contextually selected domain of people. Nevertheless, our taste judgments differ vastly and statistical data on people's tastes usually do not help settle disputes.

4.11 Objections

4.11.1 Belief and Sequenced Worlds Content

Objection. The notion of belief in sequenced worlds content is one on which thinkers do not just locate themselves but the entire conversational group of which they are a

member. But clearly, this is not a plausible notion of belief outside of conversational contexts.

Response. In solitary thinking and soliloquy, in which the thinker is not addressing anyone in a second-personal way, we can think of belief as locating the ‘group’ that consists just of the thinker herself. That is, her belief content is a set of sequenced worlds whose sequences have a single center – i.e., it is a centered worlds content. So in solitary contexts, belief naturally amounts to self-location.

4.11.2 On Believing in Context

Objection. The previous response entails that belief in a conversational context and belief in a solitary context come apart. For instance, Ben’s belief *that I am hungry*, prior to a conversation with Anna, has centered worlds content p : $\{\langle w, t, x \rangle: x \text{ is hungry in } w \text{ at } t\}$. But p is not the content he communicates to Anna on the sequenced worlds model when he says ‘I am hungry.’ Instead he communicates q , given the conversational sequence $\langle \text{Ben}, \text{Anna} \rangle$: $q = \{\langle w, t, \langle x_1, x_2 \rangle \rangle: x_1 \text{ is hungry in } w \text{ at } t\}$. To see that $p \neq q$, note that by the equivalence established in SEQUENCED WORLDS BELIEF AND CENTERED WORLDS BELIEF, $q = q'$:

$$q' = \{\langle w, t, x \rangle: \text{there are individuals } x_1, x_2 \text{ such that } x = x_1, x \text{ is addressing } x_2, \text{ and } x_1 \text{ is hungry in } w \text{ at } t\}$$

If we uphold the Lockean picture on which the same content plays the role of SPEECH ACT CONTENT and MENTAL CONTENT, then Ben is not communicating the content of the belief he has prior to conversation. Instead, he has to subtly modify the beliefs he brings to the conversation if he wants to communicate them. But that is implausible.

Response: It is true that on the sequenced worlds model, Ben’s belief contents undergo a subtle transformation when he enters a conversational context. I cannot here offer a complete theory of belief update over time. But one thing can be said: *Given the beliefs Ben acquires by entering a conversation*, his belief in p is equivalent to his belief in q' . Ben’s overall belief state in w at t , $\text{Dox}_{\langle w, t, \text{Ben} \rangle}$, can be modelled as the set of centered worlds compatible with what he believes in w at t . If Ben in w_1 and at some time t_1 prior to conversation believes *de se* that he is hungry, his belief state $\text{Dox}_{\langle w_1, t_1, \text{Ben} \rangle}$ contains only centered worlds that are elements of p . If at t_2 , Ben enters a conversation with Anna by telling her, ‘I am hungry,’ he acquires the *de se* belief *that I am addressing someone* – call its content r . So $\text{Dox}_{\langle w_1, t_2, \text{Ben} \rangle}$ contains only centered worlds which are members of r . But given $\text{Dox}_{\langle w_1, t_2, \text{Ben} \rangle}$, the belief in p and the belief in q' are equivalent in the following sense: updating $\text{Dox}_{\langle w_1, t_2, \text{Ben} \rangle}$ with p would have the same result as updating with q' . More generally, for any doxastic state that contains only r -centered worlds, its intersection with p equals its intersection with q' . The

only difference between p and q' is that the latter requires that the center is addressing someone. Since $\text{Dox}_{\langle w_1, t_2, \text{Ben} \rangle}$ already contains only centered worlds in which the center is addressing someone, updating with p as well as updating with q' will have the effect of eliminating all centered worlds except those in which the center is hungry. Hence there is nothing mysterious about the contents of Ben's beliefs changing when he enters a conversation. The cognitive significance they have for him is the same in the conversational context.

4.11.3 Objecting on Someone Else's Behalf

Objection. On the sequenced worlds view, it seems that I can object to updating the common ground with the content of a bare taste claim on someone else's behalf. For instance, on the view Bill's rejoinder in the following conversation between Ben, Sal and Jim should be a natural conversational move.

- (4.29) Ben: This liquorice is tasty.
 Sal: ? I disagree/No. Jim doesn't like it.

Response. Sal's rejoinder is odd for two reasons: (i) It is generally odd to speak for people who are present in the conversation when it comes to value judgments. There is some oddity in Sal's reply even in the following conversation between Ben, Sal, and Jim.

- (4.30) Ben: Jim finds this liquorice tasty.
 Sal: ? I disagree/No. Jim finds it disgusting.

(ii) It is odd to speak on other people's behalf when they are present because the epistemic access we have to their tastes is much less immediate than their own access. Note that it is much less odd to object on behalf of someone who, for whatever reason, is less well positioned to express their taste judgments. For instance, it is not so odd to object on a child's behalf – especially when agreement or disagreement as a group matters for practical purposes.

4.11.4 Sameness of Content

Objection. On the sequenced worlds model, assertions of the same sentence by the same person in different conversations may express different contents where intuitively, they are the same. If Ben says 'I am hungry' in a conversation with one addressee, it will express a different content than an assertion of the same sentence by Ben in a conversation with two addressees. The first will express a pair-centered content, the second a content whose worlds are centered on a sequence with three individuals. But it seems that what he said is the same on both occasions.

Response. It is true that sequenced worlds speech act content depends on facts about the context, namely the number of participants in the conversation. As a result, assertions of the same sentence by the same speaker may express different sets of sequenced worlds in contexts that differ only with respect to the number of interlocutors. But while intuitions about ‘what is said’ are themselves highly context-dependent and can track different aspects of meaning,⁴⁴ there is one aspect of meaning which is stable across the different contexts of Ben’s assertion, and which is closely related to the content represented as sets of sequenced worlds. To see this, let us take a step back. A sequenced worlds content, just like a possible worlds content, can be understood as a way of distinguishing between alternative possibilities. Possible worlds content is a way of distinguishing between ways the world might be. Sequenced worlds content is a way of distinguishing between ways a number of people might be. Both types of content are partitions of a space of possibilities. For sequenced worlds content, that space of possibilities itself may differ in kind from context to context, depending on the number of people whose possibilities are represented. It can thus yield different sets of sequenced worlds – sets of worlds with sequences of different length. If we understand content as a way of distinguishing between group possibilities, then the content of Ben’s assertion is stable across contexts. Technically, this stable aspect of meaning is a function from a conversational context that includes an ordered n -tuple of participants to a set of sequenced worlds. All of Ben’s assertions have in common the function that takes a conversational sequence (of any length) to a set of sequenced worlds such that the center representing the speaker is hungry in the world at the time.

4.11.5 ‘Tasty to us’

Objection: On the sequenced worlds view of speech act content, an assertion of ‘Liquorice is tasty to us’ expresses the same speech act content_{SW} as an assertion of ‘Liquorice is tasty.’ In a conversation between two people, they both express (4.31).

(4.31) $\{\langle w, t, \langle x_1, x_2 \rangle \rangle: \text{liquorice tastes good to } \langle x_1, x_2 \rangle \text{ in } w \text{ at } t\}$

But an assertion of ‘Liquorice is tasty to us’ does not have the same communicative effects as an assertion of ‘Liquorice is tasty.’ For instance, in a conversational context in which it is common belief that liquorice fails to taste good to at least one of the addressees, the assertion ‘Liquorice is tasty to us’ seems infelicitous. But a speaker to whom liquorice tastes good can still felicitously assert ‘Liquorice is tasty’ in that context. So it seems that the two assertions have different truth conditions.

⁴⁴Among the vast literature, see for instance Lewis (1980), Cappelen and Hawthorne (2009, ch.2) and Stojanovic (2012).

Response. I will suggest that in contexts in which ‘us’ picks out the conversational group, assertions of ‘Liquorice is tasty’ and ‘Liquorice is tasty to us’ do in fact have the same assertability conditions, contrary to appearance. To explain the perceived difference in appropriate assertability in the objection’s example, however, I will first have to say much more about norms of assertion and the pragmatics of disputes about taste. The full response will therefore have to wait until section 5.6.

4.11.6 Negated Taste Claims

Objection. The semantics for PPTs, together with the definition of *SPEECH ACT CONTENTS_{SW}* on the sequenced worlds model of communication, makes the wrong predictions for negated taste claims. Consider the following assertion by Ben:

- (4.32) a. This liquorice is not tasty.
 b. [[This liquorice] [is not tasty *PRO_C*]]

In the conversational context $\langle w_c, t_c, \langle \text{Ben}, \text{Anna} \rangle \rangle$, Ben’s assertion expresses the content *LIQUORICE*:

- (4.33) *LIQUORICE*: $\{ \langle w, t, \langle x_1, x_2 \rangle \rangle : \text{the liquorice does not taste good to } \langle x_1, x_2 \rangle \text{ in } w \text{ at } t \}$

LIQUORICE is true at a triple $\langle w_1, t_1, \langle y_1, y_2 \rangle \rangle$ just in case it is not the case that the liquorice is tasty to *both* y_1 and y_2 in w_1 at t_1 . Consider the situation in which Ben does not like the liquorice, but Anna does. In this situation, *LIQUORICE* is true, and thus Anna should accept Ben’s proposal to update with *LIQUORICE*. But intuitively, Anna should not accept Ben’s assertion, since she does like liquorice. It seems that Ben and Anna should agree on (4.32a) only if the liquorice is not tasty to *either* of them. So *LIQUORICE* cannot be what Ben expresses in (4.32a), contrary to what is predicted on the sequenced worlds view.⁴⁵

Response. In fact, the view does not predict that (4.32a) always expresses *LIQUORICE*. (4.32a) is ambiguous and has two readings. On the first, the negation takes wide scope over *PRO_C*, as in (4.32b). On this wide scope reading, (4.32a) expresses *LIQUORICE*. On the second reading, negation takes narrow scope. This reading is obtained when *PRO_C* is moved to the front of the sentence:

⁴⁵This objection bears some similarity to an objection Lasnik (2005, 650-2) raises against a form of indexical contextualism that says that an assertion of ‘Liquorice is not tasty,’ at least in some contexts, expresses the content that liquorice is not tasty to the group that includes speaker and hearer. Indexical contextualists may avail themselves of the reply I give below depending on the syntactic and semantic details of their view.

(4.34) PRO_C this liquorice is not tasty

On this narrow scope reading, (4.32a) expresses *LIQUORICE'*.

(4.35) *LIQUORICE'*: $\{\langle w, t, \langle x_1, x_2 \rangle \rangle: \text{to } \langle x_1, x_2 \rangle, \text{ the liquorice does not taste good in } w \text{ at } t\}$

LIQUORICE' is true at a triple $\langle w_1, t_1, \langle y_1, y_2 \rangle \rangle$ just in case the liquorice tastes good to *neither* y_1 nor y_2 in w_1 at t_1 . Thus *LIQUORICE'* is not true in the situation in which Ben does not like the liquorice but Anna does. On the narrow scope reading, the sequenced worlds model correctly predicts that Anna should not accept Ben's assertion (4.32a).

It is plausible to think that PRO_C can be moved to the front of the sentence, just like overt experiencer arguments can be moved to the front:

(4.36) To Ben and Sal, this liquorice is tasty.

(4.37) To Ben and Sal, this liquorice is not tasty.

(4.38) To everyone in their right mind, this liquorice is not tasty.

For (4.37) and (4.38), the default readings are narrow scope readings. For instance, (4.38) is naturally understood as true just in case for every x in their right mind, the liquorice does not taste good to x . So it is plausible to assume that (4.34), where PRO_C is moved to the front, expresses *LIQUORICE'*.

Finally, narrow scope readings are also the strongly preferred readings when the overt experiencer argument has *not* been moved to the front. Consider (4.39):

(4.39) ? This liquorice isn't tasty to Ben and Sal, but it is tasty to Ben.

(4.39) is clearly odd at first sight. But if the first conjunct, 'This liquorice isn't tasty to Ben and Sal,' is read with negation taking wide scope, the whole sentence may still be true – in a situation where the liquorice tastes good to Ben but not to Sal. However, the oddity is easily explained on the reading of the first conjunct on which negation takes narrow scope. In this case, the second conjunct, 'it is tasty to Ben,' is a flat contradiction of the first conjunct.

In sum, it is plausible to assume that for Ben's assertion of (4.32a), 'This liquorice is not tasty,' the default reading is the narrow scope reading, for which the sequenced worlds model predicts the speech act content *LIQUORICE'* and makes the right predictions about the acceptance conditions for the hearer.

4.12 Conclusion

I have argued that the conflict between a Lewisian view of belief as self-location and the received Lockean picture of communication can be resolved by conceiving of the contents of mental attitudes and speech acts as sets of sequenced worlds – possible worlds ‘centered’ on a sequence of individuals at a time. Sequenced worlds content is the kind of centered information that is transferred from speaker’s head to hearer’s head in successful communication. Communication, on the sequenced worlds view, is the project of distinguishing between possible ways the group of interlocutors might be and involves the coordination of participants’ individual perspectives. The point of assertions about matters of taste is to reach a joint perspective. The aim of *de se* assertions is to establish the speaker’s individual possibilities.

There are other solutions of the conflict. But all of them must give up either the self-locating account of belief or the Lockean picture of communication. The major theoretical benefit of the sequenced worlds view is that it provides a unified account of thought and communication which preserves the insights of both the received and simple Lockean picture of communication and the self-locating account of belief.

4.A Diagonalisation and Context-Sensitivity

In section 4.9, speech act content is the result of (taking the m -inverse of the result of) diagonalising the sentence's semantic value: $\lambda c. \llbracket \Phi \rrbracket^{c,c}$. As a consequence, not only the contribution of first- and second-person pronouns to speech act content is affected, but also the contribution of *any* expression whose extension varies with the context of use (e.g., demonstratives and indexicals). In sections 4.8 and 4.9, we assumed Kaplanian entries for all of the following expressions.

Given a context of use $c = \langle w, t, \langle x_1, x_2 \rangle \rangle$ that determines a unique location l_c of c ,

$$(4.40) \llbracket \text{I} \rrbracket^{c,i} = x_1$$

$$(4.41) \llbracket \text{you} \rrbracket^{c,i} = x_2$$

$$(4.42) \llbracket \text{here} \rrbracket^{c,i} = l_c$$

$$(4.43) \llbracket \text{now} \rrbracket^{c,i} = t$$

$$(4.44) \llbracket \text{this} \rrbracket^{c,i} = \text{the object demonstrated in } c$$

We can, however, give different entries for these expressions so that diagonalisation affects only the contribution of first- and second-person pronouns but not other context-sensitive expressions. There are at least two ways to do this. To keep it simple, I will exemplify the difference for the entries for 'I' and 'this.'

1. Let a context of use be $c' = \langle w, t, \langle x_1, x_2 \rangle, s \rangle$, where s is the object salient in c' .

$$(4.45) \llbracket \text{I} \rrbracket^{c',i} = x_1$$

$$(4.46) \llbracket \text{this} \rrbracket^{c',i} = s_c$$

(4.47) Speech act content of a sentence Φ at a context of use c :

$$\lambda \langle w, t, \langle x_1, x_2 \rangle \rangle. \llbracket \Phi \rrbracket^{\langle w, t, \langle x_1, x_2 \rangle, s_c \rangle, \langle w, t, \langle x_1, x_2 \rangle \rangle}$$

(4.47) λ -abstracts not over all features of the context of use but only over $\langle w, t, \langle x_1, x_2 \rangle \rangle$, and Φ is evaluated with respect to a context including s_c , the object that is salient *at the context of use* c . Since 'this' takes as its referent s , the diagonal speech act content of an assertion of 'I like this' at c is (4.48):

$$(4.48) \{ \langle w, t, \langle x_1, x_2 \rangle \rangle: x_1 \text{ likes } s_c \text{ in } w \text{ at } t \}$$

2. Let a context of use be $c = \langle w, t, \langle x_1, x_2 \rangle \rangle$. The interpretation function maps an expression ϕ , a context c , an index i , and an assignment function g to extensions. g is a contextually determined function that takes numerical indices to objects. The demonstrative 'this' is encoded at LF with a numerical index, e.g., *this*₅.

$$(4.49) \llbracket \mathbf{I} \rrbracket^{c,i,g} = x_1$$

$$(4.50) \llbracket \text{this}_5 \rrbracket^{c,i,g} = g(5)$$

(4.51) Speech act content of a sentence Φ at a context of use c :

$$\lambda c'. \llbracket \Phi \rrbracket^{c',c',g_c}$$

(4.51) λ -abstracts over the entire context. But the semantic value of ‘this₅’ does not depend on the context that is being abstracted over. Instead, it gets its value from g , which is determined by the context of use c . So the speech act content of an assertion of ‘I like this’ at c is (4.52):

$$(4.52) \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : x_1 \text{ likes } g_c(5) \text{ in } w \text{ at } t \}$$

SEQUENCED WORLDS, EAVESDROPPING, RETRACTION

5.1 Introduction

It is popular among linguists and philosophers with relativist leanings to understand content, in particular the content of beliefs about subjective matters such as personal taste and of sentences expressing such beliefs, as sets of centered worlds – possible worlds ‘centered’ on an individual inhabiting the world at some time.¹ Many philosophers and linguists also believe that *de se* attitudes, attitudes about oneself, have centered worlds content. In chapter 4 we saw that centered worlds content is in conflict with the received Lockean picture of communication, on which communication involves the transmission of information from speaker’s head to hearer’s head. Chapter 4 developed a solution to the conflict in terms of sequenced worlds content – sets of possible worlds ‘centered’ on a group of individuals – which reconciles the Lockean picture of communication with key aspects of the self-location account of attitudes underlying the centered worlds approach. On the emerging picture of communication, to engage in conversation is to distinguish between alternative ways that the conversational participants might be, where this does not require that all share the ways they *individually* might be.

Chapter 4 also outlined a compositional semantics for predicates of personal taste and first- and second-personal pronouns that draws upon Stephenson (2007a) and Kaplan (1989) in key respects. But this semantics and the sequenced worlds model leave many questions open. Crucially, the model as developed does not yet make predictions about the data that has been of central concern in debates between relativists and contextualists about predicates of personal taste. So what are the views’ predictions about eavesdropping, retraction, and disagreement? Is the fully developed sequenced worlds view a form of contextualism or relativism about predicates of personal taste?

Chapters 5 and 6 are devoted to answering these questions. I will begin by showing that the sequenced worlds model, as it stands, is neutral between versions of contextualism and relativism about assertions regarding matters of taste. More precisely, it is neutral between nonindexical contextualism and truth relativism. I will then look at empirical data from eavesdropping and retraction in chapter 5 and from disagree-

¹See for instance Egan et al. (2005), Egan (2007, 2010a), Lasersohn (2005) and Stephenson (2007a) for centered worlds versions of relativism.

ment cases and monadic truth ascriptions in chapter 6. The empirical data will help us decide between the contextualist version and the relativist version of the sequenced worlds model. To account for the pragmatic phenomena on the model, we will have to develop it further in crucial respects. While the model can be developed in ways that service either nonindexical contextualism or truth relativism, it exposes and highlights some unfortunate consequences of nonindexical contextualism. Chapters 5 and 6 thus contain an argument in favour of sequenced worlds relativism.

5.2 Contextualism and Relativism

The sequenced worlds model, with the semantics and the notion of SPEECH ACT CONTENT_{SW} given in chapter 4, can be completed in different ways to result in either a version of nonindexical contextualism or truth relativism regarding predicates of personal taste.

Nonindexical contextualism and truth relativism about predicates of personal of taste (PPTs) can each be characterised by three theses.² The views share the first two theses.

INVARIANT CONTENT

On its default first-personal use, a sentence involving a PPT expresses a content that is neutral with respect to the experiencer's perspective or standard of taste.³

INVARIANT CONTENT sets apart nonindexical contextualism and truth relativism from indexical contextualism. On a simple version of indexical contextualism about PPTs, an utterance of 'Liquorice is tasty' by Anna expresses the content *that liquorice is tasty to Anna*. On the contrary, nonindexical contextualism and truth relativism predict that her utterance expresses the content *that liquorice is tasty*, in which no value is assigned to

²The terms 'nonindexical contextualism' and 'truth relativism' are here used in MacFarlane's (2009; 2012) sense, although I introduce them without using MacFarlane's notions of indexicality, use-sensitivity and assessment-sensitivity.

³INVARIANT CONTENT does not exclude PPTs from inducing some variability of content with context. Many predicates of personal taste are gradable adjectives. For instance, an activity can be more fun than another activity; a food can be somewhat tasty, or very tasty. Nonindexical contextualists and relativists can allow sentences involving PPTs to express contents that are specific with respect to a degree element which can vary with context of use ('tasty to degree x'), as long as the degree element is determined independently of the perspective of the experiencer. See Kennedy (2000, 2007) for discussion of the degree analysis of gradable adjectives. Similarly, they can allow a comparison class to be determined by the context of use, such that 'tasty' expresses 'tasty for an x,' where x is a contextually salient comparison class (e.g., the class of all cookies), as long as this comparison class is determined independently of the perspective of the experiencer. See Richard (2008) for a relativist view on which a comparison class is determined by the context of use. See, e.g., Kamp (1975), McConnell-Ginet (1973), and Klein (1980) for the partial predicate analysis.

an experiencer argument. INVARIANT CONTENT also distinguishes nonindexical contextualism and truth relativism from content relativism (sometimes also called indexical relativism), according to which what an utterance of a sentence involving a PPT expresses is dependent on the interpreter's context. On this view, a single utterance of 'Liquorice is tasty' by Anna may express the content *that liquorice is tasty to Ben* relative to Ben's context of interpretation, and the content *that liquorice is tasty to Sal* relative to Sal's context of interpretation. In what follows, I will use 'relativism' to mean *truth relativism*, and the explicit 'content relativism' to refer to content relativism.⁴

SENSITIVITY

Propositional truth depends not just on the world (and time) of evaluation but also on the individual(s) or standard of taste relative to which the proposition is evaluated.

SENSITIVITY allows for the truth of contents expressed by first-personal uses of taste claims to vary with the taste preferences of the individual whose perspective is relevant to the evaluation of the claim.

Nonindexical contextualism and truth relativism differ with respect to the third thesis, the definition of truth. According to nonindexical contextualism, the truth of sentences in context depends solely on features of that context.

NONINDEXICAL CONTEXTUALIST TRUTH

A sentence Φ containing a PPT, as used in c , is true iff Φ is true at the context c and the index determined by c .⁵

According to truth relativism, the truth of sentences in context is relative to the perspective of an assessor.

RELATIVE TRUTH

A sentence Φ containing a PPT, as used in a context c , is true relative to an assessor iff Φ is true at the context c and the index determined (at least in part) by the assessor's perspective.

Nonindexical contextualism about PPTs is the view that encompasses INVARIANT CONTENT, SENSITIVITY and NONINDEXICAL CONTEXTUALIST TRUTH. According to the view,

⁴See Cappelen (2008b,a) for an exposition of the content relativist view and for a summary of content relativisms regarding conditionals (Weatherson, 2009), epistemic modals, and indexicals (Egan, 2009). Cappelen's main case for content relativism rests on predicates used in instructions, imperatives, and legal texts. See MacFarlane (2012, 106) to situate content relativism in MacFarlane's taxonomy of theories.

⁵Cf. Kaplan's definition of truth in context (1989, 522)

a sentence like ‘Liquorice is tasty,’ on a first-personal use, it will express the experiencer-neutral content *that liquorice is tasty*, whose truth depends on the perspective of the individual salient at the context of its use. Once the context of use is fixed, the truth of a sentence, or content, does not further vary with anyone’s perspective.

Truth relativism about PPTs is the view that encompasses INVARIANT CONTENT, SENSITIVITY and RELATIVE TRUTH. To get a relativism with bite, the thesis should be added that one and the same sentence containing a PPT, on a single first-personal use in a context c_1 , may be true relative to the perspective P_I of one assessor and false relative to the perspective P_I of another assessor.⁶

INVARIANT CONTENT, SENSITIVITY and RELATIVE TRUTH (or NONINDEXICAL CONTEXTUALIST TRUTH) are to be understood as constraints on semantic and pragmatic theorising. Truth relativism (nonindexical contextualism) as defined here is not itself a semantic or pragmatic theory. In particular, the theses defining truth relativism (nonindexical contextualism) are neutral on and compatible with different semantic implementations.⁷ So in order to see that the sequenced worlds model is compatible with both nonindexical contextualism and truth relativism, we simply need to show that the sequenced worlds model – in particular, the semantics for PPTs of section 4.8 and the definition of SPEECH ACT CONTENTS_{SW} of section 4.9 – validate INVARIANT CONTENT, SENSITIVITY and NONINDEXICAL CONTEXTUALIST TRUTH or RELATIVE TRUTH. First, then, note that INVARIANT CONTENT fits with SPEECH ACT CONTENTS_{SW} because the diagonal content of a sentence Φ at any context of use c is $\lambda c'. \llbracket \Phi \rrbracket^{c', c'}$, which does not depend on c . Second, SENSITIVITY is compatible with our semantics for first-personal uses of PPTs, since PRO_C takes as its value the sequence coordinate of the index, and the truth of sequenced worlds speech act contents crucially depends on the values of the sequence.

Third, the sequenced worlds model is also compatible with either NONINDEXICAL CONTEXTUALIST TRUTH or RELATIVE TRUTH. To see this, it will be helpful to first clarify the relation between the compositional semantics and these notions of truth and then show that the sequenced worlds semantics of chapter 4 allows the formulation of either notion of truth. We are then in a position to show that the sequenced worlds model is compatible with NONINDEXICAL CONTEXTUALIST TRUTH and also with RELATIVE TRUTH (but not together).

The notion of truth is crucial for the connection between the compositional semantics and the use of expressions in the language for which the semantics is given.

⁶Without this thesis, relativism is compatible with the view that there is only one valid perspective that assessors may adopt, so that a single taste claim will never be true from one assessor’s perspective but false from another assessor’s perspective.

⁷The theses defining relativism are compatible with all major truth relativist views on taste, such as Egan et al. (2005), Egan (2010a), Kölbel (2002), Lasersohn (2005), MacFarlane (2012), Richard (2008) and Stephenson (2007a).

To see why it is important, let us take a step back. Truth-conditional theories of semantic meaning go hand in hand with the idea that to know the meaning of a sentence requires knowing the conditions under which the sentence is true. This knowledge enables speakers to *use* sentences in contexts appropriately. But intensional compositional semantics of the sort we have been considering do not themselves provide a notion of truth that tells us when a sentence, as used in a context, is true. The semantics formulated in section 4.8 provides a notion of *truth at a context and index* for sentences of the language: A sentence Φ is true at a context c and index i iff $\llbracket \Phi \rrbracket^{c,i} = \mathbf{1}$. (Given a sentence as input, the semantics will output a truth value relative to a context and an index.) This notion of truth at a context and index, however, does not in itself determine when a sentence, as used in a context, is true. For the notion does not itself specify a value for the index, given a context in which the sentence is used. Yet truth at a context and index can be used to define a notion of truth that is relevant for speakers' use and assessment of sentences in context. And this second definition of truth in terms of truth at a context and index is the locus where nonindexical contextualists and truth relativists differ. NONINDEXICAL CONTEXTUALIST TRUTH and RELATIVE TRUTH are different ways of connecting the compositional semantics to language use.⁸

The sequenced worlds compositional semantics of section 4.8 allows the formulation of a nonindexical contextualist notion of truth and a truth relativist notion of truth. The nonindexical contextualist version of the sequenced worlds semantics will have TRUTH IN CONTEXT.

TRUTH IN CONTEXT

A sentence Φ is true in the context c iff $\llbracket \Phi \rrbracket^{c,i_c} = \mathbf{1}$,
 where i_c is the index $\langle w_c, t_c, \langle x_c^1, \dots, x_c^n \rangle \rangle$ whose coordinates are determined by c . (Given the formal representation of c as a sequenced world, $i_c = c$.)

The truth relativist version of the sequenced worlds semantics will have TRUTH AT A CONTEXT AND FROM AN ASSESSOR'S PERSPECTIVE.

TRUTH AT A CONTEXT AND FROM AN ASSESSOR'S PERSPECTIVE

A sentence Φ at context c is true from an assessor's perspective P iff for $n \geq \mathbf{1}$,
 $P = \langle w, t, \langle x_1, \dots, x_n \rangle \rangle$ and $\llbracket \Phi \rrbracket^{c,i_P} = \mathbf{1}$, where $i_P = \langle w, t, \langle x_1, \dots, x_n \rangle \rangle$.

We are now in a position to see that the sequenced worlds model of communication is compatible with NONINDEXICAL CONTEXTUALIST TRUTH and RELATIVE TRUTH. We simply have to show that the model is compatible with TRUTH IN CONTEXT as well as

⁸MacFarlane (2003, 2012) calls the extra step, in which the pragmatically relevant notion of truth is defined, the 'postsemantics,' to distinguish it from the 'semantics proper', the recursive definition of truth at a context and index in terms of which the former notion is defined.

with TRUTH AT A CONTEXT AND FROM AN ASSESSOR'S PERSPECTIVE. In the model as presented in chapter 4, truth only occurs in the group-centric norms of assertion and acceptance, which permit assertion (acceptance) of a sentence only if (if) the expressed content is *true from the conversation's perspective*. The content a sentence Φ expresses in a context c is true from the conversation's perspective P just in case Φ is true at c and the index determined by P . But this is just to say that Φ is true in the context c : The conversation's perspective, as defined in chapter 4, just is c , which is formally represented as a sequenced world.⁹ Thus, the index as determined by c is identical to the index as determined by the conversation's perspective. The sequenced worlds model is compatible with TRUTH IN CONTEXT.

The group-centric norms are also compatible with the relativist notion of truth. For a sentence Φ in a context c to be true from the conversation's perspective is for it to be true from the assessor's perspective in the special case where the assessor's perspective is identical to the conversation's perspective. Of course, the distinctive difference of the relative notion of truth from the standard TRUTH IN CONTEXT thus plays no role in the norms of assertion and acceptance. But this does no harm to the norms' *compatibility* with relativism. Relativism does not require relativistic norms of assertion and acceptance.¹⁰

The sequenced worlds model, with the proposals for the semantics of PPTs and pronouns and for speech act content, can be completed with TRUTH IN CONTEXT as well as TRUTH AT A CONTEXT AND FROM AN ASSESSOR'S PERSPECTIVE. It is thus neutral between nonindexical contextualism and truth relativism.

5.3 On Two Implementations of Relativism

It is worth briefly pausing over the different ways that INVARIANT CONTENT, SENSITIVITY, AND RELATIVE TRUTH have been implemented in the relativist literature, and to situate the sequenced worlds model among the approaches. A natural way of classifying the approaches divides them in two groups, depending on the following interrelated theoretical choices: (i) the kinds of coordinates in the index, (ii) the technical representation of mental content and speech act content, and (iii) in the definition of RELATIVE TRUTH, the determination of the coordinates in the index by the context of use or the context of assessment.

⁹In section 4.7.2, THE CONVERSATION'S PERSPECTIVE P^N was defined as follows: The perspective of a conversation is $\langle w_c, t_c, \langle x_1, \dots, x_n \rangle \rangle$, where w_c and t_c are the world and time at which the conversation takes place, and the sequence of individuals $\langle x_1, \dots, x_n \rangle$ is determined by the conversational sequence for the conversation. It represents the individual perspectives of all conversational participants x_1, \dots, x_n in w_c at t_c .

¹⁰See MacFarlane (2012, ch.5) for a discussion of the relation between relative truth and truth norms of assertion, which comes to the same conclusion.

One option is the centered worlds relativism advocated by Egan et al. (2005) and Egan (2007, 2010a) (see also Egan (2011)). On this view, the index of the intensional semantics consists of a triple of a world, time, and an individual. The contents of attitudes and assertions are sets of centered worlds. In evaluations of contents in context for truth and falsity, all of the coordinates in the index are determined by the assessor's location – the index just is the centered world representing the location of the assessor. (In the terminology introduced by John MacFarlane, all coordinates of the index are determined by the *context of assessment*.)

The other option, advocated by MacFarlane (2005a,b, 2011a, 2012), takes a standard of taste in the index. The index of the intensional semantics consists of a tuple of a world and a standard of taste (and perhaps a time). The contents of attitudes and assertions are sets of world-standard tuples $\langle w, s \rangle$ (or world-time-standard triples). In evaluations of contents in context for truth and falsity, the world (and time) coordinate in the index is determined by the context of use – for instance, the context in which the sentence whose content is getting evaluated was used – whereas the standard of taste is determined by the context of assessment. A content p as used in a context of use c_U and assessed at a context of assessment c_A is true iff p is true at $\langle w_{c_U}, s_{c_A} \rangle$.¹¹

Not all existing views in the recent literature that have received the label 'relativism' perfectly fit the mould of one or the other kind of relativism. On some views, the index contains a standard of taste but the assessor's perspective determines all coordinates of the index (Kölbel, 2009). On other views, centered worlds content is coupled with the thesis that assessments for truth and falsity are made relative to the world and time of the context of utterance (Lasersohn, 2005; Stephenson, 2007a).¹²

¹¹On Richard's (2004; 2008) versions of relativism, propositional truth is also relativised to a standard of taste. In contrast to MacFarlane, he does not make explicit the role of the context of assessment in determining the standard of taste coordinate of the index.

¹²Kölbel (2009) advocates a semantic implementation with a tuple of world and standard of taste as the index, yet his position is most naturally read as taking the assessor's context to determine both the world and the standard of taste coordinate. The relativism advocated in Kölbel (2002, 2004) is similar to centered worlds relativism, although propositions are here taken as true or false relative to just 'perspectives,' where perspectives are understood as functions that assign truth values to propositions.

Lasersohn (2005, 2008, 2009) advocates centered worlds relativism insofar as his indices are triples of a world, time, and an individual, and speech act contents are sets of centered worlds. However, he defines a notion of (pragmatically relevant) truth which, in MacFarlane's terminology, assigns truth to a content in context just in case the content is true at the world and time of the context of use and the individual determined by the context of assessment. (Formally, his view bears resemblance to nonindexical contextualist views since he defines a Kaplanian notion of *truth in context* (2005, 666). But Lasersohn treats contexts, the formal object c in the semantics, not as concrete speech situations but as just that: formal objects, fixing values not only for standard features of the speech situation but also a 'judge' (2005, 669–70). Given a particular speech situation, there are as many corresponding contexts as there are potential judges. In evaluating contents for truth and falsity, an assessor picks one such context – usually the one that has herself as the judge.) Lasersohn sides with centered worlds relativism on (i) and (ii) but with MacFarlane on (iii).

Stephenson (2007a,b) also endorses a centered worlds version of (i) and (ii). She does not explicitly

The sequenced worlds account is closer to the first option than the second. It is similar to centered worlds relativism regarding (i) and (ii), since it maintains that the index contains a world, time, and *sequence* of individuals, and endorses sequenced worlds content where centered worlds relativism has centered worlds content. And it is closer to centered worlds relativism than to MacFarlane's relativism on (iii), since the assessor determines the entire index in evaluations for truth and falsity – sometimes selecting the assessor's individual perspective (eavesdropper and other inter-conversational assessments) and sometimes selecting the perspective of the conversational group of which the assessor is a member (intra-conversational assessments).

We will discuss the differences in empirical predictions between the two versions of relativism in chapter 6, sections 6.5 and 6.6. These differences concern tense and the assessment of past and future assertions as well as modality and transworld assessments. As we will see, sequenced worlds relativism sides with centered worlds relativism in their predictions about tense and cross-temporal assessments.

5.4 Eavesdroppers

The neutrality of the sequenced worlds model between nonindexical contextualist and relativist views about predicates of personal taste is a strength of the model. But how are we to decide between the nonindexical contextualist version of the sequenced worlds model and the relativist version? Setting aside worries about the coherence of the notion of relative truth, the decision is a largely empirical matter.¹³ It turns on key empirical data from eavesdropping and retraction cases for which the views differ in their predictions. In this section, I will exemplify the differences as they arise for eavesdropping cases. After reviewing the norm(s) of assertion in section 5.5, I will turn to retraction cases in section 5.7.

Consider (5.1):

- (5.1) [Sal is secretly listening in on the following conversation between Ben and Anna.]
- | | |
|-------------------|---------------------------------------|
| Ben: | Liquorice is tasty. |
| Anna: | That's right. It is tasty. |
| Sal [to himself]: | That's false. Liquorice tastes awful. |

state a definition of (pragmatically relevant) truth. But her remarks on retraction and the assessment of past assertions in (2007a, 517) suggest that she endorses a definition of truth that resembles MacFarlane's approach: A content p as used in a context of use c_U and assessed by judge j at a context of assessment c_A is true iff p is true at $\langle w_{c_U}, t_{c_U}, j_{c_A} \rangle$.

Stephenson's and Lasersohn's systems differ from Egan's system in their predictions of cross-temporal and trans-world assessments of assertions – that is, of assertions made at other times or in counterfactual worlds. This difference will be discussed in chapter 6, section 6.5 on tense and modality.

¹³See chapters 2 and 5 of MacFarlane (2012) for some worries about the coherence of relative truth and a defence of MacFarlane's version of truth relativism.

Relativism predicts that (5.1) can be perfectly acceptable. It is intuitively correct for Anna to say ‘That’s right’ (if liquorice tastes good to her) and it is also intuitively correct for the eavesdropper Sal to say to himself ‘That’s false’ (if liquorice does not taste good to him).¹⁴ According to relativism, Sal’s assertion of ‘That’s false’ in (5.1) is acceptable because it is false relative to his perspective that liquorice is tasty.¹⁵

It seems that relativism’s predictions coincide with intuitive judgments. Assuming that liquorice does not taste good to Sal – after all, he says that it tastes awful – it seems correct for him to dismiss Ben’s assertion that liquorice is tasty as false.

Nonindexical contextualism predicts that Sal’s last assertion is false. Sal assesses Ben’s assertion, which is correct (relative to any perspective) because it is true in Ben’s context of use. While Sal is within his linguistic rights to say ‘Liquorice tastes awful,’ he cannot correctly say that Ben’s assertion is false. On the contrary, nonindexical contextualists predict that Sal could truly assert ‘Liquorice tastes awful. Still, Ben’s assertion is true.’

Relativism seems better positioned to make correct predictions regarding eavesdropping data. However, my main goal here is to show that no matter one’s choice between nonindexical contextualist and relativism, one can and should adopt the sequenced worlds model. I will briefly outline how to account for the nonindexical contextualist interpretation of the data on the sequenced worlds model and then turn to the relativist interpretation.¹⁶

In order to make predictions that differ from relativism’s predictions, nonindexical contextualists need to say that eavesdroppers assess the utterances themselves for truth and falsity. So if Sal is said to be wrong in his assessment of Ben’s assertion, that is because he is mistaken about the assertion’s truth in its context. After all, it follows from TRUTH IN CONTEXT, defined for sequenced worlds content, that ‘Liquorice is tasty’ is true in the context in which liquorice tastes good to Ben and Anna.¹⁷

It is important to register a second option for nonindexical contextualists. They could predict that Sal’s assertion is true by claiming that ‘that’ in (5.1) picks out the proposition expressed by Ben’s assertion rather than the assertion itself. Ordinary truth ascriptions could uniformly be treated as ascriptions of *propositional* truth. Then Sal’s assertion is true because the assertion ‘The proposition that liquorice is tasty is false’

¹⁴Note that eavesdropper cases have mainly been advanced to support relativism about epistemic modal claims, but the same pattern of relativist predictions arises for third-party assessments of taste claims. See Egan (2007, 2–5), Egan et al. (2005), and MacFarlane (2011a, §3.1) for discussion of eavesdropper cases regarding epistemic modals.

¹⁵The monadic, object language predicates ‘true’ and ‘false,’ on the relativist view, are disquotational in the sense that for any context *c* and perspective *P*, the equivalence schema *It is true that* Φ *iff* Φ is true at *c* and from *P*. I will discuss the object-language monadic truth predicate in more detail in section 6.5.1.

¹⁶For some doubts about eavesdropping data from PPTs, see for instance Stephenson (2007a, 518 n.27).

¹⁷To be precise, it is the sentence used by Ben in the context which is true according to TRUTH IN CONTEXT. It is plausible to assume that if utterances are the kinds of things that are truth-apt, their truth values correspond to those of the sentences used in the context of utterance.

is true in Sal's context of use, provided only that 'true' is governed by the equivalence schema *It is true that Φ iff Φ* . I am generally sympathetic to the idea that utterance truth is a theoretician's notion and that ordinary truth ascriptions are ascriptions of propositional truth. However, if nonindexical contextualists give up the claim that the empirical data yield ascriptions of utterance truth, their view ends up making the same predictions as truth relativism on eavesdropping. As we will see, the option of making the same predictions as relativism is available to nonindexical contextualists for cases of retraction and disagreement as well. So the second option will be a recurring theme. However, if they choose the second option for every set of empirical data, their view becomes empirically indistinguishable from relativism. In that case, the decision between truth relativism and nonindexical contextualism will have to turn on non-empirical issues.¹⁸ In what follows, I will generally canvass the more interesting option that sets apart nonindexical contextualism from relativism and will note the second option only in passing.

Things get more interesting on the relativist version. Eavesdroppers are by definition third parties who are *not* members of the conversation. The sequenced world model, as developed so far, is an account of conversation only. It allows for more than one way to explain eavesdropper assessments. Let me spell out the three most promising.

1. Fully deferential eavesdropper. In trying to understand a claim made in some conversation to which he is not party, the eavesdropper could reconstruct the common ground as the shared presuppositions of the *actual* members of the conversation, himself excluded. In his own assessments, he is fully deferential to what is acceptable in the conversation. His assessments reflect what actual members of the conversation should accept. That is, the eavesdropper assesses taste claims relative to the conversation's perspective. Obviously, this will not yield relativism-friendly predictions, as a liquorice-hating eavesdropper would judge the claim 'Liquorice is tasty' true if it was true from the perspective of everyone in the conversation.

2. Inclusive/partially deferential eavesdropper. An eavesdropper could reconstruct a common ground from the shared presuppositions of the actual members of the conversation and himself. His assessments reflect acceptance of expressed contents to the common ground that includes his presuppositions. On this way of understanding the eavesdropper, he is correct in rejecting 'Liquorice is tasty' and judging it false, as relativism predicts, because it is not true from the conversation's perspective, which includes his own, that liquorice is tasty. A problem with this explanation of what eavesdroppers do is that it would predict that an eavesdropper feels addressed by some assertions involving the words 'all of you,' since he includes himself among the conversational participants. It might also seem implausible that eavesdroppers, in their reconstruction of

¹⁸See MacFarlane (2009, §7) and Cappelen and Hawthorne (2009, 20-4) for discussion

the common ground, judge as if they are party to the conversation.

3. Egocentric eavesdropper. Eavesdroppers may be judging assertions from their own, individual perspective. That is, they are not in the business of reconstructing a conversational common ground at all. If they are judging assertions relative to an individual perspective, they cannot be judging the sequenced worlds content that is added to the common ground if accepted in conversation. That is because a sequenced worlds content, with sequences of more than one individual, will require a perspective of more than one individual for its evaluation. However, there is a content nearby that eavesdroppers may plausibly be judging: the Kaplan horizontal. Recall that the Kaplan horizontal is a function from the index to an extension:

$$\text{Kaplan horizontal of } \Phi \text{ at } c: \quad \lambda i. \llbracket \Phi \rrbracket^{c,i} = \\ \{ \langle w, t, \langle x_1, \dots, x_n \rangle \rangle : \llbracket \Phi \rrbracket^{c, \langle w, t, \langle x_1, \dots, x_n \rangle \rangle} = 1 \}$$

If the index is a centered world, this content can be assessed from an individual perspective.

The Kaplan horizontal of ‘Liquorice is tasty’ at c , given an index consisting of a world, time, and individual is the set $\{ \langle w, t, x \rangle : \text{liquorice tastes good to } x \text{ in } w \text{ at } t \}$. This content is true from the perspective of an assessor at $\langle w_1, t_1, a \rangle$ iff liquorice tastes good to a in w_1 at t_1 . This yields relativism-friendly predictions: In (5.1), Sal may assess Ben’s assertion as false because he judges the horizontal content from his own perspective, from which liquorice is not tasty.

Importantly, taking the Kaplan horizontal as the object of eavesdroppers’ assessments also makes the right predictions for *de se* and *de te* assertions. We have given pronouns their standard Kaplanian semantics. So an assertion of ‘I am hungry’ in a context in which Ben is the speaker expresses the Kaplan horizontal $\{ \langle w, t, \langle x_1, \dots \rangle \rangle : \text{Ben is hungry in } w \text{ at } t \}$. An assertion of ‘It’s your turn’ in a context where the intended addressee is Anna expresses the Kaplan horizontal $\{ \langle w, t, \langle x_1, \dots \rangle \rangle : \text{it’s Anna’s turn in } w \text{ at } t \}$. There is thus no risk that because eavesdroppers judge the content relative to their own perspective, they are misled in resolving the reference of pronouns, indexicals and demonstratives.

This last option, the egocentric eavesdropper, seems by far the best if we want to make the relativist’s predictions regarding eavesdropping data on the sequenced worlds model. But it involves introducing a second notion of content, defined in standard ways by the semantics but different from $\text{SPEECH ACT CONTENTS}_{\text{SW}}$. This may seem ad hoc. Why have a different type of content for eavesdroppers’ evaluations?

Eavesdroppers make inter-contextual assessments. If they know that they do not share the context of the conversation with speaker and audience, their proper understanding of utterances will reflect that fact. For example: they will not be inclined to think that ‘here,’ as used by the speaker, refers to where they themselves are standing,

when their location is at a significant distance from the conversation's location. They do not see themselves as involved in the central business of conversation, which is establishing the live possibilities of the conversational group by adding information to the common ground. So the sequenced worlds content cannot be the piece of information they gain if they understand the assertion. The piece of information they gain has to contain all contextual information required for the resolution of indexical expressions – yet be apt for assessment from the eavesdropper's own perspective. The Kaplan horizontal is just this piece of information.

We will return to the evaluation of content for truth and falsity in chapter 6, section 6.5.1.

5.5 Norms of Assertion

In section 5.2 above, we saw that both nonindexical contextualism and truth relativism are compatible with the group-centric norm of assertion formulated in chapter 4 and repeated here.

GROUP-CENTRIC NORM OF ASSERTION

ASSERT_G A speech act content p is appropriately assertable in context c only if p is true from the conversation's perspective in c .

In chapter 4, I mentioned that the group-centric norm makes some predictions that might seem too strong. For talk about taste, the group-centric norm entails that, for instance, 'This cookie is tasty' is assertable only if the cookie tastes good to speaker *and audience*. That is because all interlocutors have to be correctly located in the content, which is true from the conversational perspective just in case the cookie tastes good to each one of them. But why should a speaker have to make sure that she and her audience have a common outlook on taste in order to guarantee that her assertion about the cookie's tastiness is appropriate? The subjectivity of taste claims might seem to be better captured by an egocentric norm that only requires that the cookie be tasty to the speaker. However, we already saw in chapter 4 that egocentric norms are not equipped to explain why in some situations, claims about what is tasty to the speaker are acceptable while bare taste claims are not.

In this section, I will make the positive case for the group-centric norm while giving the intuitions underlying an egocentric norm their rightful place. As we will see, two norms of assertion – a strong group-centric norm and a weak speaker-oriented norm – hold sway over discourse, each related to a different conversational purpose. Here is the basic picture. Conversations are cooperative enterprises with the goal of locating the conversational group, i.e. reducing the group-possibilities in the context

set. When bare taste assertions are made, this goal requires that participants agree on the tastiness of the food in question (or agree to disagree). But while the maximally cooperative, group-centric communicative purpose of bare taste claims is to establish a shared perspective on the tastiness of the food, bare taste claims also serve the more speaker-oriented purpose of giving voice to the speaker's own perspective. Each of these two purposes gives rise to a norm of assertion, which is conditional on the purpose. Judgments about the appropriateness of assertions may reflect either of the norms, depending on the purpose guiding the judgment.

The plan for the section is as follows. I will first show what explanatory work is done by the strong group-centric norm of assertion. I will then turn to intuitive judgments about the appropriateness of bare taste assertions that are weaker than those guided by the strong norm. This will require discussing the expressive-persuasive nature of bare taste assertions and how it is accounted for on the sequenced worlds model. At the end of the section, we will be in a position to state the two norms of assertion.

So let us start with the strong group-centric norm of assertion, $ASSERT_G$, and the conversational goal of establishing a shared perspective on the tastiness of a food in question. Disputes about taste often become unreasonable when it is clear that no agreement can be reached. There is a sense in which bare taste claims, but not explicitly relativised taste claims, are pointless and uncooperative conversational moves when it has already been established in conversation that speaker and audience do not share tastes.¹⁹ It is often, but not always, unreasonable to keep insisting that some food is tasty when one's interlocutor has made it plain that she strongly disagrees with that judgment. In this kind of situation, however, it *is* reasonable to retreat to the claim that the food is tasty to oneself. For illustration, consider again the following example.

- (4.6) a. Ben: Schnitzel is tasty.
 b. Anna: No, it's not tasty! It is bland.
 c. Ben: Well, it's tasty to *me*, at least.
- (4.7) a. Ben: Schnitzel is tasty.
 b. Anna: No, it's not tasty! It is bland.
 c. Ben: ? Well, it is tasty.

The strong group-centric norm $ASSERT_G$, but no egocentric norm, explains the difference in assertability between (4.6c) and (4.7c). The group-centric norm predicts that (4.7c) is not appropriate to assert in this kind of situation because the asserted content is not true from the conversation's perspective. But it makes no such prediction for

¹⁹By 'bare taste claims' I mean first-personal uses of a PPT that is covertly relativised to PRO_C (cf. section 4.8). There are of course other uses of PPTs on which they are covertly relativised to a contextually salient experiencer ('exocentric uses') or on which an experiencer variable is bound by a quantifier. In this section, I put such uses aside.

(4.6c) because the asserted content – the set of pair-centered worlds such that Schnitzel tastes good to the speaker-center – is true from the conversation's perspective in the case where Schnitzel tastes good only to Ben. The group-centric norm captures the reasonableness of bare taste assertions, because their appropriateness conditions reflect the conditions of fully cooperative communicative success, which consists in an update of the common ground that entails that all interlocutors agree on the tastiness of Schnitzel. When in such situations we judge that a bare taste assertion is uncooperative and inappropriate, our judgments are guided by the fully cooperative, group-centric communicative purpose of bare taste claims.

It is a consequence of the strong group-centric norm of assertion that if there is significant divergence in our views on matters of taste, many of our taste assertions are likely to be inappropriate. But very often, especially at the beginning of a conversation about matters of taste, it seems perfectly appropriate to make a bare taste claim such as 'This cookie is tasty,' even when someone among our audience as a matter of fact disagrees. How can we explain such judgments of conversational appropriateness?

Bare taste claims are aimed at establishing a shared perspective. But they also serve the purpose of voicing our own individual perspective. Under normal circumstances, I want my audience to share *my* perspective, and for that I need to put my perspective out there, in the hope that my audience will agree. In many cases, this hope is well-founded. Our perspectives are very often very similar. It is very often reasonable to assume that we are alike in our perspectives on the tastiness of the food in question, be it because it is reasonable to assume that as humans, we share a basic physiological make-up, or because we are similar in our dispositions to enjoy foods according to their taste, or because we belong to a community of values whose members arrive at similar evaluative judgments due to normative pressure towards the coordination of attitudes. Even when there is resistance from my audience that indicates they do not share my perspective, it might be reasonable – up to a point – to sustain the assumption of relevant similarity because there is good reason to think that they might *come to* share my perspective as a result of the conversation. Where the purpose of voicing one's perspective – with an eye to persuading the audience of adopting the perspective – is in the foreground, assertions seem appropriate only if they correctly voice the speaker's perspective and there is some chance that the audience can be persuaded of adopting the perspective. Appropriateness in this sense is captured by the weak norm of assertion that is tied to the more speaker-oriented purpose of voicing one's perspective.

Before we can state the weak norm of assertion, we need to get clearer on the expressive-suggestive nature of bare taste assertions. This requires making precise the assumption of relevant similarity on the sequenced worlds framework. For conversational participants to assume that they are similar with respect to their perspectives on the tastiness of some food is for them to presuppose that they have a *joint perspective*

on the sequenced worlds content p , which says that the food is tasty.

JOINT PERSPECTIVE ON P

$\langle w, t, \langle x_1, \dots, x_n \rangle \rangle$ is a joint perspective on a sequenced worlds content p iff for all individuals x_i and $x_j \in \{x_1, \dots, x_n\}$: either both $\langle w, t, \langle x_1, \dots, x_i, x_j, \dots, x_n \rangle \rangle \in p$ and $\langle w, t, \langle x_1, \dots, x_j, x_i, \dots, x_n \rangle \rangle \in p$, or both $\langle w, t, \langle x_1, \dots, x_i, x_j, \dots, x_n \rangle \rangle \in \neg p$ and $\langle w, t, \langle x_1, \dots, x_j, x_i, \dots, x_n \rangle \rangle \in \neg p$.²⁰

For a pair-centered content p , this means that the sequenced world $\langle w_1, t_1, \langle \text{Ben}, \text{Anna} \rangle \rangle$ is a joint perspective on p just in case either both $\langle w_1, t_1, \langle \text{Ben}, \text{Anna} \rangle \rangle \in p$ and $\langle w_1, t_1, \langle \text{Anna}, \text{Ben} \rangle \rangle \in p$, or both $\langle w_1, t_1, \langle \text{Ben}, \text{Anna} \rangle \rangle \in \neg p$ and $\langle w_1, t_1, \langle \text{Anna}, \text{Ben} \rangle \rangle \in \neg p$. Where p is a content expressed by a bare taste claim, this intuitively says that Anna and Ben have a joint perspective on the tastiness of some food in w_1 at t_1 just in case the food tastes good either to both of them or to neither of them in w_1 at t_1 . A *presupposition* (in the sense defined in section 4.7.1) of *joint perspective on p* is in place in a conversation with participants x_1, \dots, x_n just in case the context set contains only joint perspectives on p . For a conversation between Ben and Anna this means that a presupposition of joint perspective on a pair-centered content p is in place just in case for every sequenced world $\langle w, t, \langle x_1, x_2 \rangle \rangle$ in the context set, either both $\langle w, t, \langle x_1, x_2 \rangle \rangle \in p$ and $\langle w, t, \langle x_2, x_1 \rangle \rangle \in p$ or both $\langle w, t, \langle x_1, x_2 \rangle \rangle \in \neg p$ and $\langle w, t, \langle x_2, x_1 \rangle \rangle \in \neg p$.

Provided that speakers know their own taste and the context set contains the conversation's perspective (the 'actual' sequenced world), an assertion will not be inappropriate (in either weak or strong sense) in a conversation in which a correct presupposition of joint perspective is in place.²¹

Let us now move on to the expressive-suggestive nature of bare taste assertions. It is crucial to realise that the point of bare taste assertions is never just to state one's perspective. We noted that there is a difference between asserting 'This cookie is tasty' and 'This cookie is tasty to me.' The latter is a statement of one's perspective, and it can function as a 'partial retraction' of one's bare taste claim. The former cannot function

²⁰ Note that the negation, $\neg\Phi$, of a bare taste sentence Φ on a first-personal use has two readings depending on the scope of the negation. Where Φ is the sentence 'X is tasty,' the default wide scope reading says, very roughly: it is not the case that X is tasty to all of the conversational participants. The default narrow scope reading says: X is not tasty to either of the conversational participants. For the definition of JOINT PERSPECTIVE ON P to deliver the intuitively correct result, $\neg p$ must be the content expressed by the narrow scope reading of $\neg\Phi$. Thus, where Φ is a bare taste sentence, $\lceil \neg p \rceil$ does not denote the complement set of p . See section 4.11.6 for discussion of negated taste claims.

²¹ See Egan (2007, 2010a) and López de Sa (2008) for two accounts on which a presupposition of relevant similarity is a necessary condition for the felicity of assertions of bare taste sentences.

in this way. So what is it about the bare taste assertion that distinguishes it from the mere statement of one's perspective?

Emotivists and others have long noted that evaluative claims have a persuasive, or recommending, force.²² They recommend a particular attitude towards the object, event, or action in question. On the sequenced worlds model, it is not hard to see how this could be so. If Ben asserts that liquorice is tasty, he proposes to add to the common ground the content that liquorice tastes good to all participants. For his assertion to be appropriate (in the weak sense), a presupposition of joint perspective has to be in place. If no such presupposition is yet in place and liquorice does not taste good to the addressee, she faces a choice. She can either reject the claim or she can *accommodate* the presupposition of joint perspective. In the default case where the common ground is common belief, she accommodates the presupposition if she comes to believe that food of the relevant kind either tastes good to both the speaker and herself or to neither of them (and believes that the speaker believes this and so on). Knowing from the speaker's assertion that liquorice tastes good to the speaker, she will sincerely believe that it tastes good to both of them only if she changes her individual perspective on the tastiness of liquorice, thus bringing it about that the taste claim is true. The kind of accommodation is just what accommodation of any type of presupposition (on the Stalnakerian model) amounts to, viz. adding the missing presupposition to the common ground.²³ It is a peculiarity of the sequenced worlds framework that adding the presupposition of joint perspective may involve changing one's own perspective: For the conversation's perspective to be a joint perspective on the asserted content, the addressee must like what the other participants like. If she does not do so already, she has to change her taste or that of the other participants. Thus, the persuasive force of bare taste assertions amounts to the potential need for accommodation on the hearers' part, which they bring about by changing their perspectives.²⁴

We can thus explain why even when it is understood that the audience disagrees with the speaker about the tastiness of some food, it may still be sensible for the speaker to insist on her judgment as long as she has reason to think that her audience is in a position to accommodate. And even when she has little reason to think that her audience will in fact accommodate, the practical pressure of having to coordinate her perspective with her hearers' perspectives – for instance, when they have to take a collective decision on which food to order in a restaurant – may provide sufficient reason to insist on a

²²See, for instance, Stevenson (1963, ch.2)

²³It is widely (though not uncontroversially) assumed that a rule of accommodation applies to speech acts that carry presuppositions: If a speech act requires presupposition *q* to be appropriate and *q* is not yet presupposed, then *ceteris paribus* the presupposition *q* comes into existence (cf. Lewis (1979b, 340)).

²⁴Relativists like Egan (2010a, 273), MacFarlane (2007a, 20), and Recanati (2007, 93 n.35) have noted the link between the process of accommodation (in Lewis' sense) and the purpose of establishing a shared viewpoint. For Richard (2008, 101), the process of 'accommodation and negotiation' can be found wherever we use vague scalar expressions that give rise to relative truth.

bare taste claim in light of opposition.²⁵

Let me summarise. Judgments about the propriety of bare taste assertions may be guided by different conversational purposes. On the one hand, they may be guided by maximal cooperativeness – a property an assertion possesses if everyone in the conversation can appropriately accept it. These judgments are accounted for by the strong group-centric norm of assertion. They track reasonableness – what a dispute about matters of taste lacks when ‘it makes no sense’ to keep disputing. On the other hand, judgments may track a much lower threshold of appropriateness. In that case, they are guided by the speaker-oriented purpose of voicing one’s perspective and persuading one’s audience of sharing one’s perspective. An assertion counts as appropriate in this weaker sense only if it correctly represents the speaker’s perspective and there is a chance that the hearers may be persuaded – that is, the hearers are in a position to accommodate in such a way as to bring about the asserted content’s truth from the conversational perspective.

In talk about objective matters, these two purposes do not come apart. But in talk about subjective matters, the changes an assertion proposes to make to the common ground may be appropriate relative to the speaker-oriented purpose, yet fail to be appropriate with respect to maximal cooperativeness. Judgments may be guided by the strong group-centric norm of assertion or by the weak speaker-oriented norm of assertion.

STRONG GROUP-CENTRIC NORM OF ASSERTION

ASSERT_G A speech act content p is appropriately assertable in context c only if p is true from the conversation’s perspective in c .

WEAK SPEAKER-ORIENTED NORM OF ASSERTION

ASSERT_W A speech act content p is appropriately assertable in context c only if p correctly locates the speaker and the hearers in c are in a position to accommodate in such a way as to bring about p ’s truth from the conversational perspective at c .

A speech act content_{SW} p correctly locates a speaker S just in case p contains a triple consisting of the speaker’s actual world $@$, her current time t , and a sequence with S in the position that represents S relative to the conversational sequence: Given the conversational sequence $\langle S, \dots \rangle$, there is a triple $\langle @, t, \langle S, \dots \rangle \rangle$ such that $\langle @, t, \langle S, \dots \rangle \rangle \in p$.

²⁵The discussion in this section owes much to Egan (2010a), with whose general outlook on the reasonableness of disputes about taste it is in large agreement. Of course, no claim is made about the converse agreement.

5.6 The Objection from ‘tasty to us’

The sequenced worlds model predicts that an assertion of ‘Liquorice is tasty’ may have the same content as an assertion of ‘Liquorice is tasty to us’ in the same context. Some readers may find this objectionable for the following reason. In a conversation between two people, assertions of ‘Liquorice is tasty’ and ‘Liquorice is tasty to us’ both express (5.2).

$$(5.2) \quad \{\langle w, t, \langle x_1, x_2 \rangle \rangle\}: \text{liquorice tastes good to } \langle x_1, x_2 \rangle \text{ in } w \text{ at } t\}$$

But intuitively, an assertion of ‘Liquorice is tasty to us’ is about what tastes good to the group, whereas an assertion of ‘Liquorice is tasty’ is not. This difference shows in the different assertability conditions of the assertions. For instance, in a conversational context in which it is common belief that liquorice fails to taste good to at least one of the addressees, the assertion ‘Liquorice is tasty to us’ seems infelicitous. But a speaker to whom liquorice tastes good can still felicitously assert ‘Liquorice is tasty’ in that context. This strongly suggests that the two assertions have different truth conditions.

In section 4.11.5, where we first encountered this objection, I postponed the reply until this chapter. The discussion of norms of assertion now puts us in a position to see how two assertions like the above may seem to come apart in their assertability conditions. It will be helpful to first restate the objection in a rigorous way. I will then make a few clarificatory remarks before I explain why the two assertions may seem to have different assertability conditions.

Here is the step-by-step reconstruction of the objection.

1. Let c be a conversational context in which it is common belief that liquorice fails to taste good to one of the addressees. The sequenced worlds content expressed by an assertion of ‘Liquorice is tasty’ in c = the sequenced worlds content expressed by an assertion of ‘Liquorice is tasty to us’ in c = $\{\langle w, t, \langle x_1, \dots \rangle \rangle\}$: liquorice tastes good to $\langle x_1, \dots \rangle$ in w at t [Premise]
2. For any c' , if two assertions made in c' express the same content (have the same truth-conditions), then they have the same assertability conditions in c' . [Premise]
3. So in c , an assertion of ‘Liquorice is tasty’ and an assertion of ‘Liquorice is tasty to us’ have the same assertability conditions. [from 1, 2]
4. But the two assertions do not have the same assertability conditions in c . The assertion of ‘Liquorice is tasty to us’ is infelicitous and the assertion of ‘Liquorice is tasty’ is felicitous. [Premise]
5. Contradiction [3, 4]

6. Hence premise 1 is false: the content expressed by an assertion of ‘Liquorice is tasty’ in $c \neq$ the content expressed by an assertion of ‘Liquorice is tasty to us’ in c , *pace* the predictions of the sequenced worlds model. [from 1, 5]

The objection crucially relies on the claim about the sequenced worlds model in premise 1 and the principle linking truth conditions and assertability conditions in 2. But notice, first, that the conversational context c is not sufficiently specified to guarantee the truth of premise 1. ‘Tasty PRO_C’ and ‘tasty to us’ express the same content only on one of several possible readings of ‘us.’ The first-person plural pronoun ‘we’/‘us’ can pick out any group that includes the speaker. In particular, it can pick out groups including none of the addressees, some or all of the addressees.²⁶ It is only in contexts in which ‘us’ picks out the group consisting of speaker, all addressees, and no one else that ‘tasty PRO_C’ and ‘tasty to us’ express the same content in conversation. So the context c has to be a context that triggers this contextual interpretation if it is to establish the truth of premise 1. But this use of ‘tasty to us’ seems rare. Typically, PPTs are explicitly relativised to present people to mark a difference between them; hence the use of ‘tasty to me’ when retreating from a bare taste claim in the face of opposition. Likewise, the more natural use of ‘tasty to us’ is the exclusive reading, which underlines that some food tastes good to some group including the speaker, even if not to (all of) the addressees. So the scope of cases of which premise 1 is true is limited.

It is also worth noting, second, that premise 2 is far from obvious. Indeed it is routinely dismissed in accounts of the communicative effects of assertions that appeal to pragmatic implicatures, semantic or pragmatic presuppositions, or information structural properties like focus. In the example above, however, I do not think that any of these phenomena is responsible for the perceived difference in assertability conditions. So I will not dispute premise 2 here.

Instead, I maintain that premise 4 is false. The assertions of ‘Liquorice is tasty’ and ‘Liquorice is tasty to us’ do have the same assertability conditions in a suitable context c in which ‘us’ picks out the conversational group. Both assertions are not appropriately assertable in c according to the strong group-centric norm because it is not true that liquorice tastes good to all of the conversational participants. But there may very well be good reason to think that agreement is still possible because the disagreeing addressee is in a position to accommodate. So both assertions are appropriately assertable according to the weak speaker-oriented norm of assertion. Our impression that the assertions

²⁶More precisely, ‘us’/‘we’ allows of inclusive and exclusive readings. On the exclusive reading, the referent is a group that includes the speaker but excludes the addressee (‘I and others but not you’). On the inclusive reading, ‘tasty’ and ‘tasty to us’ do not express the same content in conversation. On inclusive readings, the addressee is included in the group referred to by ‘we’/‘us.’ Inclusive readings can further be distinguished. On so-called 1+2 readings (the numbers refer to first and second person), the group referred to consists only of speaker and addressee(s). On 1+2+3 readings, the group referred to consists of speaker, addressee(s) and third parties.

come apart in appropriate assertability is due to the fact that the difference in overt linguistic material makes different purposes and thus different norms of assertability salient. The speaker-oriented purpose of a bare taste assertion, even in a situation in which it is understood that someone in the conversation disagrees with the claim, is to voice the speaker's perspective and persuade the hearers of adopting that perspective. As long as there is a chance that hearers can be persuaded, the assertion satisfies the weak norm. Our judgments of appropriate assertability of 'Liquorice is tasty' are guided by the speaker-oriented weak norm. In contrast, the explicit relativisation 'to us' in 'Liquorice is tasty to us' makes the group's perspective on the tastiness of liquorice salient and draws attention to the purpose of maximal cooperativeness, which is geared at getting the group's perspective right. In *c*, where it is understood that the perspectives of the members diverge, 'to us' makes salient that no joint perspective is in place. Thus 'to us' draws attention to the falsity of the assertion's content and the failure of the strong norm. As a result, the speaker-oriented purpose gets trumped in salience, and our judgments are guided by the strong group-centric norm.²⁷

5.7 Retraction

Let us now return to a second set of empirical data that will help us decide between a nonindexical contextualist and a relativist completion of the sequenced worlds model. The model is compatible with the predictions of both relativists and nonindexical contextualists regarding cases of retraction, although we will see that it fits much better with relativist predictions. I will start by presenting some retraction data (§5.7.1). In order to explain the different cases of retraction on the common ground model, we will first have to address the issue of how the passage of time in conversation affects the common ground (§5.7.2). After explaining retraction on the sequenced worlds model of the common ground (§5.7.3), I will show that the model can be complemented by either relativist or nonindexical contextualist norms for retraction to yield the respective predictions (§5.7.4). For the retraction of assertions, it will not come as a surprise that there are two norms for retraction, corresponding to the two norms of assertion.

²⁷Recanati (2008) endorses what he calls 'moderate relativism' about aesthetic predicates on which "It is beautiful' means something like *It is beautiful for us*, that is, for the community to which the speaker and his audience belong.' (59) He discusses a problem similar to the objection above: Why, in light of opposition, do speakers sometimes not retreat to a weaker explicit statement about their own aesthetic perspective but keep asserting 'It is beautiful'? Recanati suggests that the disputants appeal to a community standard which they try to shape with their assertions, with 'one foot in the future' (quote from an unpublished manuscript by Johan Brännmark that Recanati cites). What Recanati's account does not explain is why it is significantly less natural to try and shape one's community's standards by making the explicitly relativised aesthetic assertion with the same content. It is an advantage of the pragmatic account I favour that it can account for the difference between bare uses of PPTs and uses on which the predicate is explicitly relativised to the conversational group.

5.7.1 The Data and Its Interpretation

MacFarlane (2012, 12–13) argues that first-personal uses of PPTs give rise to a particular pattern of retraction:

When I was a kid, I once told my mother, ‘Fish sticks are tasty.’ Now that I have exposed my palate to a broader range of tastes, I think I was wrong about that; I’ve changed my mind about the tastiness of fish sticks. So, if someone said, ‘But you said years ago that fish sticks were tasty,’ I would retract the earlier assertion.

Retraction of an assertion is a speech act by which one withdraws the original assertion, performed either explicitly by saying ‘I retract that’ or ‘I take that back,’ or implicitly by acknowledging, for instance, the falsity of one’s original assertion (cf. MacFarlane (2012, section 4.6)). Relativism, so MacFarlane, predicts the obligation to retract assertions involving PPTs (on first-personal uses) in cases like the above. Provided that ‘Fish sticks are tasty’ is true relative to young MacFarlane’s context of assessment, it was appropriate for him to assert it. Now that his taste has changed, however, this very assertion is false relative to his current context of assessment. It is thus obligatory for adult MacFarlane to retract this assertion. Relativism makes this prediction given the following norm for retraction.²⁸

(5.3) An agent in context c_2 is required to retract an (unretracted) assertion of p made at c_1 if p is not true as used at c_1 and assessed from c_2 .²⁹

Nonindexical contextualism, on the contrary, predicts that MacFarlane is not required to retract. Given TRUTH IN CONTEXT, young MacFarlane’s assertion of ‘Fish sticks are tasty’ is true and remains true. Provided the following truth norm for retraction, then, MacFarlane is not required to retract his (still) true original assertion.

²⁸Not all relativists predict retraction data based on this conversational norm for retraction. As we will see below, for instance, Stephenson (2007a) rejects retraction data for PPTs. On her view, speakers evaluate an assertion a with respect to their perspective *at the time the assertion a was made*. Changes in their perspective since the time of assertion are thus irrelevant to their evaluation. With speakers still judging their previous assertions true, no obligation to retract is predicted.

²⁹MacFarlane (2012, 128). According to the norm in (5.3), falsity from one’s current perspective (context of assessment) is sufficient for the obligation to retract. It is not necessary. Other facts are also in themselves sufficient for the obligation to retract. For instance, when it turns out that a speaker did not and still does not have appropriate evidence for her assertion, she is required to retract her original assertion. (This norm is likely if the corresponding norm of assertion that says that one may make an assertion only if one has appropriate evidence is likely.) Other reasons sufficient for an obligation to retract are connected to an assertion’s violation of general norms of conduct such as norms of politeness, etiquette, ethical norms or particular institutional norms as they are in place, e.g., in court. MacFarlane’s discussion of (5.3) is compatible with further norms for retraction holding sway over discourse. His claim is that if a truth norm is uniquely *constitutive* of assertion, then the corresponding truth norm for retraction is constitutive of the speech act of retraction.

- (5.4) An agent in context c_2 is required to retract an (unretracted) assertion of sentence Φ made in context c_1 if Φ is not true in c_1 .³⁰

It is worth noting a consequence of the nonindexical contextualist's view of retraction. On the view, it should be felicitous for MacFarlane to say any of the following:

- (5.5) I won't take that back. But of course, fish sticks aren't tasty.
 (5.6) My assertion from back then is true. Still, fish sticks aren't tasty.
 (5.7) My assertion from back then is true. But I now reject what I said.

Nonindexical contextualism's predictions of felicity are certainly counterintuitive. Retraction cases seem to favour relativism. Nonetheless, it is worth pausing over the data. Stephenson (2007a, 517), for instance, observes that in the following case, Sam's retraction 'seems odd and pathologically meek':

- (5.8) Mary: How's the cake?
 Sam: It's tasty.
 Sue: No it isn't, it tastes terrible!
 Sam: #Oh, then I guess I was wrong.

Stephenson concludes from this case that 'predicates of personal taste do not give rise to the same retraction phenomena [as epistemic modals]' and accordingly provides no account of retraction for taste claims (Stephenson, 2007a, 517).³¹ Similarly, Moltmann (2010, 218) maintains that 'retraction does not seem to hold for sentences with predicates of [personal] taste.'

But the conclusion that it is never obligatory or even felicitous to retract assertions about taste is too quick. Even within one and the same conversation, retraction of assertions involving PPTs can be appropriate, as in (5.9).

³⁰Nonindexical contextualists *could* make the same predictions as relativists by adopting a different norm for retraction:

- (i) An agent in context c_2 is required to retract an (unretracted) assertion of the content p made at c_1 if p is not true in c_2 .

By (i), MacFarlane is required to retract because the content that his original assertion expressed is not true in his current context. See section 5.4 for discussion.

³¹For some doubts about the relativist's predictions regarding the retraction of epistemically modal sentences, see von Fintel and Gillies (2008) and Schaffer (2011), and regarding the retraction of knowledge ascriptions, see Greenough (2011).

- (5.9) Ben: This liquorice and mint tea is tasty! I've never tried this combination before.
- Anna: No, it's not! The anise, minty and sweet flavours together are disgusting. And the aftertaste is the worst.
- Ben: Hmmm. Let me have another sip then... You're right. I take that back. It's not tasty at all.

So why does retraction seem appropriate in (5.9) but not in (5.8)? What distinguishes cases like (5.8) from cases like (5.9) is that in the latter, but not the former, the speaker's change of perspective on the tastiness of the food between his original assertion and his later retraction is plausible. In (5.8), it is 'pathologically meek' for Sam to change his mind about the cake's tastiness simply because Sue opposes his judgment. (Sam is not given any further motivation to change his mind.) In (5.9), on the contrary, Anna gives Ben good reasons to change his mind by drawing his attention to flavours he may not yet have appreciated. The fact that Ben replies by saying 'Let me have another sip' suggests that he had not tried enough of the tea to be completely settled in his opinion.³² Similarly, the naturalness of retraction in MacFarlane's fish sticks example is proportional to the plausibility of him having changed his perspective on fish sticks since his childhood.^{33 34}

³² Susceptibility to reasons and change of perspective seem perhaps more commonplace with judgments about what is fun or entertaining than with judgments about tastiness. In (5.9), one might also think that the tea at no point tasted good to Ben and that he was wrong in asserting 'This liquorice and mint tea is tasty' in the first place. On this interpretation, Ben does not change his perspective during the conversation, but simply comes to realise what his perspective really is. However, if we assume that Ben's first assertion was based on a pleasurable phenomenal experience of drinking the tea, then this interpretation of (5.9) presupposes a metaphysics of taste on which something can fail to taste good to an agent at a time even if the agent has a pleasurable phenomenal experience that disposes them to judgments of tastiness. Limitations of space and expertise prevent me here from settling the metaphysics and physiology of taste. Let it suffice that there can be cases in which a speaker genuinely changes their perspective on the tastiness of food within the period of a single conversation.

³³ One may find the obligation to retract incurred by MacFarlane rather weak for another reason, however. It is not clear just how much we are committed to what we said decades ago, and there seems to be no need to retract something to which we are not anymore committed. One may find that one's childhood claims carry very short-lived commitments and that in MacFarlane's example it is rather odd that someone will take him up on his childhood assertion by saying 'But you said years ago that fish sticks were tasty.' If, however, that challenge is taken seriously, then the commitments taken on by the childhood assertion are accepted as still being in place and the requirement to retract as binding.

³⁴ There are of course also cases of retraction of tastiness assertion in which no change of perspective takes place, but in which the speaker admits that what she said was not even true from the perspective she had as she was saying it. Consider (i).

- (i) [Context: Sal has cooked dinner for his friends Anna and Ben.]
- Ben: Your lasagne tastes awful, Sal.
- Anna: No, it doesn't, and you know that. You're just being mean.
- Ben: God... You're right, Anna. I'm sorry, Sal. I take that back. I'm in a bad mood today. Your lasagne tastes good, as usual.

There are, finally, cases in which speakers ‘partially retract,’ as Pearson (forthcoming, §4.6) observes.

- (4.6) a. Ben: Schnitzel is tasty.
 b. Anna: No, it’s not tasty! It is bland.
 c. Ben: Well, it’s tasty to *me*, at least.

In (4.6), Ben gives up on his original assertion and retreats to the weaker claim ‘it’s tasty to *me*,’ which should be acceptable to Mary even if the cake does not taste good to her. The phenomenon of ‘partial retraction,’ however, is different from the above cases of retraction in a number of ways. In (4.6), Ben is probably not prepared to assert the negation of his previous claim. Nor does it seem that Ben has changed his mind about the cake’s tastiness. Moreover, in (4.6) Ben does not seem to be required to (fully) retract.

The data about retraction is thus a little more complex than it might at first appear. We should expect an account of retraction to explain the difference between the cases and make the intuitively correct predictions for each.

5.7.2 The Common Ground and the Passage of Time

Retraction is the speech act of withdrawing a speech act made at an earlier point in the same or a previous conversation. Before we can explain retraction on the common ground model, we need to get clear on the temporal dimension of discourse on the common ground model. The view that speech act content and mental content are sets of sequenced worlds is a form of temporalism. Propositional temporalism is the view that some contents, or propositions, are time-neutral: their truth value can change with the time of evaluation. For instance, the content expressed by the present tense sentence ‘Bill is sitting’ may be true when evaluated with respect to t_1 but false with respect to t_2 .³⁵ It is easy enough to see that sequenced worlds content is temporalist content. Consider the content that Ben’s assertion of ‘Bill is sitting’ expresses at t_1 in a conversation with two participants:

(5.10) *SIT*: $\{\langle w, t, \langle x_1, x_2 \rangle \rangle$: Bill is sitting in w at t

Here Ben concedes that he said something that was not true from his own perspective at the time. Anna’s opposition is not just based on her own perspective, but calls on Ben to be sincere, that is, to not say what is not true from his perspective. MacFarlane’s relativism and nonindexical contextualism equally predict the felicity of Ben’s retraction in (i).

³⁵The alternative view is propositional eternalism, according to which all contents are time-specific: their truth value does not change over time; contents have their truth value eternally. Propositions conceived of as sets of possible worlds are eternalist proposition. Stalnaker’s model of the common ground assumes eternalist possible worlds content. As the problem I am about to present concerns only temporalism, the amendment of the common ground model I will suggest is not required on Stalnaker’s original model.

Suppose in the world w_1 Bill is sitting at t_1 but not at t_2 . Then *SIT* will contain the triple $\langle w_1, t_1, \langle x_1, x_2 \rangle \rangle$ but not $\langle w_1, t_2, \langle x_1, x_2 \rangle \rangle$, and will thus be true of w_1 at t_1 (and of $\langle x_1, x_2 \rangle$) but not of w_1 and t_2 .

The temporalist nature of sequenced worlds content raises a problem for the view that in conversation, the effect of a successful assertion is the addition of its speech act content to the common ground (or the intersection of the context set with its speech act content). Suppose as above that at t_1 , Bill is sitting in w_1 , the world of Ben's conversation, but not at t_2 . At t_1 , Ben truly asserts 'Bill is sitting.' The assertion is accepted and its speech act content is added to the common ground. The context set will now only contain $\langle w, t, \langle x_1, x_2 \rangle \rangle$ -triples such that Bill is sitting in w at t . Suppose Ben at the later time t_2 truly asserts 'Bill is standing.' Now the context set will be intersected with *STAND*:

(5.11) *STAND*: $\{ \langle w, t, \langle x_1, x_2 \rangle \rangle : \text{Bill is standing in } w \text{ at } t \}$

Since the context set after t_1 contains only Sitting-Bill triples, intersecting it with *STAND* yields the empty set and 'crashes' the context. The model predicts that Ben has contradicted himself. But obviously, Ben's later assertion is compatible with his earlier assertion, and conversation will proceed smoothly.

The solution to the problem is to recognise that the common ground needs to be updated with the flow of time. If participants accept the content of an assertion of 'Bill is sitting' at one point in time, they are not thereby disposed to accept the content of another assertion of 'Bill is sitting' one minute later. What their earlier acceptance seems to dispose them to is acceptance at the later time of the content which, very roughly, an assertion of 'Bill was sitting one minute ago' would express. Provided that participants keep track of the passage of time and presuppose that they all keep track of time,³⁶ the context set is constantly being updated according to the following rule.³⁷

Let t_i be a moment in time, where the index i is a rational number. Moments in time are ordered by the relation *earlier than or simultaneous with*, which we can express by the relation \leq on the moments' indices. The context set is updated as follows.³⁸

³⁶More precisely, the requisite assumption is that participants are capable of keeping track of time, and presuppose that they are so capable, in the minimal sense that they know how much time has passed between their current moment and salient events in the past (they know the *relation* between their present and past events). It need not be assumed that they can place their present time '*de re*'/absolutely in time, e.g. by knowing 'what time it is' – by knowing the number conventional clocks ascribe to their present time.

³⁷'Constantly' – the context set needs to be updated at least as often as conversational moves are made.

³⁸The temporal update rule as formulated does not take into account the information increase that often goes hand in hand with the passage of time. For instance, if speakers come to silently presuppose that the goat has now left the room, this presupposition needs to be added to the common ground independently.

TEMPORAL UPDATE_{CS}

$\forall t_1, t_2$ through which a conversation with context set CS passes and $\forall \langle w, t, \langle x_1, \dots, x_n \rangle \rangle \in CS_{t_1}$: CS_{t_2} contains all $\langle w, t', \langle x_1, \dots, x_n \rangle \rangle$ such that $t' = t + (t_2 - t_1)$, where $(t_2 - t_1)$ is the time interval between moments t_1 and t_2 .

Intuitively, TEMPORAL UPDATE_{CS} takes every world-time-individuals triple that is an element of the context set at t_1 and adds the unit of time that has passed to the time of every triple. So if, for instance, time passes by 1 unit and $\langle w, t_{45}, \langle x_1, \dots, x_n \rangle \rangle$ is in the context set, then the temporal update delivers $\langle w, t_{46}, \langle x_1, \dots, x_n \rangle \rangle$. Participants who take the live possibilities for themselves as a group to be *having features xyz* will a minute later take their live possibilities to be *having had features xyz a minute ago*.³⁹

On the representation of conversation in terms of common ground, the set of pre-supposed contents, the rule corresponding to TEMPORAL UPDATE_{CS} is as follows:

TEMPORAL UPDATE_{CG}

$\forall t_1, t_2$ through which a conversation with common ground CG passes, $\forall p \in CG_{t_1}$ and $\forall \langle w, t, \langle x_1, \dots, x_n \rangle \rangle \in p$: CG_{t_2} contains all and only those q such that $\forall \langle w, t', \langle x_1, \dots, x_n \rangle \rangle \in q, t' = t + (t_2 - t_1)$.

Intuitively, TEMPORAL UPDATE_{CG} takes all the sequenced worlds contents p in the common ground at t_1 and adds the unit of time that has passed to the time of each world-time-individuals triple in every content, delivering a common ground containing contents q . The temporal update of a single content p is the result of applying TEMPORAL UPDATE_{CG} to a common ground containing p as its sole member.

An example will illustrate the temporal update of the context set. At t_1 Ben says to Anna, pointing to a photograph of a man: ‘This man is arriving in London. I will meet him in 30 minutes.’ If Ben’s assertion is successful, the context set, given the conversational context $\langle w_c, t_1, \langle \text{Ben}, \text{Anna} \rangle \rangle$, will contain only $\langle w, t_n, \langle x_1, x_2 \rangle \rangle$ -triples such that the man who is being demonstrated is arriving in London at t_n and x_1 meets that man 30 minutes after t_n in w . Applying TEMPORAL UPDATE_{CS} for the passage of 1 minute (which we will represent by the addition of 1 integer to moments’


³⁹Not all presuppositions concern only the present moment. Many are tenseless (‘2 plus 2 equals 4’), concern past or future moments, or concern the present moment as part of a larger interval in time (‘Bill is a Member of Parliament’). I will have to leave the discussion of the aspectual and tense features of present tense sentences for another occasion. Suffice it to say that an assertion like ‘Bill is a Member of Parliament’, if understood to make a claim concerning an interval of time that stretches at least n moments into the future, will express a set of sequenced worlds such that for every of its world-time-individuals triples $\langle w, t_i, \langle x_1, \dots, x_m \rangle \rangle, \langle w, t_{i+n}, \langle x_1, \dots, x_m \rangle \rangle$ is such that Bill is (still) a Member of Parliament in w n moments after t_i . In this case, the assertion may n moments later still be reaffirmed by asserting the present tense sentence ‘Bill is (still) a Member of Parliament.’

indices) to the context set at t_1 , the resulting context set at t_2 will contain all and only triples $\langle w, t_{n+1}, \langle x_1, x_2 \rangle \rangle$. Intuitively, what these triples have in common is that the man who is demonstrated one minute before t_{n+1} is arriving in London one minute before t_{n+1} and x_1 meets the man 29 minutes after t_{n+1} in w .⁴⁰

5.7.3 Retraction and the Common Ground

We are now in a position to account for retraction on the sequenced worlds model. Let us start by distinguishing different kinds of cases of retraction. Inter-conversational retraction is retraction of a speech act made in a previous conversation, as in MacFarlane's example. Intra-conversational retraction is retraction of a speech act made in the same conversation. Cases of intra-conversational retraction can further be divided. Call those cases in which the challenge happens as an immediate reaction to the assertion and the retraction follows as a direct result, perhaps after some negotiation, *negotiation and withdrawal*. (5.8) and (5.9) above are cases of negotiation and withdrawal. Call those cases in which challenge and retraction happen only later, after the assertion was accepted and the conversation moved on, *later correction*. Figure 5.1 summarises the different kinds of retraction cases.

Figure 5.1: Retraction



5.7.3.1 Conversational Commitments and the Proposal Nature of Assertion

The sequenced worlds model is based on a Stalnakerian account of the assertion and its effect *within* a conversation. The model does not account for the relationship between

⁴⁰TEMPORAL UPDATE_{CS/CG} does not rule out that conversational participants often retain the information contributed by an assertion of a present tense sentence in a different sense of 'retain.' For instance, if someone asserts 'Tesco on Market Street is open,' it may be a shared expectation that Tesco will continue to be open for a while. In that case, at a later point in the conversation, it will for instance be felicitous to presuppose that Tesco is (still) open. We can pragmatically explain the availability of the content of the original assertion, in addition to the temporal update of this content, by appeal to the shared expectation that needs to be in place if the content is later available. In such cases, the common ground contains the content such that if the Tesco store is open, it will remain open for a while.

speech acts taking place in different conversations. Thus, we will have to look beyond the model for inter-conversational retraction cases. For intra-conversational cases, the distinction between negotiation and withdrawal and later correction is important. On Stalnaker's model of conversation, assertions are proposals to add contents to the common ground, the shared presuppositions of speaker and hearers. Stalnaker (1978, 86) says about the proposal nature of assertion: 'To make an assertion is to reduce the context set in a particular way, *provided that there are no objections from the other participants in the conversation.*' (emphasis D.K.) Thus, if retraction happens as a result of immediate challenges, the original assertion is unsuccessful: its content is not added to the common ground. This is the case of negotiation and withdrawal. Retractions that happen only after the proposal was successful and the assertion's content was added to the common ground are cases of later correction.

The proposal nature of assertion also allows for a distinction in the conversational commitments that a speaker incurs in making an assertion. By making a proposal to add a content p to the common ground, a speaker incurs a commitment to accepting p as true (for the purposes of the conversation) – in short, a commitment to p .⁴¹ Speakers do not forfeit this commitment even when the assertion is unsuccessful and p is not added to the common ground. Call this commitment a *speaker commitment*. If the assertion of p is successful, i.e. accepted by all participants as true (for the purposes of the conversation), (accepting) p (as true for the purposes of conversation) becomes a *joint commitment* of speaker and hearers and ceases to be a speaker commitment. The joint commitments in a conversation are all and only the propositions in the common ground. The sum total of a speaker's commitments are her speaker commitments plus the joint commitments. Both commitments are public commitments. Retraction is the explicit withdrawal from a commitment. In intra-conversational cases of negotiation and withdrawal, retraction is the withdrawal from a speaker commitment. In intra-conversational cases of later correction, retraction is the withdrawal from a joint commitment.⁴²

⁴¹The incurred commitment may be a stronger commitment than accepting the content as true (for the purposes of conversation). For instance, it may be the commitment to giving evidence for the content's truth. I have assumed that the norm of assertion states a requirement of truth. Correspondingly, the commitment incurred by an assertion is to the content's truth. If the reader's judgments lean towards a stronger commitment, she might also be inclined to endorse a stronger norm of assertion. For discussion, see footnote 22 in section 4.7.2.

⁴²The proposal nature of assertion is emphasised on many pragmatic models of discourse (see for instance Clark and Schaefer (1989); Clark (1992); Farkas and Bruce (2010); Ginzburg (1996)). Farkas and Bruce (2010) propose a richer structure of 'context,' which contains the common ground – the shared 'discourse commitments' of all participants as well as a set of individual discourse commitments of each participant (those not yet in the common ground), a 'Table' of the issues under discussion, and a 'projected set,' which records the potential changes to the common ground that are proposed by a speech act. An assertion has the effect of adding its content to the speaker's individual discourse commitments, of placing it on the Table, and of creating a projected set that is, roughly, the common ground plus the asserted content. Hearers are prompted to confirm (or reject) the assertion, to the effect that the projected set becomes

5.7.3.2 Intra-conversational Cases of Retraction

The most frequent and the most natural kind of retraction is perhaps negotiation and withdrawal. It is worth illustrating it with an example. (5.9), repeated here, is a case of negotiation and withdrawal.

- (5.9) Ben: This liquorice and mint tea is tasty! I've never tried this combination before.
 Anna: No, it's not! The anise, minty and sweet flavours together are disgusting. And the aftertaste is the worst.
 Ben: Hmm. Let me have another sip then... You're right. I take that back. It's not tasty at all.

In (5.9) Anna rejects Ben's proposal to add the content of 'This liquorice and mint tea is tasty' to the common ground. Hence, the content is not added to the common ground and no joint commitment is incurred. Ben subsequently withdraws his proposal by saying 'I take that back.' He thus withdraws from his speaker commitment to the content of his original assertion. Since Ben, after retracting his claim, also affirms the negation of his original assertion by saying 'It's not tasty at all,' the negation of the original content is added to the common ground (assuming that Anna acts consistently and does not reject this proposal).⁴³ During the negotiation, which could be longer than the three turns in (5.9), many other contents may of course be added to the common ground.

It is a little less clear how to account for cases of later correction on the sequenced worlds model, cases in which an assertion is retracted after it has been accepted by all conversational participants and the conversation has moved on. MacFarlane (2011c) suggests that on Stalnaker's account, later correction might be thought of as 'a new proposal to undo the changes that were made... a proposal to *subtract* that information [the information that got added by the assertion].' He raises the following difficulty for this way of understanding later correction:

the new common ground, and the content disappears from the Table and the speaker's set of individual discourse commitments.

The notion of a speaker commitment to the content of an assertion is broadly in agreement with Farkas and Bruce's notion of 'individual discourse commitment,' and the main points in this section are compatible with Farkas and Bruce's model, which provides a more fine-grained account of the dynamics of (dis)agreement and retraction of assertions. For the discussion of the dynamics of (dis)agreement on the sequenced worlds model, see section 6.3. An important difference between the sequenced worlds model and Farkas and Bruce's discourse model is that their model assumes eternalist possible worlds contents, and hence has no need for a temporal update of the common ground.

⁴³Retracting one's original assertion is not equivalent to affirming the negation of the content originally expressed. One may, for instance, retract because one prefers to suspend judgment on the matter, thus neither affirming nor rejecting the original content. However, assertions of 'I take that back' or 'I was wrong' may imply that the speaker endorses the negation of the original content.

This idea is not easy to integrate with a representation of the common ground as a set of possible worlds: clearly, subtracting a proposition from such a set does not amount to taking the *union* of the proposition and the set. But if we represent the common ground as a set of propositions, we can view the effect of retraction as simple subtraction. (MacFarlane, 2011c, 90)

The problem MacFarlane raises is this: if the effect of assertion is understood as the result of intersecting its content with the context set (whether that be a set of possible worlds or sequenced worlds), there is no straightforward set-theoretic operation on the context set that amounts to the effect of later correction. Taking the union of the context set and the content will not do because this will, among other things, smuggle back in a number of possibilities that have already been ruled out. For instance, the content expressed by ‘Sam is home’ will contain a possibility in which Sam is at home and the conversation happens among gods under the tree Yggdrasil, a possibility that is presumably not a live possibility in most conversations.

MacFarlane suggests that on the representation of the common ground as a set of contents, the effect of later correction amounts to the simple operation of removing the asserted content from the common ground. But this operation does not adequately capture the effect of later correction either. In many conversations, later correction has a domino effect on some of the contents that were added after the retracted assertion. Suppose we are solving a murder case, reasoning with various pieces of evidence we provide in conversation. I say that the murder required physical strength. You say that the butler suffers from advanced osteoporosis. We conclude that the butler is not the murderer, and shift our attention to the gardener. Later, I retract my assertion by saying ‘Wait. Let me take that back. I got it mixed up. It’s the butler’s *wife* who suffers from osteoporosis.’ Clearly, the conversation does not just proceed by erasing from the common ground the single piece of information added by my original assertion. We will also remove the conclusion we drew on its basis, and potentially other information that was offered later on.

What exactly the effect of later correction is depends on a number of contextual and psychological factors such as the participants’ memory and reasoning skills and their assumptions about these skills in each other. One limit case may be the removal of just the originally asserted content. But in some situations, it might not even be clear to the participants what the effect of a particular act of retraction is. Thus the other limit case is the complete breakdown of the common ground with the need to explicitly establish, for every (still relevant) assertion made after the retracted assertion, whether it is still accepted by all participants.

The Stalnakerian sequenced worlds model does not offer a precise account of later correction, at least not without major extension. But in the absence of a detailed story

about all the contextual factors pertaining to the exact effect of retraction, a very simple account may do. According to this account, a later correction is an (implicit) assertion of a content incompatible with the originally asserted content (e.g., its negation). When this new content is added to the common ground, the common ground becomes inconsistent. (The intersection of the context set with the new content is the empty set.) Thus the conversation's context set 'crashes' and needs to be repaired. In repairing, a conversational participant makes her presuppositions based on her beliefs about which information she can still take for granted, her beliefs about which information the other participants still take for granted, her beliefs about what their beliefs are about her beliefs, and so forth.⁴⁴ This process is very much like that at the beginning of a conversation, except that participants now have more clues as to what the others already take for granted.

5.7.3.3 Inter-conversational Cases of Retraction

In inter-conversational retraction cases, as in MacFarlane's fish sticks example, a speaker retracts an assertion she made in a previous conversation, typically in reaction to an explicit challenge. Three things are important for the understanding of inter-conversational retraction on the sequenced worlds model.

First, inter-conversational retraction is required, and plausible, only if the assertion's content is under discussion in the current conversation and if speaker and hearers take the speaker to still be committed to what she asserted in the previous conversation. That is, they treat the content of the original assertion (not quite its content, see points two and three below) either as a speaker commitment of the speaker or a joint commitment in the current conversation. Of course, not everything one ever asserted is thereby a commitment in a current conversation. Whether a speaker's previous commitment becomes a current commitment in conversation depends in part on whether it is a presupposition in the conversation that the speaker made the assertion. Often, this presupposition is added as the result of a challenge – as in MacFarlane's example when someone says 'But you said years ago that fish sticks were tasty' – whereby the commitment incurred by the earlier assertion becomes a current commitment of the speaker.

Second, the current speaker commitment is not exactly a commitment to the content the previous assertion expressed in its conversational context. Recall that conversation is the joint project of establishing the conversational group's possibilities. In establishing the possibilities of one's current conversational group at the time of speaking, we are clearly not committed to accepting as current live possibilities the very possib-

⁴⁴A detailed account of individual speakers' process of revising their presuppositions is beset by the problems known from the work on belief revision. At the same time, logics of belief revision provide a fruitful starting point for such an account (see e.g. van Ditmarsch et al. (2008, ch. 3), Gärdenfors (1988)).

ilities a potentially different conversational group took to be their live possibilities at an earlier time. Rather, we treat earlier conversations in the same way as eavesdroppers treat a conversation they overhear. In both cases, the piece of information gained by remembering or overhearing a speech act is ‘de-contextualised’ information. As we saw in section 5.4, for instance, eavesdroppers are not inclined to think that ‘here,’ as used by the speaker, refers to the shared location; rather, they learn something about the conversation’s location. Similarly for temporal indexicals like ‘now’ and ‘today’, pronouns like ‘I’ and ‘you,’ and other indexical expressions. The Kaplan horizontal of the original assertion of a sentence Φ at the context $c - \lambda i. \llbracket \Phi \rrbracket^{c,i}$, a function from an index to a truth value – delivers just this piece of information.

Third, the current speaker commitment is not exactly a commitment to the Kaplan horizontal of the previous assertion either. Why not? The Kaplan horizontal, just like $\text{SPEECH ACT CONTENT}_{\text{SW}}$, is a temporally unspecific content. So just like the common ground – which can be understood as the participants’ shared conversational commitments – requires updating with the passage of time, the commitment incurred from a previous assertion needs to be the temporal update of the assertion’s Kaplan horizontal. Technically, the speaker’s current commitment is the result of applying the rule $\text{TEMPORAL UPDATE}_{\text{CG}}$ of section 5.7.2 to that Kaplan horizontal, where the sequence of individuals in the index contains as many slots as there are participants to the current conversation. Let me illustrate this with an example. Ben, in a conversation with Anna and Sal, is reminded that a week ago, he said ‘This man was at the airport a day ago,’ pointing at a photograph of the spy Ortcutt. The Kaplan horizontal of his assertion for an index with a sequence of three individuals is *ARRIVE* (given here as the characteristic set).

(5.12) *ARRIVE*: $\{\langle w, t, \langle x_1, x_2, x_3 \rangle \rangle : \text{Ortcutt was at the airport in } w \text{ one day before } t\}$

Intuitively, what Ben’s previous assertion commits him to in his current conversation is the content that Ortcutt was at the airport a week and a day ago. We get this content by applying $\text{TEMPORAL UPDATE}_{\text{CG}}$ to the Kaplan horizontal – that is, by adding the number of temporal units that have passed between Ben’s assertion and his current conversation to the time in each triple in *ARRIVE*.

Putting the pieces together, the commitment incurred by a speaker’s previous assertion is the temporal update of the assertion’s Kaplan horizontal. This commitment may be a speaker commitment, amounting to a standing proposal to add the temporal update of the assertion’s Kaplan horizontal to the common ground. If the speaker commitment has not yet been added to the common ground, the speaker’s retraction is a retraction from this speaker commitment and resembles the case of negotiation and withdrawal. The commitment may also be a joint commitment. If all participants

presuppose that the speaker made the previous assertion and presuppose the temporal update of the Kaplan horizontal, the speaker's retraction is a retraction from the joint commitment and resembles the case of later correction.⁴⁵

5.7.3.4 Generics and Retraction

So does the account of retraction on the common ground model give us the right results for all the cases introduced in section 5.7.1? We will only be able to give a full answer to this question once we have introduced the norms for retraction on the sequenced worlds view in the next section. But one important complication needs to be discussed first. Consider again the fish sticks case. Adult MacFarlane asserts that fish sticks are not tasty, and his interlocutor challenges him by saying 'But years ago you said that fish sticks are tasty.' He thus reminds MacFarlane of the commitment incurred by his childhood assertion. We want to say that the temporal update of the Kaplan horizontal of this assertion is incompatible with the speech act content of the current assertion of 'Fish sticks aren't tasty,' so MacFarlane has two contradictory speaker commitments. MacFarlane is a cooperative conversational participant and resolves the incompatibility by withdrawing from his speaker commitment to the temporal update of his previous assertion's Kaplan horizontal.

The trouble is that *if we treat the assertion of 'Fish sticks are tasty' like any other claim about the present*, it turns out that the temporal update of its Kaplan horizontal is *not* incompatible with the speech act content of MacFarlane's later assertion of 'Fish sticks aren't tasty.' If we treat the earlier assertion as a claim about just the time when it was made, the temporal update of its Kaplan horizontal is a set of sequenced worlds which have in common that years ago fish sticks tasted good to the centers. This content, however, *is* compatible with the speech act content of MacFarlane's later assertion, which is the set of sequenced worlds such that fish sticks do not (currently) taste good to the centers. There are a lot of sequenced worlds of which both of these contents are true – for instance those whose centers have changed their mind on the tastiness of fish sticks. Thus, MacFarlane should be able to respond to the challenge by saying 'Look. Years ago fish sticks were tasty. But I have changed my mind. They're not anymore.'

I think that this reply is not utterly unnatural, but there is certainly a preference for retraction. To see why, it is important to notice first that MacFarlane's assertion involves the bare plural 'fish sticks' and thus makes a *generic* statement about fish sticks. His assertion roughly means that fish sticks are usually/typically/generally tasty. With other generic statements, similar effects regarding their temporal dimension can be observed. Consider an assertion of 'Firemen are fearless,' which is challenged years later. In response to the challenge, the speaker can either retract and say 'I take that

⁴⁵The speaker may want to retract not because of pressure from her hearers but because she discovered defeating evidence or because she does not believe anymore what she said.

back. Firemen aren't fearless,' or she can clarify and say 'Look. Firemen used to be fearless. But now they're not anymore.' The latter response seems preferable, but I take it that both are available.

What is it that explains the availability of either response in both generic statements, yet a preference for retraction in the taste case and a preference for clarifying and holding one's ground in the firemen case? In the literature on generics, it has been recognised that generic statements do not concern only the present moment. That is, 'Firemen are fearless' does not just express a generalisation about *current* firemen. Ariel Cohen (1999a) points out that even if at the moment of speaking, all Supreme Court judges as a matter of fact have a prime Social Security number, the following generic statement would not be true.

(5.13) Supreme Court judges have a prime Social Security number.

Cohen observes that 'the truth, indeed the acceptability of [(5.13)] requires that Supreme Court judges have a prime Social Security number not just at present, but *that this property be expected to hold in the future with some regularity*' (Cohen, 1999a, 225, my emphasis). How this future-directed aspect of generics is best captured in the semantics is an intricate question, but the details need not concern us here.⁴⁶ What matters for our purposes is that generic statements require for their truth that the said property is possessed not only by (some/many/most) present instances of the kind but also by future instances, insofar as the future is a regular continuation of the present.

Thus a speaker who is challenged on her generic assertion of 'Firemen are fearless' made in the past will retract if she accepts that the actual history since her assertion has been a regular continuation of the past and contains too many fearful firemen for the generic to be true. If, however, a speaker stands her ground and clarifies, it is because she does not accept that the actual history has been a regular continuation of the past and so does not think that later instances disconfirm her generic statement. Typically, a speaker will provide some justification about why the actual history has not been a regular continuation, for instance by pointing to a significant change in firemen's properties that is responsible for their loss of fearlessness.

The same holds in the fish sticks case. If adult MacFarlane accepts the challenge, it is because he accepts that the actual history has been a regular continuation of the past, so fish sticks' not being tasty is in conflict with the past claim that they are tasty (and are to be expected to remain tasty). If, on the contrary, he stands his ground and clarifies, it is because he rejects that the actual history has been a regular continuation of the past. To justify that it is not, he may point to some relevant change of properties

⁴⁶Cohen's own proposal is that generics express probability judgments and involve quantification over 'admissible histories' in the metalanguage, where admissibility requires that histories are continuations of the actual history that preserve the relevant regularities.

of fish sticks or a change of properties of his own that are relevant to his experience of taste.

Finally, the difference in preference – retracting or clarifying – between generics involving PPTs and other generics can be explained by the fact that in generic statements of both sorts, the kind to which some property is attributed is attended to. This makes salient the various properties of its instances, which may, or may not, have changed since the generic statement was made. It is thus natural to clarify by pointing to some relevant property of the kind that has changed (and thus makes the actual history an irregular continuation of the past), both in taste and other cases. But generic statements in which an unrelativised PPT occurs with silent PRO_C do not draw attention to the experiencer's perspective – only the food is explicit in the sentence. This explains why in cases where it is the speaker's *perspective* that has changed, clarifying is less frequent. We should thus expect that where the speaker has changed her mind about some food's tastiness on account of a perceived change in the food itself, she will be more inclined to clarify than in cases where it is her perspective that has changed.

In sum, it is the future-directed feature of generics that accounts for the incompatibility of MacFarlane's commitments in the fish sticks case.⁴⁷

There is a residual worry. What about similar cases of non-generic taste assertions? For instance, suppose young MacFarlane asserted 'This dish is tasty.' It might be thought that adult MacFarlane may still retract this claim if he has changed his mind on the tastiness of that particular piece of food. But obviously ascriptions of tastiness to *particulars* come with no expectation that the particular food will remain tasty.

Note that if MacFarlane's assertion is understood as a claim about the *kind* of food he had on that occasion, the explanation appealing to features of generic statements is again available. But if we focus on the reading which says of the particular token dish that it is tasty, I must admit that my intuitions about the naturalness of retraction are not very clear. It seems to me that the more natural reading is about the kind. But where the particular-reading is salient, retraction would involve the past tense assertion 'That dish was not tasty.' (If one retracts by saying, 'That dish isn't tasty,' it is clearly the kind-assertion that is retracted.) But it strikes me that this past tense assertion sits oddly with the admission that the dish did taste good to the speaker at the time. Rather, it seems natural to say 'That dish *was* tasty. I just wouldn't find it tasty nowadays.' It must be admitted, however, that intuitions may differ here. We will come back to issues arising from the interaction of PPTs with time in section 6.5.3 below.

⁴⁷I am thankful to Bernard Nickel and Rachel Sterken for helpful discussion on the temporal dimension of generics.

5.7.4 Norms for Retraction

This outlines the account of retraction on the common ground model of conversation. What still needs to be addressed is the question under which conditions retraction is normatively required of cooperative conversational participants.

The norms for retraction of the relativist and the nonindexical contextualist above mirrored their egocentric norms of assertion. The relativist's norms say, very roughly, that an assertion is permissible only if it is true from the asserter's perspective at the time of asserting; the retraction of that assertion is required if the assertion is not true from the retractor's perspective at the time it is reconsidered. In section 5.5 on assertion above, we have seen that judgments of conversational propriety can be guided by two different conversational purposes, namely the group-centric purpose of maximal cooperativeness and the more speaker-oriented purpose of voicing one's perspective and persuading one's audience of sharing one's perspective. The different judgments correspond to two different norms of assertion, repeated here.

STRONG GROUP-CENTRIC NORM OF ASSERTION

ASSERT_G A speech act content p is appropriately assertable in context c only if p is true from the conversation's perspective in c .

WEAK SPEAKER-ORIENTED NORM OF ASSERTION

ASSERT_W A speech act content p is appropriately assertable in context c only if p correctly locates the speaker and the hearers in c are in a position to accommodate in such a way as to bring about p 's truth from the conversational perspective at c .

Judgments about the obligation to retract previous assertions mirror the judgments about appropriateness of assertion. Corresponding to a strong and a weak norm of assertion, there is a strong and a weak norm for retraction. These norms also register the different conversational purposes that may guide judgments regarding the obligation to retract. Consider first the strong norms for retraction, formulated here as a pair of norms, one for intra-conversational retraction and one for inter-conversational retraction. (Their normative impact is the same.)

STRONG GROUP-CENTRIC NORMS FOR RETRACTION

RETRACT_{S_I} An agent at t_2 is required to retract an assertion of p made in the same conversation at t_1 if the temporal update of p is not true from the conversational perspective at t_2 .

RETRACT_{S2} An agent at t_2 is required to retract an assertion of p made in a previous conversation at t_1 if the temporal update of the assertion's Kaplan horizontal is not true from the conversational perspective at t_2 .

RETRACT_{S1} and RETRACT_{S2} entail that an assertion of a bare taste sentence must be retracted if it is met with a sincere challenge by one of the conversational participants. This might seem too strong. However, there is a sense in which a proposal to add a content to the common ground which faces this kind of resistance should be dropped. If the temporal update of (the Kaplan horizontal of) the assertion's speech act content_{SW} p is not true from the current conversational perspective because at least one of the participants disagrees with the speaker on the tastiness of the food in question, it is not appropriate to add that content to the common ground, and consequently, the proposal to add that content should be dropped to allow for the fully cooperative progression of the conversation.

In most cases of retraction, however, intuitive judgment tracks the speaker-oriented purpose of conversation. The judgments reflect the prescriptions of a weaker pair of retraction norms, mirroring the weak norm of assertion, which permits assertion only where the asserted content correctly locates the speaker and hearers can accommodate in such a way as to bring about the content's truth from the conversational perspective.

WEAK NORMS FOR RETRACTION

RETRACT_{W1} An agent at t_2 is required to retract an assertion of the speech act content p made in the same conversation at t_1 if (i) the temporal update of p is not true from the conversational perspective at t_2 and (ii) hearers at t_2 are in no position to accommodate in such a way as to bring about the truth of the temporal update of p .

RETRACT_{W2} An agent at t_2 is required to retract an assertion of p made in a previous conversation at t_1 if (i) the temporal update of the assertion's Kaplan horizontal is not true from the conversational perspective at t_2 and (ii) hearers at t_2 are in no position to accommodate in such a way as to bring about the truth of the temporal update of the assertion's Kaplan horizontal.

Let us see what the predictions are of the weak and strong norms for retraction regarding the data with which we started out in section 5.7.1. In MacFarlane's fish sticks case, it is not true from the conversational perspective of adult MacFarlane that fish sticks are tasty, because fish sticks do not taste good to MacFarlane anymore. This is a case of inter-conversational retraction. Both the weak norm RETRACT_{W2} and the strong norm RETRACT_{S2} predict that MacFarlane should retract, since the temporal update of

the Kaplan horizontal of his childhood assertion ‘Fish sticks are tasty’ is not true from his current conversational perspective for the reason that fish sticks fail to be tasty to MacFarlane himself (see section 5.7.3.4 for details). This prediction coincides with the relativist’s prediction.

In Stephenson’s cake case, repeated here, retraction seemed ‘pathologically meek and odd.’

- (5.8) Mary: How’s the cake?
 Sam: It’s tasty.
 Sue: No it isn’t, it tastes terrible!
 Sam: #Oh, then I guess I was wrong.

In section 5.7.1, we explained the oddity of the retraction with the implausibility of a change of mind on the part of Sam simply on account of Sue’s resistance to accept his assertion. (5.8) is an intra-conversational case of negotiation and withdrawal. The strong norm RETRACT_{S_I} predicts that Sam should retract, since the fact that the cake does not taste good to Sue makes Sam’s assertion false from the conversational perspective. The weak norm RETRACT_{W_I} , however, does not predict an obligation to retract, since Sam’s hearers *are* in a position to accommodate in such a way as to bring about the truth of the temporal update of the content expressed by Sam’s assertion. All Sue needs to do is change her mind about the taste (or act as if she changed her mind and accept the content of Sam’s assertion merely for the purposes of the conversation). Provided that Sam does not change his mind, relativism also predicts that he is not obliged to retract. (5.8) is thus evidence that the weak norms for retraction are needed on the common ground model to make the predictions of relativism.

In the liquorice and mint tea case, repeated here, retraction did seem natural to the same degree that it seemed plausible that Ben changed his mind about the tea between his first assertion and his later retraction.

- (5.9) Ben: This liquorice and mint tea is tasty! I’ve never tried this combination before.
 Anna: No, it’s not! The anise, minty and sweet flavours together are disgusting. And the aftertaste is the worst.
 Ben: Hmmmm. Let me have another sip then... You’re right. I take that back. It’s not tasty at all.

(5.9) is an intra-conversational case of negotiation and withdrawal. Provided that Ben did change his mind, both RETRACT_{W_I} and RETRACT_{S_I} predict that Ben should retract. That is because from the later conversational perspective in (5.9), the temporal update of the content of Ben’s assertion is not true, due to Ben’s change of mind, and the hearer Anna is in no position to accommodate in such a way as to bring about the truth of

the temporal update of the content of Ben's assertion from the current conversational perspective; after all, she cannot change Ben's taste preferences single-handedly. Relativism also predicts that Ben is required to retract, since it is not true from his current context of assessment that the tea is tasty.

Finally, recall (4.6), a case of 'partial retraction.'

- (4.6) a. Ben: Schnitzel is tasty.
 b. Anna: No, it's not tasty! It is bland.
 c. Ben: Well, it's tasty to *me*, at least.

We noted above that (4.6) differs from other retraction cases in crucial ways. For instance, it would be odd for Ben to retract by asserting the negation of his original assertion. And it seems implausible that Ben's partially retracts as a result of changing his mind about the taste of the cake. Provided that Ben does not change his mind, relativism predicts that he need not retract. The weak intra-conversational norm for retraction, $\text{RETRACT}_{\text{WI}}$, equally predicts that Ben is not required to retract, as his hearer Mary *is* in a position to bring about the relevant content's truth from their current conversational perspective. But what does Ben's partial retraction amount to? There is a sense in which Ben has withdrawn from his proposal to add the content of his original assertion to the common ground in reaction to Mary's rejection. He settles for the weaker claim, which is acceptable for Mary even if cake does not taste good to her. On the common ground model, partial retraction in (4.6) amounts to the withdrawal from a speaker commitment, without the endorsement of the negation of the original assertion. Partial retraction is predicted by the strong norm $\text{RETRACT}_{\text{SI}}$, since the temporal update of the content of Ben's assertion is not true from the later conversational perspective.

In sum, the relativist's predictions are captured by the weak norms for retraction. In most cases, the predictions of the strong norms for retraction coincided with those of the weak norms. But in Stephenson's cake case, the strong norms predicted an obligation to retract, contrary to intuitive judgments. It is in cases of 'partial retraction' that the strong norms do explanatory work.

Nonindexical contextualism makes the prediction that in none of the cases discussed the speaker is required to retract. That is because the view is naturally combined with a norm for retraction that prescribes retraction only if, roughly, the assertion was not true from the speaker's perspective at the time of assertion. It is possible to state a weak norm for retraction that delivers roughly similar predictions.

NC WEAK NORM FOR RETRACTION

RETRACT_{WNC} An agent at t_2 is required to retract an assertion of the speech act content p made at t_1 if (i) p was not true from the conversational perspective at t_1 and (ii) hearers at t_1 were in no position to accommodate in such a way as to bring about the truth of the temporal update (of the Kaplan horizontal) of p .

To illustrate, MacFarlane is not required to retract his childhood assertion of ‘Fish sticks are tasty’ according to this norm for retraction because the content of his assertion was true from the conversational perspective at the time of assertion.⁴⁸ In all but Stephenson’s cake case, however, the nonindexical contextualist’s predictions are counterintuitive. As we saw in section 5.7.1, the nonindexical contextualist’s pragmatics predicts the felicity of the following assertions by adult MacFarlane.⁴⁹

(5.5) I won’t take that back. But of course, fish sticks aren’t tasty.

(5.6) Although what I said back then is false, I won’t take it back.

(5.7) My assertion from back then is true. But I now reject what I said.

On the common ground model of communication, there is a simple explanation of why these assertions are odd. By not retracting, the speaker indicates that he stands by his conversational commitments, viz. the temporal update of the content of ‘Fish sticks are tasty.’ But by following this up with ‘fish sticks aren’t tasty’, the speaker makes a contradictory proposal, the proposal to intersect the context set with a content that

⁴⁸Or, if the asserted content was not true from the conversation’s perspective because fish sticks did not taste good to one of the hearers, clause (ii) of the norm would most probably have failed: his hearers *were* in a position to accommodate in such a way as to bring about the content’s truth from the conversation’s perspective.

⁴⁹Nonindexical contextualism *could* make the same predictions as relativism regarding the obligation to retract. At the cost of losing distinctiveness, nonindexical contextualists could adopt norms equivalent to the relativist’s norms, using the notion of TRUTH IN CONTEXT. The strong and weak retraction norm for intra-conversational retraction would look as follows:

RETRACT_{NCS_t} An agent at t_2 is required to retract an assertion of p made in the same conversation at t_1 if the temporal update of p is not true in the context at t_2 .

RETRACT_{NCW_t} An agent at t_2 is required to retract an assertion of the speech act content p made in the same conversation at t_1 if (i) the temporal update of p is not true in the context at t_2 and (ii) hearers at t_2 are in no position to accommodate in such a way as to bring about the truth of the temporal update of p .

This set of norms, however, still leaves nonindexical contextualism with a range of counterintuitive predictions. For instance, adult MacFarlane could felicitously assert the following:

(i) I take that back. Still, my assertion is true.

(ii) I was wrong, fish sticks aren’t tasty. Still, my assertion is true.

contains only possibilities in which fish sticks taste good to none of the conversational participants.⁵⁰ However, since the previous assertion's commitment has been kept, the speaker is also making a proposal to intersect with possibilities in which fish sticks taste good to all of the participants. The two proposal cannot both be accepted without 'crashing' the context set – reducing it to the empty set.

Thus, the sequenced worlds model of the common ground sits oddly with the nonindexical contextualist's predictions about retraction. But this oddity is not due to peculiar theoretical choices on the sequenced worlds model. Rather, the oddity reflects the counterintuitiveness of some consequences of nonindexical contextualism – as witnessed by (5.5)–(5.7). These are consequences that nonindexical contextualism has independently of its implementation in a sequenced worlds framework. The bad fit with nonindexical contextualism regarding retraction indicates that a relativist version of the sequenced worlds model should be favoured.

5.8 Conclusion

Chapter 4 developed the fundamentals of the sequenced worlds view. At the beginning of this chapter, we saw that the view is neutral between nonindexical contextualism and truth relativism about predicates of personal taste. However, the subsequent discussion of eavesdropping and retraction data showed that the predictions of relativism provide support for a relativist version of the sequenced worlds view. Among other things, the discussion of pragmatic phenomena of conversation showed that eavesdroppers, who are not part of the conversational project of locating the conversational group, receive and evaluate the information carried by an assertion's horizontal content rather than diagonal content; that two norms of assertion are in place depending on the purpose(s) of the conversation; and that over the course of a conversation, the temporalist common ground is in constant temporal update in order to preserve the information accumulated in the conversation.

⁵⁰That is, assuming that 'Fish sticks aren't tasty' receives a narrow scope reading for the negation. See chapter section 4.11.6 on the interaction of negation with PPTs.

SEQUENCED WORLDS, DISAGREEMENT, TRUTH ASCRIPTIONS

6.1 Introduction

Chapter 4 introduced the sequenced worlds view. Chapter 5 showed that the view is compatible with nonindexical contextualism and truth relativism. At the same time, it became clear in chapter 5 that eavesdropping and retraction data lend support to a relativist version of the sequenced worlds view. In this chapter, I continue the development of the sequenced worlds view to account for disagreement cases and monadic truth ascriptions. In section 6.2, I give a sufficient conditions account of disagreement within the sequenced worlds model and show how we can make sense of the notion of faultless disagreement. In section 6.3, I provide the beginnings of an account of the dynamics of agreement and disagreement about matters of personal taste on the sequenced worlds model of common ground. The discussion of disagreement will put us in a position in section 6.4 to take up the question of the cognitive significance of *de se* and subjective attitudes. Returning to empirical questions regarding sequenced worlds relativism, I will then explore the interaction between predicates of personal taste and tense as well as modality in sections 6.5 and 6.6. In chapter 5, two major strands of relativist theories were identified, centered worlds relativism and MacFarlane's relativism, and the proximity of centered worlds and sequenced worlds relativism was noted. The discussion of tense and modality will allow us to see where sequenced worlds relativism differs from MacFarlane's relativism in empirical predictions, and will confirm the similarity of centered and sequenced worlds relativism. I will close with a summary and discussion of the roles of horizontal and diagonal content on the sequenced worlds view in section 6.7.

6.2 Disagreement**6.2.1 Introduction**

Perhaps the most frequently championed argument in favour of (truth) relativism and against (indexical) contextualism is an argument from disagreement. Contextualism,¹

¹I will use 'contextualism' to mean indexical contextualism, and will use the explicit 'nonindexical contextualism' as before to refer to nonindexical contextualism.

it is urged, makes it mysterious why Ben and Anna in (6.1) would regard themselves as disagreeing.

- (6.1) Ben: This cake is tasty.
 Anna: No/Nuh-uh/I disagree, it's not tasty!

According to a simple version of contextualism, Ben expresses the content *that the (contextually salient) cake is tasty to Ben*, and Anna expresses the content *that the cake is not tasty to Anna*. These two contents are compatible. So Anna and Ben should not take themselves to be disagreeing, and the use of disagreement markers like 'No', 'Nuh-uh', and 'I disagree' should be infelicitous. In fact, this simple contextualism predicts that Ben and Anna disagree no more in (6.1) than they do in (6.2).

- (6.2) Ben: I'm a philosopher.
 Anna: #No/Nuh-uh/I disagree, I'm not a philosopher.

Relativism, on the contrary, predicts that Ben and Anna disagree in (6.1) but not in (6.2). In (6.1), there is a single content, *that the cake is tasty*, which Ben accepts by virtue of asserting it and which Anna rejects by virtue of asserting its negation. There is no such content in (6.2). Relativism, it is concluded, saves disagreement about subjective matters.²

Some advocates of relativism make the stronger case that relativism is uniquely capable of explaining the possibility of *faultless* disagreement — 'that is, situations where one person asserts or believes some sentence content, and another person believes or asserts its negation, but where neither person has made a mistake.' (Lasersohn, 2009, 360) Thus in (6.1) Ben and Anna disagree *faultlessly* because each of them asserts a content that is true relative to himself/herself. Faultlessness is explained as truth relative to the asserter's or believer's own perspective.³

However, not all relativists put a great deal of emphasis on faultlessness (MacFarlane, 2007a; Stephenson, 2007a). Some even agree with critics of relativism and think that the combination of single content and truth from each of the parties' own perspective does not account for faultless disagreement (MacFarlane (2012, ch. 6), Richard (2008, ch. 5)).⁴

²Versions of the argument from disagreement can be found throughout the relativist literature, for instance in Kölbel (2002), Kölbel (2004, esp. 53-5), Kölbel (2009, 389-91), Lasersohn (2005), Lasersohn (2008, 306-7), Lasersohn (2009, §1), Lasersohn (2011), MacFarlane (2007a), MacFarlane (2012, ch. 6), Richard (2004), Richard (2008, chs. 4&5), Stephenson (2007a, §§2,3), and Wright (2006).

³See Kölbel (2002, 2004, 2009), Lasersohn (2005, 2009), and Wright (2006) for the argument for relativism from faultless disagreement.

⁴Anti-relativist objections to the claim that relativism is able to account for faultless disagreement are given by Moltmann (2010), Rosenkranz (2008), and Stojanovic (2007).

Initial enthusiasm has also been curbed by general worries about the prospects of providing a necessary-and-sufficient-conditions analysis of disagreement.⁵ As a result, some relativists stress eavesdropper (Egan, 2007) or retraction data (MacFarlane, 2012) over disagreement data. Others have offered independent motivation for relativist semantics from the embedding behaviour of PPTs in larger expressions, e.g., in attitude contexts and presuppositional constructions (Laserson, 2008, 2009).

I share both the worries about the project of giving an *analysis* of disagreement and the doubts about the early relativists' explanation of faultless disagreement. My goal in this section is (i) to give a sufficient-conditions account of disagreement on the sequenced worlds model, which makes the right predictions for disagreement about subjective matters and disagreement in *de se* communication; and (ii) to show in which ways a sequenced worlds relativist is (and is not) entitled to talk of 'faultless disagreement.' I will begin by stating the relativist account of disagreement on the sequenced worlds model (§6.2.2). After a note on transworld disagreement (§6.2.3), I will explain in which ways disagreements about subjective matters may be said to be faultless (6.2.4). I will close with a brief discussion of a nonindexical contextualist account of disagreement on the sequenced worlds model (§6.2.5). In the next section, we will see how some aspects of the pragmatics of agreement and disagreement in conversation can be modelled on the sequenced worlds framework (§6.3).

Before proceeding, it is important that we distinguish the two projects involved. As Cappelen and Hawthorne (2009, 60-1) point out, 'agree' has both a state and an activity meaning. The same goes for disagreement. '*A* and *B* disagree' can describe a *state* in which *A* and *B* are in virtue of having made conflicting speech acts or having conflicting attitudes; or it can describe the *activity* of *A* and *B* engaging in a discussion or dispute in which they can felicitously use expressions like 'No', 'Nu-huh', 'I disagree', or 'That's false.' The latter requires some sort of interaction between *A* and *B*, the former does not. Thus, *A* and *B* can disagree in the state sense without disagreeing in the activity sense. They may never have met, yet disagree in virtue of holding conflicting beliefs. Cleisthenes, the father of Ancient Greek democracy, and Confucius disagree on the ideal form of government, even though they presumably did not even know of each other's existence. Likewise, *A* and *B* can engage in the activity of disagreeing even when they do not disagree in the state sense. They may, for instance, have exactly the same beliefs about the moral (in)admissibility of assisted suicide, yet engage in a dispute about the topic to practice their debating skills. Speakers may also engage in a dispute (for a while) without disagreeing in the state sense because they mistakenly take the contents of their assertions to be in conflict. For instance, they may have a dispute about whether or not visiting relatives can be boring, where one of them means *that the activity of visiting relatives can be boring* and the other means *that relatives who visit*

⁵ See footnote 7 in this chapter for some of these worries.

can be boring.⁶ In section 6.2.2 and 6.2.4, I will give a sufficient-conditions account of disagreement in the state sense. This has been the main target of relativist and contextualist discussions. In section 6.3, I will address (dis)agreement in the activity sense by looking at the pragmatics of disputes about taste.

6.2.2 Disagreement on the Sequenced Worlds Model

The distinction between *intra*-conversational and *inter*-conversational phenomena has been important in the treatment of eavesdropping and retraction. It is equally important for disagreement in the state sense. A speaker may disagree with another speaker who is party to the same conversation, or she may disagree with a speaker who is party to another conversation. Moreover, as we will see in due course, another distinction is important to draw. Two agents may disagree in virtue of their assertions (whether or not belief is their attitude to the content of their assertions); or they may disagree in virtue of their beliefs (whether or not their beliefs are expressed in speech). There are thus four possibilities: (i) intra-conversational disagreement in virtue of assertions, (ii) intra-conversational disagreement in virtue of beliefs, (iii) inter-conversational disagreement in virtue of assertions, and (iv) inter-conversational disagreement in virtue of beliefs. On the sequenced worlds model, a sufficient-conditions account is given for each possibility. Let us take them in turn.

INTRA-CONVERSATIONAL DISAGREEMENT

Two parties to a conversation disagree if one asserts a speech act contents_{SW} p and the other asserts a speech act contents_{SW} q that is incompatible with p .⁷

⁶Moreover, as Cappelen and Hawthorne (2009, 60) remark, the two uses of 'agree'/'disagree' are grammatically distinct in that the event use but not the state use felicitously combines with progressive aspect. 'Ben and Anna are disagreeing about the existence of free will' only has a reading that requires Ben and Anna to be engaged in the activity of disputing to be true.

⁷INTRA-CONVERSATIONAL DISAGREEMENT, like the three other accounts of disagreement, gives only a sufficiency condition, not necessary and sufficient conditions. That is because there are many cases that plausibly count as cases of disagreement, in which there is no incompatibility between the assertions' speech act contents_{SW}. For instance, the disagreement may arise from pragmatically implicated contents. Huvenes (2012, 174) draws attention to the following case based on an example by Grice (1975):

- (i) A: Mr. X is meeting a woman tonight.
B: I disagree, he is meeting his wife.

A 's and B 's speech act contents (in Grice's terms, what each *says*) are compatible. Mr. X's meeting a woman does not preclude that it is his wife he is meeting. Yet B 's use of 'I disagree' does not seem infelicitous. The disagreement can be explained if B is understood to be disagreeing not with what A said but with what he implicated. According to Grice, A under normal conditions implicates that the woman to be met is not X's wife, a claim with which the content of B 's assertion is incompatible.

Huvenes (2012, 167) points to another type of case, in which the speech act contents are compatible. (The following case differs from Huvenes' own example.)

- (ii) Sal: I like *The Sopranos*.
Ben: Oh, no/I disagree. I hate *The Sopranos*.

INTRA-CONVERSATIONAL DISAGREEMENT relies on the notion of incompatibility. But what is it for two sequenced worlds contents to be incompatible? As a first approximation, a modal analysis will do: Two sequenced worlds contents are incompatible iff they cannot both be true from the same (conversational) perspective. But a modal analysis of this kind has obvious problems. As MacFarlane (2007a, 24) points out, if one of the contents is a necessary falsehood, then no matter what the second content is, it will be impossible for both to be true (from the same conversational perspective) simply because it is impossible for the necessary falsehood to be true (from any conversational perspective). Moreover, a modal analysis of incompatibility in an account of disagreement overgenerates cases of disagreement. Suppose a physical theory *T* entails that some sentence *S*, which says something specific about the melting of ice in the Arctic, is false. Then any two speakers who utter *S* and ‘*T* is true,’ respectively, will turn out to disagree. But where the entailment is far from obvious, and the two claims intuitively about different subject matters, the speakers do not seem to disagree.

To my knowledge, there is no account of incompatibility that avoids all problems of this kind. Fortunately, the modal analysis makes the right predictions for many of the cases we are interested in. I will thus rely on it in my discussion of the cases. Officially, however, the account of disagreement in conversation will have to rely on a notion of incompatibility that is based on an intuitive notion of what it is for the truth (from some perspective) of one content to *preclude* the truth (from the same perspective) of another content.⁸ I trust that the reader is sufficiently familiar with the notion of preclusion to competently arrive at judgments of incompatibility. Again, for our cases involving predicates of personal taste, the modal analysis delivers intuitively correct results.⁹

INTRA-CONVERSATIONAL DISAGREEMENT makes the intuitively correct predictions regarding disagreements about matters of personal taste. Consider again (6.1).

Judgments on (ii) may vary, but if take Ben’s reply to be natural, it is tempting to understand Sal and Ben as disagreeing about the television series *The Sopranos*. However, Sal’s liking of it is perfectly compatible with Ben’s hating it. One way of explaining the disagreement is to point to the incompatibility of the attitudes expressed – incompatibility in roughly the sense that a single rational agent could not coherently hold both attitudes at the same time.

⁸I follow MacFarlane’s lead here in relying on a notion of preclusion that is not made fully precise. For some remarks as to how one may elucidate the notion of preclusion in terms of its formal properties, see MacFarlane (2012, 154–5).

⁹INTRA-CONVERSATIONAL DISAGREEMENT makes one idealising assumption, namely that the two assertions happen within a short enough time interval. Without this assumption, the account would predict that a speaker who says ‘The one-minute news are on’ would disagree with her interlocutor who says, more than a minute later, ‘The one-minute news are over.’ To also account for these cases, intra-conversational disagreement would have to be stated in terms of the temporal update of the first assertion’s horizontal content and the second assertion’s horizontal content. See the discussion of INTER-CONVERSATIONAL DISAGREEMENT below.

- (6.1) Ben: This cake is tasty.
 Anna: No/Nuh-uh/I disagree, it's not tasty!

In (6.1), Ben expresses the speech act contents_{SW} $CAKE_{I\phi_2}$. Anna expresses $NOT-CAKE_{I\phi_2}$.

- (6.3) $CAKE_{I\phi_2}$: $\{\langle w, t, \langle x_1, x_2 \rangle \rangle$: to $\langle x_1, x_2 \rangle$, the contextually salient cake tastes good in w at t \}
- (6.4) $NOT-CAKE_{I\phi_2}$: $\{\langle w, t, \langle x_1, x_2 \rangle \rangle$: to $\langle x_1, x_2 \rangle$, the contextually salient cake does not taste good in w at t \}

$CAKE_{I\phi_2}$ and $NOT-CAKE_{I\phi_2}$ are incompatible as there is no conversational perspective from which they are both true (for all $\langle w, t, \langle x_1, x_2 \rangle \rangle$, it cannot be the case that both $\langle w, t, \langle x_1, x_2 \rangle \rangle \in CAKE_{I\phi_2}$ and $\langle w, t, \langle x_1, x_2 \rangle \rangle \in NOT-CAKE_{I\phi_2}$). Hence INTRA-CONVERSATIONAL DISAGREEMENT predicts that Anna and Ben in (6.1) disagree.

INTRA-CONVERSATIONAL DISAGREEMENT is also correct in making no prediction of disagreement in (6.2) and other cases involving first-personal pronouns. (However, the account does not predict that speakers do *not* disagree, since it only gives sufficient conditions for disagreement, not necessary conditions.)

- (6.2) Ben: I'm a philosopher.
 Anna: #No/Nuh-uh/I disagree, I'm not a philosopher.

In (6.2), Ben expresses $PHIL_I$. Anna expresses $NO-PHIL_2$.

- (6.5) $PHIL_I$: $\{\langle w, t, \langle x_1, x_2 \rangle \rangle$: x_1 is a philosopher in w at t \}
- (6.6) $NO-PHIL_2$: $\{\langle w, t, \langle x_1, x_2 \rangle \rangle$: x_2 is not a philosopher in w at t \}

$PHIL_I$ and $NO-PHIL_2$ are not incompatible since there is a conversational perspective from which they are both true – the perspective $\langle w', t', \langle \text{Ben}, \text{Anna} \rangle \rangle$ which is such that Ben is but Anna is not a philosopher in w' at t' .

INTRA-CONVERSATIONAL DISAGREEMENT concerns the incompatibility of the parties' speech act contents_{SW}. In standard conversations in which people sincerely express their *beliefs*, they will disagree in their beliefs just in case they disagree with regard to the speech act contents_{SW} expressing those beliefs. But depending on the purpose of the conversation, speakers need not always *believe* what they assert and may still disagree.¹⁰ For instance, if the mutually recognised purpose of the conversation is to explore a hypothesis, conversational disagreement will be accompanied by disagreement in *suppositions*. If the purpose is to make up a fictional story, the attitude underlying

¹⁰Cf. section 4.7.1, footnote 20

speakers' presuppositions¹¹ is pretence and disagreement will go hand in hand with disagreement in pretence. So there can be intra-conversational disagreement without intra-conversational disagreement in belief. Likewise, there can be intra-conversational disagreement in belief without intra-conversational disagreement. Two parties to a conversation may hold relevant incompatible beliefs yet refrain from expressing them in speech acts. Thus we need to separate INTRA-CONVERSATIONAL DISAGREEMENT from intra-conversational disagreement in belief.

INTRA-CONVERSATIONAL DISAGREEMENT IN BELIEF

Two parties to a conversation, *A* and *B*, disagree in their beliefs with contents_{SW} p_A and q_B if p_A is incompatible with q_B .¹²

We also need to separate our accounts of intra- and inter-conversational disagreement. In section 5.4 on eavesdropping we saw that different contents are relevant to eavesdroppers and to conversational participants. This is because eavesdroppers do not take themselves to be part of the conversation, they are not engaged in the joint project of locating the conversational group. Rather, they are locating just themselves. The information they assess for truth is 'de-contextualised': An eavesdropper who spies on a conversation taking place hundreds of miles away will not reject an assertion of 'The money is here' simply because he thinks the money is not where he himself is. He should reject it only if the money is not in the location of the speaker. We saw that the Kaplan horizontal is just the piece of 'de-contextualised' information the eavesdropper assesses.

Eavesdropper who reject a claim, for instance by saying 'That's false', disagree with the speaker, as in (5.1) repeated here.

- (5.1) [Sal is secretly listening in on the following conversation between Ben and Anna.]
- | | |
|-------------------|---------------------------------------|
| Ben: | Liquorice is tasty. |
| Anna: | That's right. It is tasty. |
| Sal [to himself]: | That's false. Liquorice tastes awful. |

There are further cases of inter-conversational disagreement. Two speakers who are party to different conversations may be said to disagree, for instance when one asserts 'God exists' and the other asserts 'God doesn't exist.' Moreover, the same person may disagree with her past self. In MacFarlane's fish sticks example (cf. section 5.7.1),

¹¹ See the definition of SPEAKER PRESUPPOSITION_{SW} in section 4.7.1

¹² See footnote 9 in this chapter for an idealising assumption about time that INTRA-CONVERSATIONAL DISAGREEMENT IN BELIEF shares with INTRA-CONVERSATIONAL DISAGREEMENT.

MacFarlane disagrees with his childhood self about the tastiness of fish sticks.¹³ In all of these cases, it is the incompatibility of the speech acts' Kaplan horizontal that figures in the explanation of disagreement. (If we took the incompatibility of the diagonal contents to matter, we would get the wrong result that someone asserting 'It's sunny here' in Scotland would count as disagreeing with someone asserting 'It's not sunny here' in England.)

The fish sticks example also raises another important aspect. Horizontal content, just like diagonal content, is temporally neutral. But surely, I do not disagree with my former self simply because yesterday I said 'It's sunny' and now I say 'It's not sunny.' Rather, I disagree with my former self if I now say 'It wasn't sunny a day ago.' It is thus not the assertions' Kaplan horizontal whose incompatibility matters for disagreement. In the sequenced worlds account(s) of retraction, we got the right predictions by looking at the *temporal update* of the previous assertion's horizontal content (cf. sections 5.7.2 and 5.7.3). Analogously, in the sequenced worlds account of inter-conversational disagreement the temporal update of the horizontal content gives the right results.

INTER-CONVERSATIONAL DISAGREEMENT

Two speakers in different conversations at t_1 and t_2 disagree if the Kaplan horizontal of the speaker's assertion at t_2 is incompatible with the temporal update at t_2 of the Kaplan horizontal of the other speaker's assertion at t_1 .

To illustrate, INTER-CONVERSATIONAL DISAGREEMENT predicts that I disagree with my former self if I now say 'It wasn't sunny a day ago' because the temporal update of my former assertion's horizontal content, $SUNNY_{t_1}$, is incompatible with the horizontal content $SUNNY_{t_2}$ of my later assertion.¹⁴

¹³Some authors are sceptical of inter-conversational disagreement, however. For instance, Glanzberg (2007, 15) and Schaffer (2011, §4) reject clear judgments of disagreement in inter-conversational cases in which one speaker asserts 'X is tasty' and the other asserts 'X is not tasty.' According to Glanzberg, we only get clear judgments of disagreement if the two conversational contexts make it sufficiently clear that the speakers in both conversations talk about the same person or group to whom X is said to be tasty.

¹⁴By modifying the example slightly, we can see once more that the modal analysis of incompatibility as impossibility of joint truth (from the same perspective) is insufficient for accounts of disagreement. Suppose that in the above example, I say on day d_2 , the day after my first assertion on d_1 , 'It wasn't sunny yesterday.' By the Kaplan semantics for indexical expressions like 'yesterday', the horizontal content of this assertion is $SUNNY_{d_2}$:

(i) $SUNNY_{d_2}$: $\{\langle w, t, x \rangle$: it is not sunny on d_1 in $w\}$

It is not difficult to see that it is not impossible for $SUNNY_{d_2}$ and $SUNNY_{t_1}$ to be true from the same perspective. They are both true at $\langle w', t', x \rangle$, where w' is a world in which it is not sunny on day d_1 (making $SUNNY_{d_2}$ true) but in which it is sunny on day d_2 and t' is a moment during d_3 (making $SUNNY_{t_1}$ true). Hence $SUNNY_{d_2}$ and $SUNNY_{t_1}$ are not incompatible. But I clearly disagree with my former self in asserting the two contents on two consecutive days.

Note that the temporalist aspect of content is not the culprit. Similar examples arise for possible worlds content. For instance, suppose someone says 'The man in the red suit is drunk', and someone says about Sal, who is the man in the red suit, 'Sal's not drunk.' It is natural to report them as disagreeing. But on

(6.7) $SUNNY_{t_1}$: $\{\langle w, t, x \rangle$: it is sunny one day before t in $w\}$

(6.8) $SUNNY_{t_2}$: $\{\langle w, t, x \rangle$: it is not sunny one day before t in $w\}$

Inter-conversational disagreement in virtue of assertions must be distinguished from inter-conversational disagreement in virtue of belief. First, speakers may make assertions whose horizontal contents (modulo the temporal update) are incompatible yet may not believe what they say, for instance because the purpose of their respective conversations is to explore a hypothesis. They may still count as disagreeing.¹⁵ Second, two agents may disagree in belief even if they have not expressed their respective beliefs in speech.¹⁶ The account of inter-conversational disagreement in belief is as follows.

DISAGREEMENT IN BELIEF

Two agents at t_1 and t_2 , who are not party to the same conversation, disagree in belief if the Kaplan horizontal content of the agent's belief at t_2 is incompatible with the temporal update at t_2 of the Kaplan horizontal of the other agent's belief at t_1 .

There are thus four accounts of disagreement, covering disagreements in virtue of what speakers assert in the same conversation (INTRA-CONVERSATIONAL DISAGREEMENT), disagreements in virtue of the beliefs of speakers in the same conversation (INTRA-CONVERSATIONAL DISAGREEMENT IN BELIEF), disagreements in virtue of what speakers assert in different conversations (INTER-CONVERSATIONAL DISAGREEMENT), and disagreements in virtue of what agents who are not party to the same conversation believe (DISAGREEMENT IN BELIEF). Where speaker believe what they assert, intra-conversational accounts of disagreement and of disagreement in belief make the same predictions (and so do the inter-conversational accounts). But where speakers take a different attitude to what they assert, the two accounts, respectively, may make different predictions.

standard semantics for definite descriptions and proper names, it is possible for the contents expressed by these assertions to be true at the same world.

¹⁵It is an open question, I think, whether two speakers disagree when it is known that one believes what she asserts but the other makes her assertion in a conversation the goal of which is to establish common supposition. Perhaps the attitudes expressed in assertion have to be relevantly similar for the two speakers to count as disagreeing.

¹⁶There might be a difficulty about knowing what an agent believes if she has not voiced her belief. But this is irrelevant to the point. What matters is that agents can felicitously be reported to be disagreeing in virtue of conflicting beliefs (no matter what they say). For instance, if Anna believes that liquorice is tasty and Sal (who has never met nor talked to Anna) believes that liquorice is not tasty, then they can be said to disagree. (For a denial of this conditional and the view that disagreement requires the disagreeing agents to be party to the same conversation, see Schaffer (2011, §4.1).)

6.2.3 A Note on Transworld Disagreement

The two accounts INTER-CONVERSATIONAL DISAGREEMENT and DISAGREEMENT IN BELIEF have the consequence that agents in different possible worlds may be in disagreement.¹⁷ That is because the notion of incompatibility in the accounts relies on the relativist notion of truth from an assessor's perspective (cf. sections 5.2 and 5.3), on which contents are true or false relative to the assessor's *world*.¹⁸ Thus, if the horizontal content of some speaker in w_1 and the temporal update of the horizontal content of another speaker in w_2 cannot both be true relative to the (same) world of the assessor, the speakers count as disagreeing.

One may object that there are no transworld disagreements since, as MacFarlane puts it (borrowing Perry's (1986) terminology), the assertions or beliefs of agents in different worlds *concern* different worlds. MacFarlane (2007a, 23) motivates the rejection of transworld disagreement with the following example.

Consider Jane (who inhabits this world, the actual world) and June, her counterpart in another possible world. Jane asserts that Mars has two moons, and June denies this very proposition. Do they disagree? Not in any real way. Jane's assertion concerns our world, while June's concerns hers. If June lives in a world where Mars has three moons, her denial may be just as correct as Jane's assertion.

Contrary to MacFarlane, I think that our intuitive judgments about alleged cases of transworld disagreement do not in fact tell against transworld disagreement. Our judgments, I claim, can go either way, and whenever we have clear intuitive judgments, we follow counterfactual reasoning that construes the case as one of same-world disagreement. Let me explain.

There are two ways of presenting alleged cases of transworld disagreement. The first employs explicit talk of possible worlds, putting the subjects in different worlds as if they inhabited different continents. But relying on the language of modal realism in the presentation of cases has two significant disadvantages. First, it brings with it metaphysical commitments that few are prepared to accept. And second, it is doubtful whether we should trust our intuitive judgments based on modal realist talk, given

¹⁷INTRA-CONVERSATIONAL DISAGREEMENT does not have this consequence because the parties have to be part of the same conversation. I assume without argument here that one cannot be in a conversation with someone in a different possible world. In the sequenced worlds model, it is stipulated that conversations can be characterised in terms of a series of formal objects $\langle w, t, \langle x_1, \dots, x_n \rangle \rangle$, which are defined as triples of a world w , time t and individuals $x_1 \dots x_n$ *inhabiting the same world* w .

¹⁸Sequenced worlds relativism shares the aspect of world-perspectivalness with Egan's relativism but not with Lasersohn's, Stephenson's, or MacFarlane's relativism, on which propositions count as true iff they are true relative to the world and time of the speaker and the assessor/judge (or the standard of taste salient in the context of assessment). See section 5.3 for the relevant differences.

that modal realism is a controversial *theory* about the metaphysics of possibility and necessity. It is safer, then, to use language that is more entrenched in everyday use.¹⁹

The second way of presenting alleged cases of transworld disagreement uses standard counterfactual talk. Here judgments differ from case to case. In cases in which we are given information about the counterfactual assertion (or belief) as well as relevant ‘worldly’ information with respect to which the world of the counterfactual assertion differs from the actual world, judgments of no disagreement seem to prevail. In contrast, in cases in which the counterfactual world is not explicitly stipulated to be different in relevant ‘worldly’ facts, judgments of disagreement seem prominent.

Consider an example. Here is counterfactual scenario 1.

(6.9) RECESSION SCENARIO 1

In March 2012, Pier Carlo Padoanan, the OECD deputy secretary-general, said ‘The UK is back in recession.’²⁰ George Osborne, the UK’s Chancellor of the Exchequer, *could* have said (but did not), ‘The UK is not back in recession.’

Now ask yourself: Would Osborne have disagreed with Padoanan? I think one can plausibly answer yes and reason as follows: Osborne did not say it and so does not disagree with Padoanan. But had he said it, he would have disagreed with him.

Consider next counterfactual scenario 2.

(6.10) RECESSION SCENARIO 2

In March 2012, Pier Carlo Padoanan, the OECD deputy secretary-general, said ‘The UK is back in recession.’ Contrary to fact, the UK *could* be experiencing its strongest economic growth in recent years, and Osborne *could* have said (but did not), ‘The UK is not back in recession.’

Ask yourself again: Would Osborne have disagreed with Padoanan? I think in this scenario, it is much more natural to judge that they would not have disagreed.

What is important to note about these two cases is that insofar as we engage in counterfactual reasoning, our judgments are based on same-world construals of the cases. Standard accounts of counterfactual conditionals underscore this fact and support the difference in judgments. In each of the above scenarios, the described context

¹⁹MacFarlane (2007a, 23) maintains that alleged cases of transworld disagreement need not be presented in modal realist talk and provides an alternative presentation of the case using counterfactual talk. Cappelen and Hawthorne (2009, 63-6) object that this presentation does not yield intuitions of no-disagreement. In reply, MacFarlane (2012, 157-8) corrects his presentation of the case and makes appeal to the notion of a ‘counterfactual attitudes-in-context.’ I agree that this improves the presentation of cases, but I worry that pre-theoretical judgments about disagreeing with someone’s counterfactual attitude-in-context are hard to come by.

²⁰See www.guardian.co.uk/business/2012/mar/29/uk-back-in-recession-oecd

is one in which Pier Carlo Padoanan actually said ‘The UK is back in recession.’ In the first case, we then reason with the following counterfactual: ‘If George Osborne had said ‘The UK is not back in recession’, he would have disagreed with Padoanan.’ On a Lewis-Stalnaker theory, this counterfactual is true iff all the relevant worlds w that are most similar to the world of the context of use and in which Osborne made his assertion are such that he disagreed with Padoanan *in w*. The most similar worlds are all worlds in which Padoanan made his assertion, so they do turn out to disagree. In contrast, in the second case, we reason with the following counterfactual: ‘If the UK was experiencing its strongest economic growth in recent years and George Osborne had said ‘The UK is not back in recession’, he would have disagreed with Padoanan.’ This counterfactual is true iff all the relevant worlds w' that are most similar to the world of the context of use and in which the UK is going strong economically and Osborne made his assertion are such that he disagreed with Padoanan *in w'*. However, the relevant worlds that are most similar to the world of the context of use presumably are not all worlds in which Padoanan made his assertion, since in those worlds, he would have been incorrect and so may not have made his assertion. Hence in those worlds, he would not disagree with Osborne.

In sum, I believe that the only natural way of presenting alleged cases of transworld disagreement triggers judgments which are based on counterfactual reasoning that construes them as same-world disagreements. The variations in judgments about such cases can be explained by standard accounts of counterfactuals. Thus, our judgments regarding alleged cases of transworld disagreement provide no reason against disagreements between agents in different possible worlds.²¹

6.2.4 Faultlessness

Relativists like Kölbel (2002, 2004, 2009) and Lasersohn (2005, 2008, 2009) argue that relativism is uniquely positioned to explain the possibility of faultless disagreement about subjective matters. On their view, disagreement is explained by the fact that the propositions agents express or believe contradict each other; faultlessness is explained by the fact that each subject expresses or believes a proposition that is true relative to herself. In opposition to this case for relativism, however, some critics argue that

²¹If we wanted to avoid commitment to transworld disagreement, it would be easy enough to restrict the accounts to same-world cases of disagreement. The amended account of inter-conversational disagreement would be as follows.

SAME-WORLD INTER-CONVERSATIONAL DISAGREEMENT

Two speakers in different conversations at t_1 and t_2 in the same possible world disagree if the Kaplan horizontal of speaker's assertion at t_2 is incompatible with the temporal update at t_2 of the Kaplan horizontal of the other speaker's assertion at t_1 .

relativism can explain either disagreement or faultlessness, but not both.²²

I will show in this section that on the sequenced worlds view, one can make sense of the notion of faultless disagreement while at the same time acknowledging that there is more than one way in which disagreements fail to be faultless. However, in the intra-conversational cases of faultless disagreement that have attracted most attention, the notion of faultlessness on which disagreements about subjective matters can be said to be faultless is not the notion of truth relative to the speaker that relativists like Kölbel and Lasersohn have adduced. Thus the account of faultless disagreement offered here is not suited to support Kölbel's and Lasersohn's strong case for relativism. At the same time, it is not subject to the charge their explanation of faultless disagreement faces.

Note first that on the sequenced worlds view, the above explanation in terms of truth relative to the speaker is not available for all cases of disagreement. In particular, it fails for cases of INTRA-CONVERSATIONAL DISAGREEMENT and INTRA-CONVERSATIONAL DISAGREEMENT IN BELIEF. On these accounts, the relevant notion of perspectival truth is truth from the *conversational* perspective. An assertion of 'This cake is tasty' in a one-to-one conversation has the content $CAKE_{I\phi_2}$:

(6.3) $CAKE_{I\phi_2}$: $\{\langle w, t, \langle x_1, x_2 \rangle \rangle$: the contextually salient cake tastes good to $\langle x_1, x_2 \rangle$ in w at t

$CAKE_{I\phi_2}$ is true from the conversational perspective just in case the cake tastes good to both interlocutors at the time and in the world of the assertion. Its truth does not vary from one interlocutor to the other. So an explanation of faultlessness as truth from one's individual perspective is not available. If the second interlocutor asserts 'No, this cake isn't tasty', then one of them has asserted a content that is false from the conversational perspective and is thus at fault.

Not much is lost, however, for the sequenced worlds relativist who gives up the account of faultless disagreement in terms of truth relative to each speaker's perspective. As MacFarlane (2012, ch. 6) observes, 'faultlessness' can mean different things when ascribed to acts of asserting and believing. An act of assertion, for instance, may be faultless in at least any of the following ways:

1. $Faultless_1$: expressing an *epistemically warranted* content
2. $Faultless_2$: expressing a *true* content
3. $Faultless_3$: *complying with the norm(s) of assertion*

Here is what the sequenced worlds relativist can say about the multiple ways in which disagreements may or may not be faultless.

²²See for instance Moltmann (2010), Rosenkranz (2008), and Stojanovic (2007).

First, in the conversation in which Ben asserts ‘Liquorice is tasty’ and Anna asserts ‘Liquorice is not tasty,’ they may both be faultless₁ – epistemically warranted – in their assertions because they both have first-hand evidence of the taste of liquorice and they both have good reasons to believe that they are relevantly alike in their taste preferences, evidence that overrules any indication to the contrary that arises from the fact that they disagree. (Of course, they may lose these reasons at some point as the dispute continues.)

Second, Ben’s and Anna’s assertions in the same conversation can *not* both be faultless₂ – expressing a true content. We will discuss the ordinary language, monadic truth predicate in detail in section 6.5.1. For now, it will suffice to note that ‘true’ is ascribed of contents and is disquotational in the sense that it is governed by the following equivalence schema for contents: *The content that Φ is true if and only if Φ .* On the sequenced worlds view, this schema is true at context c and from a perspective just in case its right hand side is true at c and from that perspective iff its left hand side is true at c and from that perspective. Where a content varies in truth value from perspective to perspective, its ascriptions of monadic truth vary in truth from perspective to perspective, too. It is easy to see, then, that there is a perspective from which one of the assertions’ contents must be false. To an eavesdropper, for instance, liquorice will either taste good or fail to taste good. So at least one of the contents will fail to be true from the eavesdropper’s perspective. Hence, in assessing the assertions’ horizontal contents from her individual perspective, she cannot truly say of both contents that they are true.

The sequenced worlds relativist also has a theoretical notion of perspectival truth at her hands. For assertions made in the same conversation, the relevant perspective is that of the conversation. The contents of both Ben’s and Anna’s assertions cannot both be true from the conversational perspective. Whatever their respective tastes, it cannot both be the case that liquorice is tasty to them and that liquorice is not tasty to them. So their assertions cannot both be faultless in the sense of expressing contents that are both *true from the conversational perspective*.

For assertions made in different conversations, however, it is the speakers’ individual perspectives that matter. (Recall that eavesdroppers assess relative to their own individual perspective.) Speakers in different conversations who are said to disagree may both be faultless in the sense that their assertions’ horizontal contents are both *true from their respective individual perspectives*.

But Ben’s and Anna’s assertions, made in the same conversation, can *not* both be faultless in the sense of expressing a content that is true from the speaker’s perspective. This is because their speech act contents_{SW} are pair-centered contents, contents that are true or false from a pair-perspective but not from an individual perspective.

Third, their assertions may or may not both be faultless₃ – complying with the

norm(s) of assertion – depending on the norm of assertion. On the sequenced worlds, the two norms of assertion register two levels of appropriateness of an assertion: appropriateness given the group-centric conversational purpose of maximal cooperativeness and appropriateness given the speaker-centric purpose of voicing one's perspective and persuading one's audience of sharing one's perspective. The first level of appropriateness is captured by the STRONG, GROUP-CENTRIC NORM OF ASSERTION, which requires the speech act content_{SW} to be true from the conversational perspective. At most one of Ben's and Anna's speech act content_{SW} can be true from the conversational perspective. Hence, their assertions cannot both be faultless in the sense of *complying with the strong, group-centric norm of assertion*.

The second level of appropriateness is captured by the weak norm of assertion, which requires only that each speech act content_{SW} correctly locate the speaker and that the hearers be in a position to accommodate in such a way as to bring about the content's truth from the conversational perspective. Provided that both Ben and Anna can change their tastes, each of their assertions meets the weak norm of assertion and is hence faultless in the sense of *complying with the weak norm of assertion*.

The ways in which intra-conversational disagreements can and cannot be faultless on the sequenced worlds view are summarised in Table 6.1.

For inter-conversational disagreements, the picture is similar. However, as we noted above, truth from the speaker's individual perspective is relevant for inter-conversational disagreement where truth from the conversational perspective is relevant for intra-conversational disagreement. Moreover, the norms of assertion are not relevant for inter-conversational disagreements. Two speakers in different conversations may be judged to disagree, and may each be faultless, without any information being available on the conversational contexts that would allow judgments regarding the appropriateness of their assertions in their respective conversations. Table 6.1 gives an overview of faultlessness in inter-conversational disagreements.

For inter-conversational disagreement in belief, the norm of belief is egocentric. In chapter 4 we saw that one should believe a content only if it is true of oneself.²³ Two believers (in different contexts) can disagree and both be faultless in the sense that each of their beliefs is *complying with the egocentric belief norm*. A summary of the ways in which believers who are not party to the same conversation can and cannot faultlessly disagree in belief is given in Table 6.1.

There are thus various ways in which disagreements in each category may be said to be faultless. The paradigm case that has received most attention in arguments for and against faultless disagreement is intra-conversational disagreement. Here, the se-

²³The norm was stated in section 4.2.3 as follows:

EGOCENTRICITY Believe *p* only if *you yourself* are correctly located by *p*.

sequenced worlds view allows for a notion of faultless disagreement on which the disagreeing parties are faultless in the sense of complying with the weak norm of assertion. The resulting picture is nuanced, though. There are equally plausible ways of understanding faultlessness – even in terms of conversational propriety – on which disagreements about subjective matters turn out not to be faultless.

Table 6.1: Faultless Disagreement

	Conversational Disagreement	Inter-conversational Disagreement	Inter-conversational Disagreement in Belief
Epistemically warranted	✓	✓	✓
True	—	—	—
True from the conversational perspective	—	N.A.	N.A.
True from the agent's perspective	N.A.	✓	✓
Complying with the strong norm	—	N.A.	N.A.
Complying with the weak norm	✓	N.A.	N.A.
Complying with ego-centric belief norm	N.A.	N.A.	✓

6.2.5 Nonindexical Contextualism and Disagreement

What should nonindexical contextualists say about disagreement? Presumably, they want to make the same predictions that relativists make on the above accounts of disagreement. They can. All that nonindexical contextualists, on the sequenced worlds view, need to do is adapt the account of incompatibility that figures in the accounts of intra-conversational and inter-conversational disagreement (in belief). The notion of incompatibility received a modal analysis that made reference to the notion of truth. To make the same predictions as relativists, nonindexical contextualists need to replace the relativist's notion of truth from the same (conversational) perspective with their notion of truth in the same context.²⁴ The notion of incompatibility underlying the nonindexical contextualist accounts of disagreement (in belief) will then be as follows:

²⁴Recall the definition of TRUTH IN CONTEXT in section 5.2: A sentence Φ is true in the context c iff $\llbracket \Phi \rrbracket^{c, i_c} = 1$, where i_c is the index $\langle w_c, t_c, \langle x_c^1, \dots, x_c^n \rangle \rangle$ whose coordinates are determined by c .

Two sequenced worlds contents are incompatible iff they cannot both be true in the same context.

The accounts of intra- and inter-conversational disagreement (in belief) with the nonindexical contextualists' notion of incompatibility allow them to make the intuitively correct predictions about disagreement cases. However, the view is once more stuck with counterintuitive consequences that arise from the separation of utterance truth (more precisely: truth of sentences-in-context) and ordinary disquotational truth as ascribed to propositions. We have already seen that nonindexical contextualist accounts of eavesdropping and retraction make counterintuitive predictions. Making the right predictions about disagreement now has the price of predicting the following claims to be felicitous. They are about Ben's assertion of 'Liquorice is tasty' in one conversation and Anna's assertion of 'Liquorice isn't tasty' in another conversation.

(6.11) Ben and Anna disagree. Still, their assertions are both true.

(6.12) Anna: What Ben said is false and I disagree with him. But his assertion is true.

For the nonindexical contextualist who does not want to accept these consequences, there is a way to avoid them by changing the accounts of disagreement. Incompatibility can alternatively be defined as impossibility of truth in the contents' *respective* contexts of use. More precisely, two contents p_1 and p_2 , as asserted (believed) in c_1 and c_2 respectively, are incompatible iff it cannot be both that p_1 is true in c_1 and p_2 is true in c_2 . Equipped with accounts of disagreement based on this notion of incompatibility, nonindexical contextualists avoid predictions of felicity for (6.11) and (6.12). However, they do so at the cost of failing to predict disagreement in most cases of inter-conversational disagreement. As long as the contents expressed by Anna's and Ben's assertions are each true in the context in which they are each asserted, Anna and Ben are not predicted to be in disagreement. On this account of disagreement, it will be felicitous for Anna to assert (6.13).

(6.13) What Ben said is false, liquorice isn't tasty. But I don't disagree with him.

Nonindexical contextualism, on the sequenced worlds view, seems to be better off with the accounts of disagreement that give it the same empirical coverage as relativism. As with other data, however, the view will have to live with some implausible consequences, i.e. the felicity of assertions like (6.11) and (6.12).

6.3 On the Dynamics of (Dis)agreement

Disagreement as a state has received widespread attention in the literature on relativism and contextualism about subjective matters. Agreement and disagreement as a conver-

sational activity have received markedly less attention.²⁵ This is unfortunate. Getting clearer on (dis)agreement as a conversational activity promises to afford insights into the ways in which our perspective influences and is influenced by discourse and into what speakers may learn from a dispute about taste even where no agreement is reached.

The sequenced worlds view developed in this and the previous chapters has at its core a model of the common ground. This model is dynamic. It records the development of the discourse in terms of the contributions made to the common ground of shared background assumptions. The sequenced worlds model of common ground thus allows us to say more about the conversational dynamics of agreeing and disagreeing about subjective matters.

A quick review will be useful. The model, as developed in the preceding sections, consists of the *common ground* (or *context set*) – the information all speakers presuppose – as well as the individual speakers' commitments. Assertions are understood as proposals to add their content to the common ground. In asserting, speakers take on a *speaker commitment* to the truth of the asserted content. If accepted, this commitment becomes a *joint commitment* of all discourse participants; it becomes part of the common ground. If the proposal is not accepted, the speaker keeps the speaker commitment unless she implicitly or explicitly retracts her assertion. Finally, how we represent sequenced worlds content depends on the order of the conversational participants in the *conversational sequence* $\langle x_1, \dots, x_n \rangle$. To have a convenient term call the combination of the common ground (joint commitments), the individual speakers' commitments, and the conversational sequence the *conversational scoreboard*.

In what follows, I will illustrate how this sequenced worlds model illuminates the conversational dynamics of disputes about matters of personal taste. In particular, the model helps classifying disputes in terms of the different effects that they have on the conversational scoreboard and speakers' individual perspectives.

I will here not give a pretheoretical characterisation of what makes for a dispute about taste, but will rely on an intuitive grasp. A dispute may, but need not, involve the activity of disagreeing. It may, but need not, involve the use of explicit disagreement markers such as 'I disagree,' 'No,' 'That's false' and 'You're wrong'. We can define a dispute theoretically, as a conversational situation in which a proposal to add a content to the common ground is rejected by at least one interlocutor.

For all of the following cases, we stipulate that the conversational sequence is $\langle \text{Ben}, \text{Anna} \rangle$. To keep it simple, we will suppress the temporal update required over the duration of the conversation. This simplification will be harmless as none of the assertions are claims about short-lived states or events that terminate before the conversation ends.

²⁵Though see the integration of larger passages of conversation in the explanation of disagreement in the state sense in Schaffer (2011).

Let us start with a case of *Agreement by Persuasion*, which we have already encountered.

- (5.9) Ben: This liquorice and mint tea is tasty! I've never tried this combination before.
 Anna: No, it's not! The anise, minty and sweet flavours together are disgusting. And the aftertaste is the worst.
 Ben: Hmmmm. Let me have another sip then... You're right. I take that back. It's not tasty at all.

In (5.9), Anna disagrees with Ben's first assertion by making a counter-proposal to add a content to the common ground that is incompatible with the content of Ben's assertion about the tastiness of the tea. Ignoring their other contributions, at this point they have incompatible speaker commitments but no joint commitment regarding the tastiness of the tea. Eventually, they come to agree on the tastiness of the tea because Anna successfully persuades Ben to retract his previous claim and share her view, which she expressed by saying 'No, it's not.' Again ignoring other contributions to the common ground, the effect of the dialogue on the conversational scoreboard is that Anna and Ben have a joint commitment to the truth of the content in (6.14), and the contents of their assertions about the tastiness of the tea are removed from their respective lists of speaker commitments.

- (6.14) $\{\langle w, t, \langle x_1, x_2 \rangle \rangle$: the contextually salient liquorice and mint tea tastes good to $\langle x_1, x_2 \rangle$ in w at t

The next dialogue is a case of *Change of Perspective*.

- (6.15) Ben: This liquorice is tasty.
 Anna: Excuse me? You hate anise.
 Ben: Let me try another one then. ... You're right. It doesn't taste good to me at all.

In (6.15), Anna rejects Ben proposal to add the content *that this liquorice is tasty* to the common ground because she doubts Ben's sincerity. Since Ben retracts his first assertion by saying 'You're right' and concedes that the liquorice does not taste good to him, they end up agreeing. Regarding the tastiness of liquorice, the effect of the dialogue on the common ground is an update with (6.16).

- (6.16) $\{\langle w, t, \langle x_1, x_2 \rangle \rangle$: the contextually salient liquorice does not taste good to x_1 in w at t

Neither Ben nor Anna have any remaining speaker commitments. Note that Anna and Ben only agree on whether the liquorice tastes good to Ben. The dialogue, and the resulting common ground, do not establish anything regarding Anna's perspective on the taste of the liquorice. She might naturally continue the conversation with either of the following assertions, each of which is compatible with the common ground at the end of (6.15).

(6.17) Anna: I knew you don't like it. Anyway, I myself think it tastes quite good.

(6.18) Anna: I knew you don't like it. I can't stand it either.

Consider a case of *Suspension of Judgment*.

(6.19) Ben: Pffff... your coffee tastes awful.

Anna: Hold on. You just brushed your teeth. Of course you won't like coffee now.

Ben: Right. I take that back. I can't tell right now whether your coffee is tasty or not.

In (6.19), the dialogue does not result in any joint or speaker commitments regarding the tastiness of the coffee. What Anna achieves is that Ben retracts his claim and explicitly suspends judgment on the coffee's tastiness. Neither Anna nor Ben ends up with any commitments regarding the coffee's tastiness. Hence the conversational scoreboard leaves it open whether the coffee tastes good to any of them.

(6.20) is a case of *Unabated Disagreement*.

(6.20) Ben: This liquorice is tasty.

Anna: No, it tastes terrible!

Ben: But consider the alteration of sweet and salty sensations on the tongue.

Anna: That's exactly what's so disgusting. Also, the drops are too hard.

Ben: Really? Good liquorice must be hard.

Anna: ...

The dots at the end indicate that without either party retreating from their claims, the conversation either continues indefinitely, ends abruptly, or the topic is changed. In (6.20), no agreement regarding the liquorice's tastiness is reached. Hence, no joint commitment is incurred. However, since no party retracts their claims, they keep a speaker commitment to the truth of the content expressed by 'This liquorice is tasty' and 'It tastes terrible,' respectively. While they fail to agree, they come to learn something about their mutual perspectives. Provided that they assume that they are sincere and assume that they all assume that they are sincere etc., the information they thereby

acquire becomes part of the common ground. A pragmatic effect of their assertions is the update with (6.21) and (6.22).

(6.21) $\{\langle w, t, \langle x_1, x_2 \rangle \rangle$: the contextually salient liquorice tastes good to x_1 in w at t

(6.22) $\{\langle w, t, \langle x_1, x_2 \rangle \rangle$: the contextually salient liquorice does not taste good to x_2 in w at t

Unabated Disagreement must be distinguished from *Agreement to Disagree* in (6.23).

- (6.23) Ben: This liquorice is tasty.
 Anna: No, it tastes terrible!
 Ben: But ...
 Anna: Ben, it might be tasty to you, but it clearly isn't tasty to me.
 Ben: Ok, let's just agree that we don't like the same things.

In (6.23), Ben and Anna do not reach any agreement regarding the tastiness of the liquorice. And just like in (6.20), they also keep their speaker commitments and learn about each other's perspectives on liquorice. In contrast to (6.20), however, the content of Ben's final assertion is added to the common ground, which establishes an explicit agreement to accept that they have diverging perspectives.

In more fine-grained discourse models such as Farkas and Bruce's (2010), the difference between *Unabated Disagreement* and *Agreement to Disagree* can be explained more elegantly. Farkas and Bruce assume that the conversational scoreboard also has a 'Table' that keeps a record of the Question Under Discussion. As long as an item is on the Table, it remains to be resolved whether or not, roughly, the content is to be added to the common ground. To simplify, in cases of *Unabated Disagreement* the items 'Liquorice is tasty' and 'Liquorice is not tasty' remain on the Table. Cases of *Agreement to Disagree* result in the removal of the items from the Table without a change to the common ground (cf. Farkas and Bruce (2010, 102)).

The classification of disputes about taste in terms of their effects on the conversational scoreboard gives an indication of the ways in which pragmatic discourse models illuminate the dynamics of conversation. There is a variety of discourse models that work with richer conceptions of the conversational scoreboard. Many of these are based on the notion of common ground. It is thus possible to integrate the current model and its predictions for disputes about taste with more fine-grained accounts of the dynamics of discourse.²⁶

²⁶In particular, the model is compatible with the proposal by Farkas and Bruce (2010) (see footnote 42 in section 5.7.3). Discourse models based on a notion of common ground further include Asher and Lascarides (2003), Clark and Schaefer (1989), Clark (1992), Ginzburg (1996, 2012), Gunlogson (2003), and Roberts (1996, 2004). I have been first made aware of research on discourse models in pragmatics by

6.4 Believing Alike, Agreement, and Cognitive Significance

With the sequenced worlds account developed into a fuller picture, it is now time to take up a thread we left dangling in sections 4.2.3 and 4.6. Theories of mental states are in part designed to account for the state's cognitive significance: its connection to the agent's behavioural dispositions, its connection to other mental states (e.g. dispositions to draw certain inferences), and to dispositions to accept certain sentence types as true. This account will figure in the explanation and prediction of action. It explains why subjects who believe alike act alike (given similar background beliefs and desires). Consider Perry's example (ignore Perry's terminology):

When you and I entertain the sense of 'A bear is about to attack me,' we behave similarly. We both roll up in a ball and try to be as still as possible.
(Perry, 1977, 494)

Whatever a theory says about our cognitive states, it better carve out a similarity that figures in a psychological explanation of our similar actions. Likewise, this similarity must help explain why Hume and Heimson behave alike if they both believe of themselves that they are Hume. Believing alike in this sense is a matter of being in mental states with similar *cognitive significance*.

There is a second sense of believing alike. Perry continues:

When you and I both apprehend the thought that I am about to be attacked by a bear, we behave differently. I roll up in a ball, you run to get help.

In this case, we also believe alike: we both believe that I am about to be attacked by a bear. But as Perry points out, this similarity does not prompt us to behave alike. Consider Heimson and his psychiatrist:

Suppose Heimson manages to convince his psychiatrist that he is right, so that the psychiatrist also ascribes to Heimson the property of being Hume. Then Heimson and his psychiatrist share a common belief. Not in the sense in which Heimson and Hume do—the psychiatrist doesn't believe that he himself is Hume—but in another, equally legitimate sense.
(Lewis, 1979a, 537)

Tamina Stephenson's (2011) talk 'The Pragmatics of Judge Dependence.' Stephenson adapts Farkas and Bruce's model of discourse to make room for centered worlds content but does not address the dynamics of (dis)agreements about matters of taste in any detail.

Heimson and his psychiatrist believe alike without being disposed to similar behaviour. They agree that Heimson is Hume. Heimson and Hume do not agree about who is Hume. Heimson believes that he (Heimson) is Hume. Hume believed that he (Hume) is Hume and most likely had no beliefs about Heimson. Heimson's and his psychiatrist's believing alike is a matter of being in *agreement*.

A theory of doxastic attitudes should give us an account of both of these notions of believing alike or believing the same thing: having beliefs with similar cognitive significance and agreeing in belief. It should tell us what all those have in common who behave alike (given similar background beliefs and desires), and it should tell us what it is that subjects have in common who agree on something.

Here is how the sequenced worlds view of belief accounts for believing alike. As before, the sequenced worlds view distinguishes between *intra-conversational* and *inter-conversational* cases of believing alike. (The latter subsume cases of soliloquy and solitary belief.)

In inter-conversational cases, similarity in cognitive significance is explained by *sameness of diagonal content*. In section 4.11.1 we noted that agents in contexts of solitary belief and soliloquy entertain centered worlds contents. These contents correspond to the *diagonal content* of the sentences agents use to express their beliefs.²⁷ Thus, if Hume and Heimson both believe *that I am Hume*, they each have a *de se* belief with the diagonal content *HUME*. When you and I both believe *that I am about to be attacked by a bear*, we both have a *de se* belief with the diagonal content *BEAR*.

(6.24) *HUME*: $\{\langle w, t, x \rangle: x \text{ is Hume in } w \text{ at } t\}$

(6.25) *BEAR*: $\{\langle w, t, x \rangle: x \text{ is about to be attacked by a bear in } w \text{ at } t\}$

This gives us the following account of believing alike as sameness of cognitive significance for inter-conversational cases:

COGNITIVE SIGNIFICANCE (inter-conversational cases)

Two agents, who are not party to the same conversation, are in cognitively similar belief states if their beliefs have the same diagonal centered worlds content.

Believing alike as inter-conversational agreement in belief, on the sequenced world view, is explained with reference to horizontal content. We take our cue from INTER-CONVERSATIONAL DISAGREEMENT IN BELIEF in section 6.2: Heimson, who believes that he is Hume, and his psychiatrist, who agrees with Heimson that he (Heimson) is Hume, are in belief states with the same horizontal content *HEIMSON*.

(6.26) *HEIMSON*: $\{\langle w, t, x \rangle: \text{Heimson is Hume in } w \text{ at } t\}$

²⁷To get *centered* worlds content, the index needs to be a triple of a world, time, and a single individual.

For some cases in which subjects at different times are said to agree we will again have to take the temporal update of the horizontal content into account, so that we get the following account of believing alike as agreement.

AGREEMENT (inter-conversational cases)

Two agents at t_1 and t_2 , who are not party to the same conversation, agree in belief if the Kaplan horizontal content of the agent's belief at t_2 is identical to the temporal update at t_2 of the Kaplan horizontal of the other agent's belief at t_1 .

We only give a sufficient condition for believing alike as agreement because sameness of these two horizontal contents can account for many but not all cases. A temporal case is this: Ben, who believes at t_1 that it is sunny, and Anna, who believes at t_1 that it is sunny now, presumably agree that it is sunny. But their horizontal contents, given by our semantics for the sentences they would use to express their beliefs, are not identical. A modal case is this: I believe that the author of *A Treatise of Human Nature* was Scottish and you believe that David Hume was Scottish. Hence we agree. But on standard semantics for definite descriptions and proper names the horizontal contents of (the sentences we use to express) our beliefs are not exactly the same.²⁸

In intra-conversational cases, agreeing in belief is a matter of having beliefs with the same sequenced worlds content.

AGREEMENT (intra-conversational cases)

Two parties to a conversation agree in belief if their beliefs have the same sequenced worlds content.

Again, we only give a sufficient condition for agreement because not all cases of agreement are ones in which believers have the same sequenced worlds belief content. Consider Heimson talking to his psychiatrist. Heimson believes *de se* that he is Hume. His psychiatrist, having been convinced by Heimson, believes of Heimson that he is Hume, a belief he would express by saying 'Heimson is Hume.' Thus, they agree. But they do not share a sequenced worlds speech act content (unless Heimson also believes *that Heimson is Hume* or his psychiatrist also believes *that you are Hume*). Given the conversational context $\langle w_c, t_c, \langle \text{Heimson, psychiatrist} \rangle \rangle$, Heimson believes $HUME_t$ and his psychiatrist believes $HEIMSON'$.

(6.27) $HUME_t$: $\{ \langle w, t, \langle x_1, x_2 \rangle \rangle : x_1 \text{ is Hume in } w \text{ at } t \}$

(6.28) $HEIMSON'$: $\{ \langle w, t, \langle x_1, x_2 \rangle \rangle : \text{Heimson is Hume in } w \text{ at } t \}$

²⁸See also footnote 14 in section 6.2.2.

A simple fix for this case would be taking the horizontal content as the object of agreement, as in inter-conversational cases. But an account in terms of shared horizontal content does itself not provide a wide enough notion of agreement. It too fails to predict agreement in all cases where two believers agree, as we have seen with the account of inter-conversational agreement. Hence, we must content ourselves with a sufficient conditions account of agreement and note that it does not predict agreement in all cases where agents can plausibly be said to agree. Short of providing a full, necessary and sufficient conditions analysis of agreement, we will stick with the above account that gives center stage to $\text{SPEECH ACT CONTENT}_{\text{SW}}$ – the diagonal content of interlocutors' mental states and the information that gets passed from speaker to hearers.

Let us finally turn to believing alike as similarity in cognitive significance for intra-conversational cases. We observed in chapter 4, section 4.6, that believing the same sequenced worlds content does *not* make for beliefs with similar cognitive significance. If Lingens, in the conversational context $\langle w_c, t_c, \langle \text{Lingens, Ortcutt} \rangle \rangle$, believes $TIRED_I$,

$$(6.29) \quad TIRED_I: \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : x_1 \text{ is tired of reading in } w \text{ at } t \},$$

and Ortcutt believes $TIRED_I$, they are not disposed to behave alike. To account for the difference in their doxastic states, we introduced the notion of n -BELIEVING IN A CONVERSATIONAL CONTEXT. Lingens 1-believes $TIRED_I$, Ortcutt 2-believes $TIRED_I$. Intuitively, Lingens self-and-group-locates in the content in a different way than Ortcutt.

What, then, is believing alike as similarity in cognitive significance in intra-conversational cases? It is easy to see this in a simple example. Lingens and Ortcutt believe alike if they both believe *that I am tired of reading*. In the conversation, this is for Lingens to 1-believe $TIRED_I$ and for Ortcutt to 2-believe $TIRED_2$.

$$(6.30) \quad TIRED_2: \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : x_2 \text{ is tired of reading in } w \text{ at } t \}$$

$TIRED_2$ is the 2-for-1 inverse of $TIRED_I$ – $TIRED_I$ but with the first and second center of each sequenced world swapped around.²⁹ The account of believing alike as similarity in cognitive significance in intra-conversational cases is as follows.

COGNITIVE SIGNIFICANCE (intra-conversational cases)

In a conversational context $\langle w, t, \langle x_1, \dots, x_n \rangle \rangle$, x_i and x_j are in cognitively similar belief states if there is a sequenced worlds content such that x_i i -believes it and x_j j -believes its i -for- j inverse.

²⁹For the notion of an inverse, see also section 4.9

This account must be given with a proviso, however. It provides us only with cognitively *similar* belief states. Lingens and Ortcutt are not in exactly the same belief states, according to COGNITIVE SIGNIFICANCE, when Lingens 1-believes $TIREDD_1$ and Ortcutt 2-believes $TIREDD_2$. To see this, it will be helpful to look again at the equivalence between sequenced worlds belief and centered worlds belief that we stated in chapter 4.

SEQUENCED WORLDS BELIEF AND CENTERED WORLDS BELIEF

An agent A n -believes a sequenced worlds content $\{\langle w, t, \langle x_1, \dots, x_u \rangle \rangle: p(w, t, \langle x_1, \dots, x_u \rangle)\}$ in the conversational context $\langle w_c, t_c, \langle y_1, \dots, y_u \rangle \rangle$ iff

- (i) $A = y_n \in \{y_1, \dots, y_u\}$
- (ii) there are relations $R_1 \dots R_u$ such that in w_c at t_c , y_n is uniquely R_1 -related to y_1, \dots , and y_n is uniquely R_u -related to y_u (where R_n is the identity relation) and y_n 's standing in $R_1 \dots R_u$ to $y_1 \dots y_u$ establishes a conversation between $y_1 \dots y_u$
- (iii) A believes the centered worlds content $\{\langle w, t, x \rangle: \text{there are individuals } x_1, \dots, x_u \text{ such that } x \text{ is uniquely } R_1\text{-related to } x_1, \dots, \text{ and } x \text{ is uniquely } R_u\text{-related to } x_u \text{ in } w \text{ at } t, \text{ and } p(w, t, \langle x_1, \dots, x_u \rangle)\}$.

As we apply SEQUENCED WORLDS BELIEF AND CENTERED WORLDS BELIEF to Lingens' 1-believing $TIREDD_1$ and Ortcutt's 2-believing the content got from swapping the 1st and 2nd centers in $TIREDD_1$, we can see that Lingens and Ortcutt do not believe *exactly* the same centered worlds contents – $TIREDD^L$ and $TIREDD^O$ – since the relation that establishes a conversation between them need not be symmetric.

(6.31) $TIREDD^L$: $\{\langle w, t, x \rangle: \text{there are individuals } x_1, x_2 \text{ such that } x = x_1, x \text{ is addressing } x_2 \text{ and } x_1 \text{ is tired of reading in } w \text{ at } t\}$

(6.32) $TIREDD^O$: $\{\langle w, t, x \rangle: \text{there are individuals } x_1, x_2 \text{ such that } x = x_1, x \text{ is addressed by } x_2 \text{ and } x_1 \text{ is tired of reading in } w \text{ at } t\}$

The account COGNITIVE SIGNIFICANCE ignores these differences. As a result, it can at best be an account of believing alike as *similarity* in cognitive significance – an account of believing a similar thing.

The discussion shows that sameness of content, in intra-conversational cases, does not make for believing alike as similarity in cognitive significance. One might object that the sequenced worlds view thus loses the motivation of Lewis' (1979a) account of attitudes *de se*, which is to characterise states of the head in terms of the kind of attitude

and its object. Two agents are in the same belief state in case their beliefs have the same object. On the contrary, the sequenced worlds view captures cognitive significance in intra-conversational cases in terms of the object of the attitude (content), the kind of attitude (belief), *and the way the subject is related to the content by that attitude* (cf. *n*-believing).³⁰ On the view, content ‘ain’t in the head.’

I think that once the main motivation for the Lewisian view is properly understood, the objection loses its bite. The motivation is ‘to characterize states of the head; to specify their causal roles with respect to behavior, stimuli, and one another’ (Lewis, 1979a, 526). Lewis thinks this characterisation is best done in terms of the object of the attitude (content). But this seems to be a mere preference:

Mean what you will by ‘object of an attitude.’ But if you mean something that is not determined by the state of the head, and that cannot do the job of characterizing states of the head by their causal roles, then I think you had better introduce something else that can do that job. I would prefer to reserve the term ‘object of an attitude’ for that something else. (Lewis, 1979a, 526)

So whence the preference for capturing cognitive significance in terms of the attitudinal object, i.e. the content? One advantage that objects of attitudes may have over that ‘something else’ is that they can smoothly be integrated into rigorous formalisations, e.g. in Bayesian decision theory, of the common-sense psychology that describes the causal roles of states of the head (cf. Lewis (1979a, 533-4)). I take it that this is what underlies Lewis’ motivation for characterising the cognitive significance, or causal role, of mental states in terms of the *objects* of attitudes. But if this is the motivation for the self-location picture, the sequenced worlds view can happily be embraced without loss. For the notion of *n*-believing a sequenced worlds content in a conversational context can be translated into believing a centered worlds content, by SEQUENCED WORLDS BELIEF AND CENTERED WORLDS BELIEF. The resulting centered worlds contents can then play the role that Lewis intended them to play in Bayesian decision theory. The cost is only that sameness of the *object* of the attitudes – sequenced worlds contents – is not what accounts for similarity in the states’ cognitive significance. But this seems a price worth paying if the objects of the attitudes are freed from the realm of unsharable, incommunicable mental entities.

6.5 Truth Ascriptions, Tense, and Predicates of Personal Taste

On the sequenced worlds view, horizontal and diagonal content each play an important explanatory role. But with two kinds of content in the theory, we may ask: Which

³⁰Thanks to Stephan Torre and François Recanati for pressing me on this point.

of the two contents is the object of ordinary truth ascriptions? In answering this question, I will first need to state the relativist and nonindexical contextualist accounts of the monadic predicate ‘true’ (§6.5.1). I will then look at inter-conversational truth ascriptions and, in particular, at cross-temporal assessments for truth. As we will see, sequenced worlds relativism includes relativism not only about taste but also about time and worlds. In sections 6.5.2 and 6.5.3, I will explore some of the consequences for cross-temporal truth ascriptions and the role of tense, and I will defend sequenced worlds relativism about predicates of personal taste from some objections raised by MacFarlane.

6.5.1 Monadic ‘true’

Let us start by reviewing the relativist’s and the nonindexical contextualist’s accounts of the object-language, monadic predicate ‘true’ on the sequenced worlds framework. We will see that the nonindexical contextualist view has some unhappy consequences. These will be important to keep in mind for the next section on cross-temporal truth ascriptions and tense. Before turning to temporal matters, however, we need to answer the question about the objects of ordinary truth ascriptions. I will show that in intra-conversational cases, diagonal content is the object of ordinary truth ascriptions, and in inter-conversational cases, it is horizontal content.

It is plausible that the ordinary language, monadic predicate ‘true’ is attributed to contents. And doubtlessly ‘true’ is disquotational in the sense that the following equivalence schema holds.

EQUIVALENCE The content that Φ is true iff Φ .

Relativists and nonindexical contextualists can both give a semantics of ‘true’ on which EQUIVALENCE has only true instances. On the sequenced worlds view the semantics is as follows.

$$(6.33) \llbracket \text{true} \rrbracket^{c,i} = \lambda q. q(i) = \mathbf{1}, \text{ where ‘} q \text{’ is a variable over contents}^{31}$$

$$(6.34) \llbracket \ulcorner \text{that } \Phi \urcorner \rrbracket^{c,i} = \llbracket \ulcorner \text{the content that } \Phi \urcorner \rrbracket^{c,i} = \lambda \langle w, t, \langle x_1, \dots, x_n \rangle \rangle. \llbracket \Phi \rrbracket^{c,i} \quad 32$$

The semantic clause for ‘true’ says that ‘true’, as used at a context c , takes a sequenced worlds content and maps it to truth just in case the sequenced worlds content is true at

³¹‘ q ’ is thus a variable over functions of the semantic type $\langle s, \langle i, \langle e \dots \langle e, t \rangle \dots \rangle \rangle$. e is the semantic type of individuals, s stands for worlds, i for times (not to be confused with the index i in (6.33)), and t for truth values; combinations such as et , or $\langle e, t \rangle$, stand for functions from the first type (e) to the second (t). The dots above indicate that the function may take different numbers of arguments of type e , depending on the length of the sequence.

³²The assumption that $\ulcorner \text{that } \Phi \urcorner$ and $\ulcorner \text{the content that } \Phi \urcorner$ have the same semantic entry is here made merely for simplicity.

c and i . In the expression \lceil that Φ \rceil , ‘that’ combines with a sentence; the extension of the whole expressions is the sentence’s intension (a content). Given the relativist definition of TRUTH AT A CONTEXT AND FROM AN ASSESSOR’S PERSPECTIVE, for any context c and perspective P , the equivalence schema is true at c and from P just in case ‘the content that Φ is true’ is true at c and from P iff ‘ Φ ’ is true at c and from P . The latter biconditional is necessarily true since the semantic value of ‘true’ maps a content to truth at c and the index determined by P just in case the content is true at c and the index determined by P . From a perspective P , one can truly say ‘It’s true that liquorice is tasty’ just in case one can truly say from P , ‘Liquorice is tasty.’

Nonindexical contextualists can endorse the semantics for ‘true’ in (6.33). Given nonindexical contextualism’s definition of TRUTH IN CONTEXT, the equivalence schema has only true instances because the left hand side is true in any context c just in case the right hand side is true in c , which is guaranteed by the semantics of ‘true.’

As we have seen in section 5.2, relativism and nonindexical contextualism come apart in their accounts of assertion truth – more precisely, truth of sentences-in-context. For convenience, and to facilitate the use of data, I will in this section speak of the truth of assertions and will assume that the notion can be identified with truth of sentences-in-context. I should add, however, that I share the well-known doubts both about assertions, which as speech acts are a form of action, being bearers of truth and falsity and about assertion truth being identified with the truth of sentences-in-context.³³

In addition to EQUIVALENCE, sequenced worlds relativists can endorse the following plausible principles linking content truth and assertion truth/truth of sentences-in-context.

(6.35) An assertion with content p is true iff the content p is true.

(6.36) If a (declarative) sentence Φ in context c has content p , then Φ in c is true iff the content p is true.

These principles ensure that assertion truth and content truth go hand in hand. Whenever one can truly say of an assertion that it is true, one can say of its content that it is true, and vice versa.

Nonindexical contextualists, however, cannot endorse (6.35) and (6.36) together with EQUIVALENCE. For suppose Anna, to whom liquorice tastes good, asserts ‘Liquorice is tasty.’ By TRUTH IN CONTEXT, it follows that Anna’s assertion is true. Anna’s assertion expresses the content that liquorice is tasty. Call this content p . By (6.35), it follows that p is true. But suppose also that Ben, to whom liquorice does not taste good, also asserts ‘Liquorice is tasty.’ (He might be wrong about his own taste.) By TRUTH

³³Doubts about assertions being bearers of truth and falsity are voiced, for instance, by MacFarlane (2012, 66-7), doubts about assertion truth being identified with the truth of sentences-in-context by Kaplan (1989, 522).

IN CONTEXT, his assertion is not true. It also expresses p . Thus by (6.35), it follows that p is false. Contradiction.³⁴ The obvious solution for nonindexical contextualists is to give up the link between content truth and assertion truth established in (6.35) and (6.36) and to replace the principles with (6.37) and (6.38).

(6.37) An assertion with content p in context c is true iff p is true at the index determined by c .

(6.38) If a (declarative) sentence Φ in context c has content p , then Φ in c is true iff the content p is true at the index determined by c .

(6.37) and (6.38), together with EQUIVALENCE, have the implausible consequences for nonindexical contextualism that we have noted several times.³⁵ Suppose Sal, to whom liquorice does not taste good, overhears Anna's assertion of 'Liquorice is tasty.' Given EQUIVALENCE and the fact that the content expressed by the assertion is false from Sal's individual perspective, he is then correct to assert 'What Anna said is false.' However, by TRUTH IN CONTEXT (identifying assertion truth with truth of sentences-in-context), Anna's assertion is true and Sal can correctly say 'Anna's assertion is true.' ('True', here, is the predicate defined by TRUTH IN CONTEXT, which is not a disquotational.) Hence Sal is correct in asserting (6.39) and (6.40).

(6.39) What Anna said is false. But her assertion is true.

(6.40) Anna's assertion is true. But still, what she said is false.

As noted in section 5.4, nonindexical contextualists may argue that assertion truth, just like truth of sentences-in-context, is a technical notion of truth, not one used by ordinary speakers. Truth ascriptions by ordinary speakers are always ascriptions of truth to contents. Thus one need not take sentences like (6.39) and (6.40) seriously as empirical data. Rather they are consequences that semanticists have to, and can, live with.

However, this way of dealing with consequences like (6.39) and (6.40) makes the nonindexical contextualist view about PPTs and ordinary truth ascriptions empirically indistinguishable from relativism. I have throughout this chapter and the previous chapter pointed out that sequenced worlds nonindexical contextualists can always adapt features of their view to match relativism's empirical predictions. If they want to establish a theoretical alternative with empirical bite, however, nonindexical contextualists will have to accept that the infelicity of assertions of (6.39) and (6.40) counts as empirical data disconfirming their view.

³⁴This argument, and its consequences for nonindexical contextualism, can be found in a slightly different form in MacFarlane (2009, 246-7).

³⁵See sections 5.4, 5.7.4 and 6.2.5 above.

Let us now turn to the objects of ordinary truth ascriptions. Ordinary monadic truth is ascribed of contents. But on the sequenced worlds view, every assertion and every belief has two contents – horizontal content and diagonal content. These differ in truth-conditions. In consequence, it matters which content truth is ascribed of on a given occasion.

We have seen a pattern emerge in this and the previous chapter. When pragmatic phenomena such as eavesdropping, retraction, and disagreement are explained, inter-conversational cases are distinguished from intra-conversational cases. In intra-conversational cases, it is diagonal content that figures in the explanation of the phenomenon. In explanations of inter-conversational cases, it is horizontal content. It is thus no surprise that for ascriptions of monadic ‘true’, the same distinction is relevant. If a speech act content_{SW} proffered in one’s conversation is up for evaluation, devices of propositional reference such as ‘that’ in ‘that’s true,’ or ‘what you said’ in ‘what you said is true,’ pick out diagonal speech act content_{SW}. In normal conversations, in which the contents in the common ground are assumed to be true, explicit evaluations of asserted contents go hand in hand with acceptance or rejection of the proffered content. It would be positively bizarre to respond to an assertion by saying, ‘That’s false,’ and subsequently take the proffered information for granted in one’s speech acts. Acceptance/rejection and judgments of truth/falsity can line up in this way because they have the same object: diagonal speech act content_{SW}.

In inter-conversational cases, where ‘that’ picks out the content of an assertion that is made in a different conversation, or the content of a belief by someone not present in the current conversation, it is the horizontal content of the assertion or belief that is the object of truth evaluations. This is as it should be. In evaluating my own assertion of ‘I am hungry’ for truth, I am correct to assess the diagonal content *HUNGRY_D* from my perspective.

$$(6.41) \text{ } HUNGRY_D: \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : x_1 \text{ is hungry in } w \text{ at } t \}$$

However, I would be wrong in assessing *HUNGRY_D* for truth if I were eavesdropping on someone else’s assertion of ‘I am hungry.’ In that case, I recognise that ‘I’ refers to someone else – the speaker – and I judge the content true or false depending on the speaker’s appetite. I do so if I judge the horizontal content *HUNGRY_H*.

$$(6.42) \text{ } HUNGRY_H: \{ \langle w, t, \langle x_1, x_2 \rangle \rangle : s_c, \text{ the speaker in } c, \text{ is hungry in } w \text{ at } t \}$$

6.5.2 Cross-temporal Truth Ascriptions

In section 5.3, we saw that there are at least two ways of implementing relativism about matters of personal taste. According to the first option, centered worlds relativism,

contents are sets of centered worlds. The semantics' index is a triple $\langle w, t, x \rangle$ consisting of a world, time, and individual. In evaluations of contents for truth and falsity, all of the coordinates in the index are determined by the assessor's location – the index just is the centered world representing the location of the assessor. Sequenced worlds relativism is akin to centered worlds relativism, but substitutes sequenced worlds for centered worlds.

According to the second option, most prominently advocated by MacFarlane, contents are sets of triples $\langle w, t, s \rangle$ of a world, time, and standard of taste; the semantics' index is a triple $\langle w, t, s \rangle$; in evaluations of contents in context for truth and falsity, the world and time coordinates in the index are determined by the context of use – for instance, the context in which the sentence whose content is getting evaluated was used – whereas the standard of taste is determined by the context of assessment. On MacFarlane's relativism, truth of sentences-in-context is defined as follows.

MACFARLANE TRUTH

A sentence Φ as used in a context of use c_U and assessed at a context of assessment c_A is true iff the content expressed by Φ at c_U is true at $\langle w_{c_U}, t_{c_U}, s_{c_A} \rangle$.

MacFarlane-style relativism, based on MACFARLANE TRUTH, is a form of nonindexical contextualism about modality (worlds) and time. Recall from section 5.2 that (truth) relativism about X and nonindexical contextualism about X need to differ only in their definitions of truth. For the relativist, the feature in the index that matters for X is determined by the assessor's context. For the nonindexical contextualist, this feature is determined by the context of use. In MACFARLANE TRUTH, the world and time coordinates are determined by the context of use.

MacFarlane-style relativism must thus reject the plausible principles (6.35) and (6.36) linking assertion and content truth because of times and worlds. It encounters similar consequences regarding world and time that nonindexical contextualism about taste encounters for taste. Suppose Sal says at noon, 'It's noon.' His assertion is true by MACFARLANE TRUTH. But suppose an hour later, he remembers his assertion. Given EQUIVALENCE, he can then felicitously utter 'That's false.' He can do so because the temporally neutral content of his earlier assertion is not true at Sal's current context of use at 1pm. Yet Sal can also felicitously say at 1pm, 'My earlier assertion is true.' The same implausible juxtapositions of assertion and content truth are thus felicitous:

(6.43) Sal: What I said is false. But my assertion is true.

(6.44) Sal: My assertion is true. But still, what I said is false.

Predictions of felicity for (6.43) and (6.44) are avoided by the first implementation of relativism, centered worlds relativism, and by sequenced worlds relativism. Note the

explicit definition of truth of sentences in context, here given in MacFarlane's terminology of context of use and context of assessment.

EGAN TRUTH

A sentence Φ as used in a context of use c_U and assessed by individual x at context of assessment c_A is true iff the content expressed by Φ at c_U is true at $\langle w_{c_A}, t_{c_A}, x_{c_A} \rangle$.

EGAN TRUTH makes assertion truth go hand in hand with content truth and thus preserves principles (6.37) and (6.38) linking the two notions. Together with EQUIVALENCE for ordinary truth ascriptions, it predicts that Sal can assert (6.45) at 1pm.

(6.45) Sal: My assertion is false, because what I said is false: It's not noon.

It may strike the impartial reader that relativism about time is a step in the wrong direction. For although it avoids the divergence in ascriptions of assertion truth and content truth, it retains the rather implausible consequence that Sal can felicitously call what he earlier said false, despite its being true at the time of asserting. It may be pointed out that the accounts of retraction and disagreement add to the unpalatable consequences. The following claims by Sal and Anna at 1pm are predicted to be felicitous.

(6.46) Anna: What you said is false (now), but I don't disagree with you.

(6.47) Sal: What I said is false (now), but I won't take it back.

At the same time, relativism about time also has intuitive appeal. Egan remarks:

[I]t can be appropriate for listeners at different times to attribute different truth values to one and the same utterance. One possible example of this is when Blofeld says, at noon, 'it's lunchtime.' Number 2, in the room at the time of utterance, should agree that Blofeld's utterance is true. Bond, listening to the recorded conversation later that night, should probably—though I think the intuitions here are less clear—say that Blofeld's utterance is false (I don't think Bond should say that it *was* false—only that it *is* false, at the time when he's evaluating it. It would, for example, be appropriate for Bond to say 'that's false' on hearing Blofeld's recorded utterance. Certainly, it would be appropriate for him to respond 'no' if Felix asked him, 'is that true?' or 'is what Blofeld said true?'). (2007, 7 n.10)

If Egan is right, then there is intuitive appeal to the claim that (present-tense) cross-temporal assessments of assertions and contents for truth are made with respect to the assessor's *current* time. To give another illustration, Sal and Anna are listening at noon to a podcast of the morning news. The newsreader correctly announces that the truce between Syria's regime and opposition forces is still being observed. Anna, in possession of more recent information to the effect that fighting has been resumed, says in response to the announcement, 'That's false. They've started fighting again.' Not only does her remark seem apt, it also easy to see how remarking 'That's true' could be misleading for Sal, if Sal knew that Anna had more recent information.

The passage by Egan also suggests that truth ascriptions are tensed. Thus, while it is felicitous for Anna to say, 'That's false,' she would be wrong in claiming that what the newsreader announced *was then* false. If Anna wanted to comment more comprehensively on the situation, she could perfectly say 'That's false. But of course, it was true when the news were read. So the newsreader isn't wrong.' Relativism about time is naturally combined with the following principles governing the tense of truth ascriptions.

(6.48) An assertion with the content p made in the past was true iff the content p was then true.

(6.49) An assertion with the content p made in the future will be true iff the content p will then be true.

The tense in truth ascriptions also alleviates some of the oddity in our earlier examples (6.46) and (6.47). Consider Anna's and Sal's rather natural reactions at 1pm to Anna's assertion of 'It's noon' at noon.

(6.46') Anna: What you said is false now. But it was true when you said it, so I don't disagree with you.

(6.47') Sal: What I said is false (now). But it was true when I said it so I won't take it back.

These are some of the considerations that speak in favour of temporally neutral contents, relativism about time and the significance of tense in truth ascriptions. No doubt examples can be found in which 'That's true' is naturally asserted in response to a temporally neutral assertion made in the past. My intention here has not been to resolve all problems for temporally neutral contents, but rather to draw out, and make plausible, the consequences of assuming temporally neutral contents and an ordinary truth predicate governed by EQUIVALENCE, and to argue in favour of centered worlds/sequenced worlds relativism over MacFarlane-style relativism on these assumptions.

6.5.3 Tense and Predicates of Personal Taste

MacFarlane (2012) offers some empirical considerations in favour of his way of implementing relativism about taste that draw on the role of tense in claims of tastiness. To see what is at stake, let us first make the assumption, common to both kinds of relativism, that tense markers express intensional operators that shift the time coordinate in the index. Following this tradition in intensional semantics will be convenient and in agreement with the semantic system developed so far. But nothing hangs on this choice. The following discussion could equally treat tense markers as object language quantifiers.³⁶ For simplification, we will introduce three tense markers ‘PAST’, ‘PRESENT’ and ‘FUTURE’ as stand-ins for English ‘-ed’ (for simple past tense, ignoring aspect), ‘-Ø’ (for present tense), and the modal construction with ‘will’ expressing future tense in English.

$$(6.50) \llbracket \text{PAST } \Phi \rrbracket^{c, \langle w, t, \langle x_1, \dots, x_n \rangle \rangle} = \mathbf{I} \text{ iff } \exists t' \text{ such that } t' < t \text{ and } \llbracket \Phi \rrbracket^{c, \langle w, t', \langle x_1, \dots, x_n \rangle \rangle} = \mathbf{I}$$

$$(6.51) \llbracket \text{PRESENT } \Phi \rrbracket^{c, \langle w, t, \langle x_1, \dots, x_n \rangle \rangle} = \mathbf{I} \text{ iff } \llbracket \Phi \rrbracket^{c, \langle w, t, \langle x_1, \dots, x_n \rangle \rangle} = \mathbf{I}$$

$$(6.52) \llbracket \text{FUTURE } \Phi \rrbracket^{c, \langle w, t, \langle x_1, \dots, x_n \rangle \rangle} = \mathbf{I} \text{ iff } \exists t' \text{ such that } t' > t \text{ and } \llbracket \Phi \rrbracket^{c, \langle w, t', \langle x_1, \dots, x_n \rangle \rangle} = \mathbf{I}$$

$$(6.53) \llbracket \text{One year ago } \Phi \rrbracket^{c, \langle w, t, \langle x_1, \dots, x_n \rangle \rangle} = \mathbf{I} \text{ iff } \llbracket \Phi \rrbracket^{c, \langle w, t', \langle x_1, \dots, x_n \rangle \rangle} = \mathbf{I}, \text{ where } t' = t - 1 \text{ year}$$

MacFarlane’s arguments in favour of his relativism explicitly target Stephenson’s (2007a) and Lasersohn’s (2005) systems, but most of them also apply to Egan’s centered worlds relativism and the sequenced worlds relativism defended here. Where the details matter, I will point out the differences. Otherwise I will use centered worlds relativism for illustration. MacFarlane (2012, 189–90) points out that in evaluations of claims involving first-personal, non-relativized uses of PPTs, the time coordinate of the index ‘plays a double role’ on Lasersohn’s and Stephenson’s relativism (as well as on centered and sequenced worlds relativism):

It tells us not only what time-slice of the object to look at (which is important because objects can change their flavors over time), but also what time-slice of the judge to look at (which is important because judges can change their tastes over time).

³⁶For quantificational and pronominal treatments favoured in linguistic semantics, see for instance Partee (1973), Enç (1986), and Kusumoto (1999, 2005). For a philosophically oriented discussion, see Glanzberg (2011). King (2003) argues from quantificational treatments of tense to eternalist semantic content, but Ninan (2010a, forthcoming) convincingly shows that quantificational treatments of tense are compatible with temporally neutral semantic contents.

Given EGAN TRUTH, ‘This cookie is tasty’ as used in c_U and assessed by x in c_A is true just in case $\llbracket \text{This cookie is tasty} \rrbracket^{c_U, \langle w_{c_A}, t_{c_A}, x_{c_A} \rangle} = 1$ – that is, in case the time-slice of the cookie (which is salient in c_U) *at the time of* c_A tastes good to the time slice of x *at the time of* c_A .

In contrast, on MacFarlane’s view, the time coordinate plays only the first role of determining the experienced object’s time-slice. The standard of taste coordinate s is independently and directly determined by the context of assessment. There is no need to determine a judge’s time-slice. Given MACFARLANE TRUTH, ‘This cookie is tasty’ as used in c_U and assessed at c_A is true just in case $\llbracket \text{This cookie is tasty} \rrbracket^{c_U, \langle w_{c_U}, t_{c_U}, s_{c_A} \rangle} = 1$ – that is, in case the time-slice of the cookie (which is salient in c_U) *at the time of* c_U is tasty by standard of taste s . Note that what matters is not that t is determined by c_U rather than c_A but the fact that t has no impact on the determination of the standard of taste s .

The difference between MacFarlane’s relativism and centered worlds relativism manifests in distinct truth conditions of past tense and future tense sentences involving first-personal uses of PPTs. On MacFarlane’s relativism, the past tense operator, shifting the time coordinate to some time prior to the time of speaking, makes an earlier time-slice of the object relevant for truth and falsity. On centered worlds relativism, the past tense operator makes earlier time-slices of the object and of the experiencer, or judge, relevant. MacFarlane has us consider the following situation.

(6.54) HEN-OF-THE-WOODS

Suppose that one’s tastes change. At c_1 , one likes the taste of Hen-of-the-Woods, while at c_2 (one year later), one dislikes the taste—not because the taste has changed, but because one’s reactions to it have changed. (MacFarlane, 2012, 188)

Both MacFarlane’s relativism and centered worlds relativism predict that at c_2 , one can felicitously assert (6.55):

(6.55) Hen-of-the-Woods is not tasty.

In addition, MacFarlane’s relativism predicts that (6.56) is felicitously assertable in c_2 .

(6.56) Hen-of-the-Woods was not tasty a year ago.

Centered worlds relativism, in contrast, predicts that (6.56) is false from the perspective of c_2 . Instead, (6.57) is felicitously assertable in c_2 since, as stipulated, Hen-of-the-Woods tasted good to the speaker at c_1 , a year before c_2 .

(6.57) Hen-of-the-Woods was tasty a year ago.

MacFarlane admits that judgments about the felicity of (6.56) and (6.57) may differ. However, he points to three empirical predictions of centered worlds relativism that are supposed to provide reason to reject the view in favour of MacFarlane's relativism, which does not make these predictions. I quote the passage from MacFarlane (2012, 189-90) in full:

First, it [Lasersohn's and Stephenson's forms of relativism] predicts that

(6.58) Hen-of-the-Woods will still be tasty in fifty years

should entail

(6.59) Someone will be alive in fifty years,

(I assume here that it is not the case that any dead creature likes the taste of Hen-of-the-Woods, or any other food, while dead.) But this entailment seems dubious.

Second, it would seem odd to say:

(6.60) Last year Hen-of-the-Woods was tasty, but this year it isn't. It has exactly the same flavor this year that it did last year, but after my mushroom tasting class I now find it unappealing.

By contrast, it would not be at all odd to say:

(6.61) Last year Hen-of-the-Woods was tasty to me, but this year it isn't. It has exactly the same flavor this year that it did last year, but after my mushroom tasting class I now find it unappealing.

But the judge-in-index view [of Lasersohn and Stephenson] would not predict a difference in acceptability here.

Third, endorsing [(6.57) Hen-of-the-Woods was tasty a year ago] does not sit well with saying that an assertion in c_1 of

(6.62) Hen-of-the-Woods is tasty

must be retracted in c_2 . It would be odd (at the very least) to say:

(6.63) Last year I asserted that Hen-of-the-Woods was tasty. And last year Hen-of-the-Woods was tasty. Still, my assertion was not accurate and must be retracted.

In reply to these points, let me start with three cautionary remarks. First, not all predicates of personal taste are alike in their embedding behaviour under tense. Some PPTs, such as ‘tasty,’ which are ascribed of objects, exhibit all of the features of individual-level predicates (ILPs) in the sense of Carlson (1980), which express permanent and tendentially stable properties. Other PPTs, such as ‘fun,’ which are ascribed of events, satisfy most of the tests for ILPs, but exhibit some of the behaviour of stage-level predicates (SLPs), which express transient and episodic properties, in sentences with temporal modifiers.³⁷ Consider the following contrasts.

(6.64) a. SLP: Ben was drunk yesterday / last month / a year ago.
 b. ILP: *Ben was tall yesterday / last month / a year ago.

(6.65) a. This roller coaster ride was fun a few seconds ago.
 b. *This lasagne was tasty a few seconds ago / a minute ago.

(6.66) a. SLP: When Ben is in the pub, he is drunk.
 b. ILP: *When Ben is at home, he is tall.

(6.67) a. [Anna, during a roller coaster ride:] When this ride is slower, it is fun.
 b. [Anna, while eating her lasagne:] ? When this lasagne is saltier, it is tasty.

In (6.64) and (6.66), the sentences in (a) with an SLP allow temporal modification, while the sentences in (b) with an ILP do not. In (6.65) and (6.67), ‘tasty’ patterns with the ILP in (6.64) and (6.66), and ‘fun’ patterns with the SLP. That is, the sentences in (6.65a) and (6.67a) ascribing fun are acceptable, while the sentences in (6.65b) and (6.67b) ascribing tastiness are not. Note, however, that (6.67b) has an acceptable reading on which ‘this lasagne’ refers to a kind – perhaps a brand of frozen lasagne. However, the relevant contrast between ILPs and SLPs is found in predications to particulars. Even for standard examples of ILPs, kind readings such as the following are acceptable: ‘When firemen are tall, they’re better at their job.’

³⁷ For the distinction between individual-level predicates like ‘intelligent’ and stage-level predicates like ‘available’, and a variety of tests, see Chierchia (1995) and Kratzer (1995). It should be noted that the distinction is not assumed to be made in the lexicon once and for all. In many cases, setting up a pertinent context allows for SLP-readings of otherwise individual-level predicates (Kratzer, 1995, 125–6). Moreover, the decision between a state being ‘stable’ rather than ‘transient’ is often difficult, since these notions are vague.

The case of tastiness is no exception. Taste properties that determine tastiness supervene on physical properties of food. These *may* change – for instance, milk gets sour after a while as a result of fermentation processes – and thus change the tastiness. Moreover, the context can be set up in a way that allows for ‘tasty’ to behave like a stage-level predicate. In (6.65b) below, we may stipulate that lasagne drastically changes in taste as it cools down extremely quickly, giving (6.65b) a much more acceptable reading. By all grammatical tests (see, e.g. Chierchia (1995, 177–81)), however, ‘tasty’ classifies as an individual-level predicate (see Pearson (forthcoming, §3.4)).

A possible explanation of the behaviour of ‘fun’ under temporal modifiers, which preserves the thesis that ‘fun’ is an ILP, may be that fun is predicated of a *temporal part* of the event rather than the whole event in (6.65a) and (6.67a).³⁸ For our purposes, it is only crucial to note that ‘fun’ shows a different behaviour from ‘tasty’ in certain constructions with temporal modifiers. Consequently, we should understand MacFarlane’s arguments and the data to follow to be representative only of ‘tasty,’ not of PPTs in general.

My second cautionary remark is that most extant accounts of PPTs treat ‘tasty’ and ‘taste good’ as semantically on a par. However, when embedded under tense, judgments of (in)felicity may on occasion differ in strength. We will see an example in response to MacFarlane’s arguments below.

Third, the cases of PPTs embedded under past or future tense that distinguish centered/sequenced worlds relativism and MacFarlane’s relativism are cases in which it is stipulated that the food is the same but the speaker’s perspective has changed. In many of these cases, judgments tend to vary across subjects and are often not very clear. This may be due to the fact that our opinions, or the lack thereof, about the metaphysics of taste and taste experiences are relevant to the cases. I invite the reader to test their judgments on (6.68).

(6.68) [Situation: Jim just had an operation that gave him back his sense of smell. As a result, he now likes the taste of foods which he used to think tasted bland or even bad, his wife’s cooking included. Jim is euphoric. He says to his wife Allie after dinner:]

Jim: ? Allie, your dishes were always tasty!

I should thus note that the observations to follow must be understood as merely a tentative and preliminary discussion of the interaction of PPTs with tense. Let me first reply to the three points MacFarlane makes and then present some data that favours the centered/sequenced worlds approach.

First, while the entailment of (6.59) [‘Someone will be alive in fifty years’] by (6.58) [‘Hen-of-the-Woods will still be tasty in fifty years’] is far from obvious, explicit denial of the existence of any beings who could taste Hen-of-the-Woods after asserting (6.58) is also odd.

- (6.69) a. ? No one will be alive in 50 years. But Hen-of-the-Woods will still taste good.
 b. ? Even if no one will be alive in 50 years, Hen-of-the-Woods will still taste good.

³⁸See, e.g., Lasersohn (2011, 435-8) for relevant discussion

‘Tasty’/‘taste good’ seem to be different from ‘poisonous,’ for example. It is plausible to think that a particular kind of mushroom is poisonous even when it does not in fact poison anyone. It seems much less plausible that a kind of mushroom tastes good even when it does not taste good to anyone.

Second, the difference in assertability between (6.60) and (6.61) seems marginal, and further reduces when ‘tasty’ and ‘tasty to me’ are replaced by ‘taste good’ and ‘taste good to me’

(6.60’) Last year Hen-of-the-Woods tasted good, but this year it doesn’t. It has exactly the same flavor this year that it did last year, but after my mushroom tasting class I now find it unappealing.

(6.61’) Last year Hen-of-the-Woods tasted good to me, but this year it doesn’t. It has exactly the same flavor this year that it did last year, but after my mushroom tasting class I now find it unappealing.

Third, note that MacFarlane’s objection based on the oddity of (6.63), repeated below, does not in fact apply to the centered/sequenced worlds view (nor in fact Lasersohn’s and Stephenson’s view, as they presumably do not predict that (6.62) must be retracted in c_2).

(6.63) Last year I asserted that Hen-of-the-Woods was tasty. And last year Hen-of-the-Woods was tasty. Still, my assertion was not accurate and must be retracted. (MacFarlane, 2012, 189-90)

According to MacFarlane, ‘accurate’ is the property an assertion has at a context of assessment c_A iff the content it expresses is true at the context of use c_U and c_A . I prefer to use the predicates ‘correct’ or ‘true’ rather than ‘accurate’ here, so I will assume a reading of (6.63) on which one of these predicates is substituted for ‘accurate’.³⁹ Then we can see that, contrary to MacFarlane’s claim, centered/sequenced worlds relativism does not predict the felicity of (6.63). Given principle (6.48) about past tense truth ascriptions, (6.63) is odd because ‘my assertion was not correct/true’ is false. It is false because in c_1 , the context in which the assertion of ‘Hen-of-the-Woods is tasty’ was made, it *was* true.

A wider look at the interaction between PPTs and tense reveals that in many cases, centered/sequenced worlds relativism makes intuitively correct predictions where MacFarlane’s relativism does not. Consider first three cases of past tense uses. Assume once more the situation given in HEN-OF-THE-WOODS. In this situation, it would be odd to say in c_2

³⁹To reiterate the point made in section 6.5.1 about the truth-aptness of assertions, I do not believe that assertions are bearers of truth and falsity. However, where we ascribe truth to them explicitly, assertion truth boils down to truth of sentences-in-contest.

(6.70) Hen-of-the-Woods is not tasty, and never was.

Centered/sequenced worlds relativism predicts that (6.70) is false since there was a time, c_1 , when Hen-of-the-Woods tasted good to the speaker. On MacFarlane's view, (6.70) is perfectly assertable. As long as Hen-of-the-Woods does not change in taste, Hen-of-the-Woods past and present is not tasty by the single standard of taste of the speaker at c_2 .

Consider next (6.71).

(6.71) [Situation: Sal grew up in Ethiopia but left the country at age 6. As an adult, he is talking to Anna about his childhood:]

Anna: Did you eat injera, the Ethiopian flatbread?

Sal: Yes, we had injera all the time. It was tasty. Of course, I don't know whether I would like it today.

Sal's reply seems perfectly ordinary. But on MacFarlane's view, that is surprising. For his assertion of 'It was tasty' is true at Sal's context of assessment c_A just in case the past time-slices of injera are tasty by Sal's current standard of taste at c_A . But Sal goes on to express ignorance about whether or not the injera he used to have would please him now. So it is surprising that Sal would at the same time express judgments, *by his current standards of taste*, about the tastiness of the injera.

Third, consider (6.72).

(6.72) [Situation: Guests at a dinner party have been served bottles of the same dessert wine throughout the whole dinner.]

Guest: This sweet dessert wine tastes good now, but it didn't when we drank it with the main course.

(6.72) seems to be a perfectly felicitous assertion to make in the situation. This is predicted by centered/sequenced worlds relativism. But on MacFarlane's relativism, it is not. If the wine did not change its taste properties during the course of the dinner, then by the speaker's current standards of taste, it either did and does taste good or it did not and does not taste good.

As a final piece of data in favour of centered/sequenced worlds relativism, consider a case involving a future ascription of tastiness.

(6.73) [Situation: Beavis and Butthead are about to smoke a joint. They both know about the appetite-enhancing effects of marijuana. And they both know that they dislike the taste of muesli.]

Beavis: This muesli is going to be so tasty later.

Beavis' claim seems felicitous. But given that Beavis' gustatory standard at the time of speaking are such that muesli is not tasty according to it, MacFarlane predicts that (6.73) is not a good thing for Beavis to say. In contrast, it is felicitous on centered/sequenced worlds relativism in case the future time-slice of the muesli will taste good to the future time-slice(s) of Beavis (and Butthead).

These observations are not conclusive. But short of a fully worked out theory of the interaction between PPTs like 'tasty' and tense, they lend support to semantic views on which the tenses have an effect on the judge's standards of taste.

6.6 Modality, Transworld Assessments, and Predicates of Personal Taste

Modality raises parallel issues for MacFarlane's relativism and centered/sequenced worlds relativism, which I will treat briefly here. MacFarlane's relativism is a form of non-indexical contextualism regarding worlds, whereas centered/sequenced worlds relativism is a form of relativism about worlds. This has consequences for judgments about counterfactual assertions – assertions that could have been made but were not in fact made. MacFarlane must reject principles (6.35) and (6.36) connecting assertion and content truth, and must accept that a sentence like (6.74) is true.

(6.74) Suppose the moon was made of cheese and I said 'The moon is made of cheese.' That assertion is true, but what it would say is false.

On centered/sequenced worlds relativism, (6.74) is false. It is false because both the assertion and its content are false at the assessor's world (assuming that the assessor is in the actual world). A plausible principle for transworld truth ascriptions is (6.75).

(6.75) An assertion with the content p made under counterfactual circumstances would be true iff the content p would be true under those circumstances.

(6.75) licenses true assertions of sentences like (6.76).

(6.76) Suppose the moon was made of cheese and I said 'The moon is made of cheese.' That would be true. But of course it isn't in fact true.

MacFarlane's point about the double role of the time coordinate on centered/sequenced worlds relativism but not on his relativism also applies to the world coordinate. On the former view, the world coordinate determines the worldly location not only of the experienced object but also of the judge. We can thus expect differences in the views' predictions for PPTs embedded under modal operators. For instance, MacFarlane's relativism does not predict that the counterfactual in (6.77) is true.

(6.77) If no sentient beings had ever existed, nothing would be tasty.

However, on plausible assumptions about the semantics of counterfactuals, centered/sequenced worlds relativism has (6.77) come out true.⁴⁰ This may not chime with intuitive judgment and, as MacFarlane (2012, 192) points out, may play into the hands of opponents who wish to object to relativism on the grounds that it makes the property of tastiness implausibly mind-dependent. This point must be conceded. Centered/sequenced worlds relativism is, whereas MacFarlane's relativism is not, committed to the truth of (6.77) and whatever metaphysical consequences one may or may not derive from it.

Moreover, given a standard semantics for alethic modals, centered/sequenced worlds relativism is committed to the truth of (6.78).

(6.78) It is impossible that [wine is tasty, even though no one likes it.]⁴¹

Centered/sequenced worlds relativism predicts that the unembedded 'wine is tasty, even though no one likes it' is necessarily false since there is no centered world $\langle w, t, x \rangle$ such that no one in w at t likes wine and wine tastes good to x in w at t (assuming that someone's liking wine requires that wine tastes good to that person). Whatever one's judgments about (6.78), it seems a strong metaphysical commitment coming from semantic theory. But notice that opting for MacFarlane's relativism does not help avoiding potentially implausible metaphysical commitments. Consider the assertions in (6.79a) and (6.79b).

(6.79) GLUCOSE

Situation: As a matter of fact, Sal finds pure glucose powder extremely tasty.

Sal asserts:

- a. It is necessary that glucose is tasty.
- b. It is necessary that glucose is tasty to me.

MacFarlane's view predicts that Sal can felicitously assert (6.79a) but not (6.79b). Glucose has its chemical composition necessarily. In all possible worlds, glucose is $C_6H_{12}O_6$. So in all possible worlds, glucose has the same chemical composition that makes for its taste (other factors being due to the tasting subject). So given Sal's *actual* standard of taste (which is such that glucose counts as tasty), glucose in all possible worlds is tasty by the standards salient in Sal's context of assessment. Hence (6.79a) is true at

⁴⁰I will not here provide a relativistic semantics for counterfactuals. It seems plausible to assume that it will involve quantification over/shifting of entire centered worlds, not merely possible worlds. Perhaps Stephenson's (2007b, ch. 3) semantics for indicative conditionals will serve as a good starting point.

⁴¹This example is from a brief discussion by Lasersohn (2005, 663 n.13), where the point is made that the unembedded 'Wine is tasty, even though no one likes it', while false, is not necessarily false, as predicted by Lasersohn's view.

Sal's context of assessment. (6.79b), however, is false at any context of assessment because whatever the standards in that context, there is some possible worlds in which glucose is not tasty by Sal's standard in that world: it is not necessary that Sal finds glucose tasty. But both (6.79a) and (6.79b) seem equally implausible. In contrast, the CW/SW view predicts that both (6.79a) and (6.79b) are false (from any perspective $\langle w, t, x \rangle$). (6.79a) is unassertable because it is false from Sal's perspective $\langle @, t, \text{Sal} \rangle$. There is some possible world w' in which glucose fails to taste good to Sal in w' at t .

I leave it to further inquiry to decide which relativistic semantics incurs which metaphysical consequences. With many alethic modals and counterfactuals, it seems that judgments track not merely linguistic intuitions of felicity, but more theoretically founded metaphysical judgments. It would seem that commitments to some form of mind-dependence, while accompanied by less than favourable judgments of counterfactuals (6.77), may after all fit the relativist's outlook rather well.

6.7 On Content

The sequenced worlds account of eavesdropping, retraction, disagreement and truth ascriptions required a fuller development of the common ground model of communication sketched in chapter 4 and revealed a more complex overall view of content in thought and talk. On the view, two distinctions take center stage. The first is the distinction between horizontal content and diagonal content. The second is the distinction between *intra*-conversational and *inter*-conversational pragmatic phenomena. As a rule of thumb, horizontal content explains *inter*-conversational phenomena and diagonal content explains *intra*-conversational phenomena. Table 6.2 below provides a summary of the roles of horizontal and diagonal content.

The sequenced worlds view is thus a form of content pluralism – a form of content dualism, to be precise. It shares this feature with many other views.⁴² What distinguishes the sequenced worlds view from these views, however, is the way that explanatory roles are divided up between horizontal and diagonal content. The view's association of diagonal content with *intra*-conversational phenomena and of horizontal content with *inter*-conversational phenomena is pretty natural on a reasonable and widespread understanding of content. On this way of understanding contents, they are, or determine, truth conditions. But in asking, of a particular assertion, *What are the assertion's truth conditions?* we need not be asking for the *unique* truth conditions of that assertion. As Perry (2001, 80) writes,

⁴²See for example Perry's (2001) view on which there are multiple levels of content that a single utterance expresses, Chalmer's (2006) two-dimensionalism, Higginbotham's (1995) view on which there is reflexive content and content on which the reflexive tense and modal information is discarded, and Stanley (1997a,b, 2002) who distinguishes between communicated content and the content of modal evaluations.

Table 6.2: Roles of Horizontal and Diagonal Content

Diagonal content	Horizontal content
<ul style="list-style-type: none"> • speech act content/communicated content/information (ch. 4) • object of <i>intra</i>-conversational, simultaneous truth assessment (§§5.4&6.5) • object whose temporal update is retracted in <i>intra</i>-conversational retraction (§5.7) • object of <i>intra</i>-conversational disagreement (§6.2) • object of <i>intra</i>-conversational cases of believing alike as agreement (§6.4) • object of believing alike as similarity in cognitive significance (inter-conversational cases) (§6.4) 	<ul style="list-style-type: none"> • object of <i>inter</i>-conversational truth assessment (§§5.4&6.5) • object whose temporal update is retracted in <i>inter</i>-conversational retraction (§5.7) • object in <i>inter</i>-conversational disagreement (modulo temporal update) (§6.2) • object of <i>inter</i>-conversational cases of believing alike as agreement (modulo temporal update) (§6.4)

[T]he concept of ‘truth conditions of an utterance’ is a *relative concept*...

Instead of thinking in terms of *the* truth conditions, we should think about the truth conditions of an utterance *given* various facts about it... talking about *the* content of an utterance is an oversimplification.

What are these background facts relative to which an assertion has truth conditions? They are the facts that a competent hearer must know in order to understand the assertion. For instance, she must know the language of the assertion, which words are involved, their lexical meaning, and the syntax of the asserted sentence. Given these facts – language, words involved, their lexical meaning and syntax – the assertion has the truth conditions of diagonal content. How so? Suppose someone says, ‘This is Mt. Tallac.’ Given just these background facts and assuming that the speaker speaks truthfully, you will receive the following truth-conditional information: that the speaker is in a context in which the demonstrated object is Mt. Tallac. In some sense, this information is *about* the context – the situation the speaker (and hearer) is in. On the sequenced worlds view, this is the content communicated in conversation: Conversation is the project of self-and-group-location, of locating the group in a shared situation. That the shared situation is one in which Mt. Tallac is the demonstrated

object is potentially very useful information.

When we retrieve information from assertions made in conversations of which we are not members, we do not assume that we share the situation with the speaker. We are not engaged in the joint project of locating ourselves. And while it may be useful to gain information about what situation the speaker is in, what we typically want to know from her assertion is what, *given the speaker's situation*, is true of the world. If I read in an old newspaper – say the edition of 6 May 2010 – ‘Today elections are being held in the UK,’ it is surely useful to know that on the day the newspaper appeared, elections were being held. But my main interest is not under which conditions the sentence can be truly asserted. Rather, I want to know, *given* that the sentence was uttered on 6 May 2010, under which conditions it is true. I want to know the horizontal content that elections were held on 6 May 2010. These are the assertion's truth conditions given the facts that determine the language, words involved, their lexical meaning and syntax, *and the context of utterance*. The information given by the horizontal is about what else has to be true – given these facts.

6.8 Conclusion

Let me summarise. Chapter 4 started from two plausible assumptions, the Lewisian account of attitudes as self-location and the Lockean picture of communication as transfer of information from speaker's to hearer's head. It was argued that these views can be reconciled on a conception of content in terms of sequenced worlds.

Chapters 5 and 6 located the sequenced worlds view of content and communication in the debate between contextualism and relativism about predicates of personal taste. It was shown that the sequenced worlds view is compatible with nonindexical contextualist as well as truth relativist pragmatics. The discussion of crucial pragmatic phenomena such as eavesdropping, retraction, and disagreement exposed some of the shortcomings of nonindexical contextualism and thus presented an argument in favour of sequenced worlds relativism. In order to account for these pragmatic phenomena, the common ground model of communication sketched in chapter 4 was developed more fully. The resulting view is a form of content dualism. Both horizontal content and diagonal content play important roles in our cognitive and communicative lives.

It remains to be said that the sequenced worlds view developed in chapters 4 through 6 on the basis of ideas by Ninan (2010b) and Torre (2010) is incomplete. An important question, for instance, is whether the view sits comfortably with a plausible semantics of standard attitude verbs. Likewise, its relationship to mainstream semantic accounts of attitude verbs that are taken to give rise to ascriptions of *de se* attitudes – such as ‘Ben expects to win the race’ or ‘Sal wants to become president one day’ – requires attention. These questions are important, but they go beyond the current

investigation, whose focus has been the role of content in communication. They will have to be left for future work.

At the same time, the sequenced worlds view promises interesting applications beyond the perspectival attitudes we looked at here. For instance, it may be fruitfully applied to other linguistic expressions that have received relativistic treatments, in particular epistemic modals (Egan et al., 2005; Egan, 2007; MacFarlane, 2011a; Stephenson, 2007a). Moreover, the view may offer an interesting account of ‘audience-sensitive’ assertions such as the billboard inscription ‘Jesus loves you,’ which is plausibly understood as addressing each reader individually. The sequenced worlds account offers the prospects of an account on which an explosion of content – a different content for each reader – is avoided. I leave it for future work to develop the details and to evaluate the advantages of these applications.

BIBLIOGRAPHY

- Akerman, J. and P. Greenough (2010). Vagueness and non-indexical contextualism. In S. Sawyer (Ed.), *New Waves in Philosophy of Language*. New York: Palgrave Macmillan.
- Asher, N. and A. Lascarides (2003). *Logics of Conversation*. Cambridge: Cambridge University Press.
- Bach, K. (1994). Conversational implicature. *Mind and Language* 9, 124–62.
- Blome-Tillmann, M. (2008). The indexicality of ‘knowledge’. *Philosophical Studies* 138(1), 29–53.
- Blome-Tillmann, M. (2009). Contextualism, subject-sensitive invariantism, and the interaction of ‘knowledge’-ascriptions with modal and temporal operators. *Philosophy and Phenomenological Research* 79(2), 315–31.
- Brogaard, B. (2008). In defence of a perspectival semantics for ‘know’. *Australasian Journal of Philosophy* 86(3), 439–59.
- Brown, J. (2005). Adapt or die? The death of invariantism? *The Philosophical Quarterly* 55(219), 263–285.
- Brown, J. (2006). Contextualism and warranted assertibility manoeuvres. *Philosophical Studies* 130(3), 407–435.
- Cappelen, H. (2008a). Content relativism and semantic blindness. In M. García-Carpintero and M. Kölbel (Eds.), *Relative Truth*, pp. 265–86. Oxford: Oxford University Press.
- Cappelen, H. (2008b). The creative interpreter: Content relativism and assertion. *Philosophical Perspectives* 22, 23–46.
- Cappelen, H. and J. Hawthorne (2009). *Relativism and Monadic Truth*. Oxford: Oxford University Press.
- Cappelen, H. and E. Lepore (2005). *Insensitive Semantics*. Malden, MA: Blackwell.
- Carlson, G. N. (1980). *Reference to Kinds in English*. New York and London: Garland.
- Castañeda, H.-N. (1966). ‘He’: A study in the logic of self-consciousness. *Ratio* 8, 130–57.
- Castañeda, H.-N. (1967). Indicators and quasi-indicators. *American Philosophical Quarterly* 4(2), 85–100.

- Chalmers, D. (2006). Two-dimensional semantics. In E. Lepore and B. C. Smith (Eds.), *The Oxford Handbook of Philosophy of Language*, pp. 574–606. Oxford: Oxford University Press.
- Chierchia, G. (1989). Anaphora and attitudes *De Se*. In R. Bartsch, J. van Benthem, and van Emde Boas (Eds.), *Semantics and Contextual Expression*, pp. 1–31. Dordrecht: Foris.
- Chierchia, G. (1995). Individual-level predicates as inherent generics. In G. N. Carlson, F. J. P. G. N. Carlson, and F. J. Pelletier (Eds.), *The Generic Book*, pp. 176–223. Chicago: University of Chicago Press.
- Clark, H. H. (1992). *Arenas of Language Use*. Chicago: University of Chicago Press.
- Clark, H. H. and E. F. Schaefer (1989). Contributing to discourse. *Cognitive Science* 13, 259–94.
- Cohen, A. (1999a). Generics, frequency adverbs, and probability. *Linguistics and Philosophy* 22(3), 221–53.
- Cohen, S. (1987). Knowledge, context, and social standards. *Synthese* 73(1), 3–26.
- Cohen, S. (1999b). Contextualism, skepticism, and the structure of reasons. *Noûs* 33, 57–89.
- Cohen, S. (2004a). Contextualism and unhappy-face solutions: Reply to Schiffer. *Philosophical Studies* 119(1), 185–97.
- Cohen, S. (2004b). Knowledge, assertion, and practical reasoning. *Philosophical Issues* 14 (Epistemology), 482–91.
- Davis, W. (2007). Knowledge claims and context: Loose use. *Philosophical Studies* 132(3), 395–438.
- DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review* 104(1), 1–52.
- DeRose, K. (2004). The problem with subject-sensitive invariantism. *Philosophical Studies* 68(2), 346–50.
- DeRose, K. (2005). The ordinary language basis for contextualism, and the new invariantism. *The Philosophical Quarterly* 55(219), 172–198.
- DeRose, K. (2006). “Bamboozled by our own words”: Semantic blindness and some arguments against contextualism. *Philosophy and Phenomenological Research* 73(2), 316–38.
- Dimmock, P. and T. Huvenes (2012). Knowledge, conservatism, and pragmatics. Unpublished Ms.

- Dretske, F. I. (1970). Epistemic operators. *The Journal of Philosophy* 67(24), 1007–1023.
- Egan, A. (2006). Secondary qualities and self-location. *Philosophy and Phenomenological Research* 72(1), 97–119.
- Egan, A. (2007). Epistemic modals, relativism, and assertion. *Philosophical Studies* 133, 1–22.
- Egan, A. (2009). Billboards, bombs and shotgun weddings. *Synthese* 166(2), 251–279.
- Egan, A. (2010a). Disputing about taste. In R. Feldman and T. A. Warfield (Eds.), *Disagreement*, pp. 247–86. Oxford: Oxford University Press.
- Egan, A. (2010b). Three grades of self-involvement. Talk delivered at Arché/CSMN's *De Se* Workshop, Oslo, June 2010.
- Egan, A. (2011). Relativism about epistemic modals. In S. D. Hales (Ed.), *Blackwell Companion to Relativism*. Oxford: Blackwell.
- Egan, A., J. Hawthorne, and B. Weatherson (2005). Epistemic modals in context. In G. Preyer and G. Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, pp. 131–67. Oxford: Clarendon Press.
- Engç, M. (1986). Towards a referential analysis of temporal expressions. *Linguistics and Philosophy* 9, 405–26.
- Fantl, J. and M. McGrath (2002). Evidence, pragmatics, and justification. *The Philosophical Review* 111(1), 67–94.
- Fantl, J. and M. McGrath (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Farkas, D. F. and K. B. Bruce (2010). On reacting to assertions and polar questions. *Journal of Semantics* 27(1), 81–118.
- Gärdenfors, P. (1988). *Knowledge in Flux: Modeling the Dynamics of the Epistemic States*. Cambridge, MA: MIT Press.
- Ginzburg, J. (1996). Dynamics and the semantics of dialogue. In J. Seligman and D. Westerstahl (Eds.), *Language, Logic, and Computation*, Volume 1, pp. 221–37. Stanford, CA: CSLI Publications.
- Ginzburg, J. (2012). *The Interactive Stance: Meaning for Conversation*. Oxford: Oxford University Press.
- Glanzberg, M. (2007). Context, content, and relativism. *Philosophical Perspectives* 136, 1–29.
- Glanzberg, M. (2011). More on operators and tense. *Analysis Reviews* 71(1), 12–23.

- Greenough, P. (2011). Truth-relativism, norm-relativism, and assertion. In J. Brown and H. Cappelen (Eds.), *Assertion*. Oxford: Oxford University Press.
- Grice, P. (1975). Logic and conversation. In P. Cole and J. Morgan (Eds.), *Syntax and Semantics 3: Speech Acts*, pp. 41–58. Page references are to the reprint in Grice (1989), pp. 22–40. New York: Academic Press.
- Grice, P. (1989). *Studies in the Way of Words*. Cambridge, MA, and London: Harvard University Press.
- Gunlogson, C. (2003). *True to Form: Rising and Falling Declaratives as Questions in English*. New York: Routledge.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Heim, I. (2004). Lecture notes on indexicality. Unpublished MIT Lecture Notes.
- Heim, I. and A. Kratzer (1998). *Semantics in Generative Grammar*. Malden, MA: Blackwell.
- Higginbotham, J. (1995). Tensed thoughts. *Mind & Language* 10(3), 226–49.
- Huvenes, T. T. (2012). Varieties of disagreement and predicates of taste. *Australasian Journal of Philosophy* 90(1), 167–81.
- Kahneman, D., P. Slovic, and A. Tversky (Eds.) (1982). *Judgment under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.
- Kahneman, D. and A. Tversky (1982). Judgment under uncertainty: Heuristics and biases. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), *Judgment under Uncertainty: Heuristics and Biases*, pp. 3–20. Cambridge: Cambridge University Press.
- Kahneman, D. and A. Tversky (Eds.) (2000). *Choices, Values, and Frames*. Cambridge: Cambridge University Press.
- Kamp, H. (1975). Two theories about adjectives. In E. Keenan (Ed.), *Formal Semantics of Natural Language*. Cambridge: Cambridge University Press.
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, and H. Wettstein (Eds.), *Themes from Kaplan*, pp. 481–563. New York, Oxford: Oxford University Press.
- Kennedy, C. (2000). *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. London: Routledge.
- Kennedy, C. (2007). Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* 30(1), 1–45.
- Kindermann, D. (forthcoming). Relativism, sceptical paradox, and semantic blindness. *Philosophical Studies*, 1–19.

- King, J. (2003). Tense, modality, and semantic values. *Philosophical Perspectives* 17, 195–245.
- Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and Philosophy* 4(1), 1–45.
- Kölbel, M. (2002). *Truth Without Objectivity*. London and New York: Routledge.
- Kölbel, M. (2004). Faultless disagreement. *Proceedings of the Aristotelian Society* 104(1), 53–73.
- Kölbel, M. (2009). The evidence for relativism. *Synthese* 166(2), 375–95.
- Kompa, N. (2002). The context sensitivity of knowledge ascriptions. *Grazer Philosophische Studien* 64, 79–96.
- Kratzer, A. (1995). Stage-level and individual-level predicates. In G. N. Carlson and F. J. Pelletier (Eds.), *The Generic Book*, pp. 125–75. Chicago: University of Chicago Press.
- Kripke, S. A. (2011). Frege's theory of sense and reference: Some exegetical notes. In *Philosophical Troubles. Collected Papers*, Volume Vol. 1, pp. 254–91. Oxford: Oxford University Press.
- Kusumoto, K. (1999). *Tense in Embedded Contexts*. Ph. D. thesis, University of Massachusetts at Amherst, Amherst.
- Kusumoto, K. (2005). On the quantification over times in natural language. *Natural Language Semantics* 13, 317–57. 10.1007/s11050-005-4537-6.
- Lasersohn, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28, 643–86.
- Lasersohn, P. (2008). Quantification and perspective in relativist semantics. *Philosophical Perspectives* 22(1), 305–337.
- Lasersohn, P. (2009). Relative truth, speaker commitment, and control of implicit arguments. *Synthese* 166, 359–374.
- Lasersohn, P. (2011). Contexts, relevant parts, and (lack of) disagreement over taste. *Philosophical Studies* 156(3), 433–49.
- Lewis, D. (1975). Languages and language. In K. Gunderson (Ed.), *Minnesota Studies in the Philosophy of Science*, Volume 7, pp. 3–35. Reprinted in Lewis (1983b), pp. 163–88. Page references are to the reprint. Minneapolis: University of Minnesota Press.
- Lewis, D. (1979a). Attitudes de dicto and de se. *The Philosophical Review* 88(4), 513–43.

- Lewis, D. (1979b). Scorekeeping in a language game. *Journal of Philosophical Logic* 8, 339–59.
- Lewis, D. (1980). Index, context, and content. In S. Kanger and S. Öhman (Eds.), *Philosophy and Grammar*. Dordrecht: Reidel.
- Lewis, D. (1983a). Individuation by acquaintance and by stipulation. *The Philosophical Review* 92(1), 3–32.
- Lewis, D. (1983b). *Philosophical Papers*, Volume 1. New York and Oxford: Oxford University Press.
- Lewis, D. (1986). *On the Plurality of Worlds*. Malden, MA: Blackwell.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy* 74(4), 549–67.
- Liao, S.-y. (2012). What are centered worlds? *The Philosophical Quarterly* 62(247), 294–316.
- Locke, J. (1690). *An Essay Concerning Human Understanding*.
- López de Sa, D. (2008). Presuppositions of commonality: An indexical relativist account of disagreement. In M. García-Carpintero and M. Kölbel (Eds.), *Relative Truth*, pp. 197–310. Oxford: Oxford University Press.
- MacFarlane, J. (2003). Future contingents and relative truth. *The Philosophical Quarterly* 53, 321–36.
- MacFarlane, J. (2005a). The assessment-sensitivity of knowledge attributions. In T. S. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology* 1, pp. 197–233. Oxford: Oxford University Press.
- MacFarlane, J. (2005b). Making sense of relative truth. *Proceedings of the Aristotelian Society* 105, 321–39.
- MacFarlane, J. (2007a). Relativism and disagreement. *Philosophical Studies* 132, 12–31.
- MacFarlane, J. (2007b). Truth and subjectivity. Unpublished Ms. Available at www.johnmacfarlane.net (May 2012).
- MacFarlane, J. (2009). Nonindexical contextualism. *Synthese* 166, 231–50.
- MacFarlane, J. (2011a). Epistemic modals are assessment-sensitive. In B. Weatherson and A. Egan (Eds.), *Epistemic Modality*, pp. 144–78. Oxford: Oxford University Press.
- MacFarlane, J. (2011b). Relativism and knowledge attributions. In S. Bernecker and D. Pritchard (Eds.), *Routledge Companion to Epistemology*, pp. 536–44. London: Routledge.

- MacFarlane, J. (2011c). What is assertion? In H. Cappelen and J. Brown (Eds.), *Assertion. New Philosophical Essays*, pp. 79–96. Oxford: Oxford University Press.
- MacFarlane, J. (2012). Assessment sensitivity. Relative truth and its applications. Unpublished Book Ms. (Draft of 8 February 2012). Available at www.johnmacfarlane.net (May 2012).
- McConnell-Ginet, S. (1973). *Comparative Constructions in English: A Syntactic and Semantic Analysis*. Ph. D. thesis, University of Rochester.
- Moltmann, F. (2010). Relative truth and the first person. *Philosophical Studies* 150(2), 187–220.
- Montminy, M. (2009). Contextualism, relativism and ordinary speakers' judgments. *Philosophical Studies* 143(3), 341–56.
- Morgan, J. (1970). On the criterion of identity for noun phrase deletion. *Chicago Linguistics Society* 6.
- Moruzzi, S. and C. Wright (2009). Trumping assessments and the Aristotelian future. *Synthese* 166(2), 309–31.
- Moss, S. (2012). Updating as communication. *Philosophy and Phenomenological Research*.
- Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy* 86(2), 279–94.
- Nagel, J. (2010). Knowledge ascriptions and the psychological consequences of thinking about error. *The Philosophical Quarterly* 60(239), 286–306.
- Ninan, D. (2010a). Semantics and the objects of assertion. *Linguistics and Philosophy* 33, 355–80.
- Ninan, D. (2010b). *De se* attitudes: Ascription and communication. *Philosophy Compass* 5(7), 551–67.
- Nolan, D. (2006). Selfless desires. *Philosophy and Phenomenological Research* 73(3), 665–79.
- Pagin, P. (2008). What is communicative success? *Canadian Journal of Philosophy* 38(1), 85–115.
- Partee, B. H. (1973). Some structural analogies between tenses and pronouns in English. *The Journal of Philosophy* 70(18), 601–609.
- Pearson, H. (forthcoming). A judge free semantics for predicates of personal taste. *Journal of Semantics*.
- Perry, J. (1977). Frege on demonstratives. *The Philosophical Review* 86(4), 474–497.

- Perry, J. (1979). The problem of the essential indexical. *Noûs* 13(1), 3–21.
- Perry, J. (1986). Thought without representation. *Proceedings of the Aristotelian Society, Supplementary Volumes* 60, 137–52.
- Perry, J. (2001). *Reference and Reflexivity*. Stanford: CSLI Publications.
- Quine, W. (1969). Propositional objects. In *Ontological Relativity and Other Essays*, pp. 139–60. New York: Columbia University Press.
- Recanati, F. (2007). *Perceptival Thought. A Plea for (Moderate) Relativism*. Oxford: Oxford University Press.
- Recanati, F. (2008). Moderate relativism. In M. Kölbel and M. García-Carpintero (Eds.), *Relative Truth*, pp. 41–62. Oxford: Oxford University Press.
- Richard, M. (2004). Contextualism and relativism. *Philosophical Studies* 119(1), 215–42.
- Richard, M. (2008). *When Truth Gives Out*. Oxford: Oxford University Press.
- Roberts, C. (1996). Information structure in discourse: Towards an integrated formal theory of pragmatics. In J.-H. Yoon and A. Kathol (Eds.), *Working Papers in Linguistics No. 49: Papers in Semantics*, pp. 91–136. Ohio State University.
- Roberts, C. (2004). Context in dynamic interpretation. In L. R. Horn and G. Ward (Eds.), *The Handbook of Pragmatics*, pp. 197–220. Oxford: Blackwell.
- Rosenkranz, S. (2008). Frege, relativism and faultless disagreement. In M. García-Carpintero and M. Kölbel (Eds.), *Relative Truth*, pp. 225–37. Oxford: Oxford University Press.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Nous* 35(4), 477–514.
- Rysiew, P. (2007). Speaking of knowing. *Noûs* 41(4), 627–62.
- Schaffer, J. (2004a). From contextualism to contrastivism. *Philosophical Studies* 119(1), 73–103.
- Schaffer, J. (2004b). Skepticism, contextualism, and discrimination. *Philosophy and Phenomenological Research* 69(1), 138–55.
- Schaffer, J. (2006). The irrelevance of the subject: Against subject-sensitive invariantism. *Philosophical Studies* 127(1), 87–107.
- Schaffer, J. (2011). Perspective in taste predicates and epistemic modals. In B. Weatherston and A. Egan (Eds.), *Epistemic Modals*, pp. 179–226. Oxford: Oxford University Press.

- Schaffer, J. and Z. G. Szabó (forthcoming). Epistemic comparativism: A contextualist semantics for knowledge ascriptions. *Philosophical Studies*.
- Schiffer, S. (1996). Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society* 96, 317–33.
- Slovic, P., B. Fischhoff, and S. Lichtenstein (1982). Fact versus fears: Understanding perceived risk. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), *Judgment under Uncertainty: Heuristics and Biases*, pp. 463–89. Cambridge: Cambridge University Press.
- Stalnaker, R. C. (1970). Pragmatics. *Synthese* 22.
- Stalnaker, R. C. (1974). Pragmatic presuppositions. In M. K. Munitz and P. Unger (Eds.), *Semantics and Philosophy*. New York: University Press.
- Stalnaker, R. C. (1978). Assertion. In P. Cole (Ed.), *Syntax and Semantics 9: Pragmatics*, pp. 315–32; page references are to the reprint in Stalnaker (1999). New York: Academic Press.
- Stalnaker, R. C. (1981). Indexical belief. *Synthese* 49(1), 129–151.
- Stalnaker, R. C. (1999). *Context and Content*. Oxford: Oxford University Press.
- Stalnaker, R. C. (2002). Common ground. *Linguistics and Philosophy* 25(5), 701–21.
- Stanley, J. (1997a). Names and rigid designation. In B. Hale and C. Wright (Eds.), *A Companion to Philosophy of Language*, pp. 555–85. Oxford: Blackwell.
- Stanley, J. (1997b). Rigidity and content. In R. Heck (Ed.), *Language, Thought, and Logic. Essays in Honour of Michael Dummett*, pp. 131–56. Oxford: Oxford University Press.
- Stanley, J. (2002). Making it articulated. *Mind and Language* 17, 149–68.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stephenson, T. (2007a). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy* 30(4), 487–525.
- Stephenson, T. (2007b). *Toward a Theory of Subjective Meaning*. Ph. D. thesis, MIT.
- Stephenson, T. (2011). The pragmatics of judge dependence. Talk at “Meaning, Context and Implicit Content” Conference, Cerisy, France. Slides available at <http://sites.google.com/site/taminastephenson> (accessed 03/2012).
- Stevenson, C. L. (1963). *Facts and Values*. New Haven and London: Yale University Press.
- Stojanovic, I. (2007). Talking about taste: disagreement, implicit arguments, and relative truth. *Linguistics and Philosophy* 30(6), 691–706.

- Stojanovic, I. (2012). The problem of de se assertion. *Erkenntnis* 76, 49–58.
- Torre, S. (2010). Centered assertion. *Philosophical Studies* 150, 97–114.
- Unger, P. (1971). A defense of skepticism. *The Philosophical Review* 80(2), 198–219.
- Unger, P. (1975). *Igorance: A Case for Skepticism*. Oxford: Clarendon Press.
- van Ditmarsch, H., W. van der Hoek, and B. Kooi (2008). *Dynamic Epistemic Logic*, Volume Synthese Library Vol. 337. Dordrecht: Springer.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. Roth and G. Ross (Eds.), *Doubting: Contemporary Perspectives on Skepticism*. Dordrecht: Kluwer.
- von Fintel, K. and A. Gillies (2008). CIA leaks. *Philosophical Review* 117, 77–98.
- von Fintel, K. and I. Heim (2011). *Intensional Semantics*. Lecture Notes. URL: <http://mit.edu/fintel/IntensionalSemantics.pdf>.
- Weatherson, B. (2009). Conditionals and indexical relativism. *Synthese* 166(2), 333–357.
- Weatherson, B. (2011). Defending interest-relative invariantism. *Logos and Episteme* 2(4), 591–609.
- Weber, C. (2011). Centered communication. Unpublished Ms.
- Weiner, M. (2006). Are all conversational implicatures cancellable? *Analysis* 66(2), 127–30.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism, and knowledge of knowledge. *The Philosophical Quarterly* 55(219), 213–35.
- Wright, C. (2006). Intuitionism, realism, relativism and rhubarb. In P. Greenough and M. P. Lynch (Eds.), *Truth and Realism*, pp. 38–60. Oxford: Oxford University Press.