



OXFORD



Relative Truth

edited by

MANUEL GARCÍA-CARPINTERO
AND MAX KÖLBEL

RELATIVE TRUTH

This page intentionally left blank

Relative Truth

Edited by
MANUEL GARCÍA-CARPINTERO
AND
MAX KÖLBEL

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

Published in the United States
by Oxford University Press Inc., New York

© the several contributors 2008

The moral rights of the authors have been asserted
Database right Oxford University Press (maker)

First published 2008

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organization. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover
and you must impose the same condition on any acquirer

British Library Cataloguing in Publication Data
Data available

Library of Congress Cataloging in Publication Data
Data available

Typeset by Laserwords Private Ltd., Chennai, India
Printed in Great Britain
on acid-free paper by
Biddles Ltd., King's Lynn, Norfolk

ISBN 978-0-19-923495-0
ISBN 978-0-19-923494-3 (Pbk.)

1 3 5 7 9 10 8 6 4 2

Contents

<i>Preface</i>	vii
<i>List of Contributors</i>	ix
1. Introduction: Motivations for Relativism <i>Max Kölbel</i>	1
I. RELATIVISM ELABORATED	
2. Moderate Relativism <i>François Recanati</i>	41
3. Semantic Relativism and the Logic of Indexicals <i>Stefano Predelli and Isidora Stojanovic</i>	63
4. Truth in the Garden of Forking Paths <i>John MacFarlane</i>	81
5. Margins for Error in Context <i>Denis Bonnay and Paul Egré</i>	103
6. Relativism, Vagueness and What is Said <i>Manuel García-Carpintero</i>	129
II. THE METAPHYSICAL SIGNIFICANCE OF RELATIVISM	
7. Relativism about Truth Itself: Haphazard Thoughts about the Very Idea <i>Crispin Wright</i>	157
8. Three Forms of Truth Relativism <i>Iris Einheuser</i>	187
III. OBJECTIONS TO RELATIVISM	
9. Assertion, Belief and Disagreement: A Problem for Truth-Relativism <i>Sebastiano Moruzzi</i>	207
10. Frege, Relativism and Faultless Disagreement <i>Sven Rosenkranz</i>	225
11. Epistemic Modals and Correct Disagreement <i>Richard Dietz</i>	239

IV. ALTERNATIVES TO RELATIVISM

12. Content Relativism and Semantic Blindness	265
<i>Herman Cappelen</i>	
13. Faultless or Disagreement	287
<i>Andrea Iacona</i>	
14. Presuppositions of Commonality: An Indexical Relativist Account of Disagreement	297
<i>Dan López de Sa</i>	
<i>Index</i>	311

Preface

In October 2004 we began organizing a workshop on “Relativizing Utterance Truth”, which was to be held 5–6 September 2005 in Barcelona. It seemed to us that there were a number of people at that time, who were working on issues to do with new ways of relativizing truth in semantics. We knew some of these people, but we thought that there might be others. So we invited a number of participants directly and also issued a call for papers. The call included this description of the workshop’s topic:

Fregean orthodoxy has it that the contents of speech (thoughts) have absolute truth-values. If one thinks one has identified the content of an utterance and the presumed content is one whose truth-value is still relative to some parameter, then one has not succeeded in identifying the content. Most philosophers of language follow this Fregean principle even today. Those who prefer not to speak of Fregean contents usually accept an analogous principle concerning utterances: that utterances of declarative sentences have absolute truth-values.

This orthodoxy has recently been challenged for a variety of different reasons. Some claim that relativizing the truth of utterances to moments of assessment is the only good way to avoid determinism. Some claim that the only way to make room for faultless disagreement is to relativize the truth of propositions. Some claim that the best semantics for epistemic modals involves relativized truth at the level of utterances. Some forms of supervaluationism about vagueness might also be seen as employing this strategy. There are further potential examples.

Thus relativizing utterance or propositional truth is a novel semantic strategy which is motivated by a variety of different phenomena. The purpose of this workshop is to bring together some proponents (and possibly opponents) in order to discuss any aspect of this topic.

The response to the call was surprising. We had been expecting to track down a few more people who might also be interested in this very specialized topic. But in the end we received far more submissions than we could accommodate and we were forced to turn away very capable people.

The workshop was held in September 2005 and was a great success. It was very much a workshop: many of the papers were trying out new ideas, and the discussions were lively. After the workshop, we thought that given the narrow focus and pioneering spirit of the workshop, it would be good to collect some of the work presented there in an anthology. We also looked for some additional contributors, who hadn’t been at the workshop. The result, after a process of collaborative editing, is the current volume. Eight of the fourteen articles have been developed out of contributions to the workshop.

We hope that this book will advance the current debate about relativism, which seems to be of interest to an increasing number of researchers in a growing number of areas.

We are grateful to a number of people for their help in putting together this book. First and foremost the contributors, who not only put a lot of work into their own contributions, but also shared the editorial burden by commenting on the work of others. We are especially grateful to John MacFarlane for his support throughout the process. A big thanks is also due to the members of *LOGOS—Logic, Language and Cognition Research Group*, who helped with the organization of the workshop, and to Alex Miller and Alessandra Tanesini, who helped with the selection of papers for the workshop. An anonymous referee for OUP provided helpful comments. Finally, we would like to thank Chiara Panizza for preparing the Index.

We gratefully acknowledge financial support from the Catalan government in the form of a research fellowship, as well as from the Spanish Ministry of Education and Science (project HUM2006–08236) and the General Directorate for Research.

M.K. and M.G.C.

List of Contributors

Denis Bonnay, Universiteit van Amsterdam

Herman Cappelen, University of St Andrews and University of Oslo

Richard Dietz, University of St Andrews

Paul Egré, Institut Nicod, Paris

Iris Einheuser, Duke University

Manuel García-Carpintero, Universitat de Barcelona

Andrea Iacona, Università dell'Aquila

Max Kölbel, University of Birmingham and Universitat de Barcelona

Dan López de Sa, Universitat de Barcelona and University of St Andrews

John MacFarlane, University of California, Berkeley

Sebastiano Moruzzi, Università di Bologna

Stefano Predelli, University of Nottingham

François Recanati, Institut Nicod, Paris

Sven Rosenkranz, Freie Universität Berlin and Universitat de Barcelona

Isidora Stojanovic, Institut Nicod, Paris

Crispin Wright, New York University

This page intentionally left blank

1

Introduction: Motivations for Relativism

Max Kölbel

The truth of an utterance depends on various factors. Usually these factors are assumed to be: the meaning of the sentence uttered, the context in which the utterance was made, and the way things are in the world. Recently, however, a number of cases have been discussed where there seems to be reason to think that the truth of an utterance is not yet fully determined by these three factors, and that truth must therefore depend on a further factor. The most prominent examples include utterances about values, utterances attributing knowledge, utterances that state that something is probable or epistemically possible, and utterances about the contingent future. In these cases, some have argued, the standard picture needs to be modified to admit extra truth-determining factors, and there is further controversy about the exact role of any such extra factors. All the essays in this volume are about this issue. It is a narrowly defined issue in the philosophy of language, but one with important connections to other areas of philosophy, such as metaethics, metaphysics, and epistemology.

In this introductory essay, I shall attempt to provide a systematic overview. I shall outline the standard approach to semantics, modifications of which have been demanded. I then look at one of the cases, matters of taste, and examine in some detail what the motivations for modifying standard semantics are. I shall then outline the parallel reasoning in a series of other cases. Finally, I shall give a brief summary of the contributions to this volume.

1.1 A STARTING POINT

There is a picture of natural language meaning that is familiar from the likes of Frege, Carnap, Searle, Kaplan, Lewis, Stalnaker, and many others. In this picture, *propositions* take centre stage. Propositions are the objects of belief, assertion, supposition, etc, and it is the expression of propositions (be it assertoric or otherwise) that linguistic communication revolves around. In central cases of communication, a speaker utters an assertoric sentence and thereby expresses a proposition. This proposition may or may not be something this speaker believes,

and may or may not be something the audience comes to believe. Though in some sense perhaps believing what one asserts and believing what others assert is the paradigm of smooth communication.

Propositions are not only the objects of thought and speech, they are also truth-bearers, and it is often thought that propositions are *absolute* truth-bearers in the sense that a proposition by itself determines a truth-value, or at least a truth-value given the way the world is. Frege seems to be the main source for this view. In “The Thought”, for example, he concludes that sentences containing indexical expressions like “yesterday” are “not the complete expression of a thought”.¹ He concludes this from the fact that the very same sentence can be true at one time and not true at another. This marks the birth of the modern theory of indexicals: the meaning of a sentence does not associate the sentence directly with a proposition. Rather, the sentence’s meaning provides a rule or function that takes one from a *context of use* to a proposition. Thus, if the sentence “Yesterday was Sunday.” is used, it will express a proposition concerning the day before the day on which the utterance takes place. Frege’s argument is an early example of what has recently been called a “context-shifting argument” (Cappelen and Lepore 2004), i.e. an argument that begins from the premiss that two utterances of the same sentence differ in truth-value when uttered in contexts that differ in certain respects, and concludes that the sentence must be context-sensitive in such a way that the proposition it expresses depends systematically on those aspects of the context. The aspects of the context on which context-sensitivity can turn can of course include the speaker of the context, the audience, its location, etc. (“my uncle”, “you”, “over here” etc.).

The basic model according to which a sentence’s meaning together with the context determines the proposition expressed is more or less accepted by most theorists.² However, not everyone has always agreed with the Fregean assumption that propositions have their (actual) truth-values absolutely. Most notably Prior (1967) and Kaplan (1977) have allowed *tensed* propositions, i.e. propositions whose truth-value varies with time. Thus, the sentence “It is Monday.” can be seen to express the same, *tensed* proposition, whenever it is used. This proposition is true on Mondays only, and false on other days. Others have toyed with the idea of otherwise perspectival propositions, such as the proposition that one’s

¹ Cf. Frege 1918: 64. For some more of Frege’s reflections on the absoluteness of truth, see 1918: 68–9 and [1915], 271.

² There is considerable controversy about extent to which the conventional meaning of a natural language sentence consists in a *specifiable* rule that takes one from context to propositional content. “Contextualists” stress the pragmatic nature of this, which may involve processes of “free enrichment” (Travis 1997, Wilson and Sperber 2001, Carston 2002, Recanati 2004), while more traditional semanticists (Kaplan 1977, Lewis 1980, Stanley 2001, Borg 2004, Lepore and Cappelen 2004, Predelli 2005) insist that the semantic content of an utterance is determined merely by a conventional semantic rule and the context of use.

own pants are on fire, as opposed to the proposition that some particular person's pants are on fire.³

Traditionally, then, the truth-value of a proposition has been thought to be relative to a possible world only, and there has been dissent from some quarters who think truth should be further relativized to times and agents. In this volume a number of new proposals of relativizations are under discussion. Thus, some argue that the truth of utterances of sentences like "Matisse is better than Picasso." depends on a standard of taste. Others argue that the truth of utterances of sentences like "Anna knows she has hands." depends on practical interests; and that the truth of utterances of "Greece might win the World Cup." depends on a state of knowledge.

Suppose we want to say that the truth-value of utterances of the sentence "Matisse is better than Picasso." varies with a standard of taste. We shall review reasons for this type of claim in more detail later on, but suppose our reason is that we have noticed that the correctness of an utterance of the sentence depends on the standard of taste of the utterer, and we are convinced that this is to do with a variation in the truth-value of the utterances in question (and not a variation in some other value, such as conversational appropriateness, sincerity etc.). There are two basic ways of construing the dependence, a more standard way and a more innovative way.

First, it may be a case of ordinary context-sensitivity or indexicality. Two utterances of "Yesterday was Sunday." will express different propositions if they are made on different days, and as a consequence they can vary in truth-value. Similarly, two utterances of "Matisse is better than Picasso.", on this more standard view, will express different propositions if they are made in suitably different contexts, and this difference in proposition expressed can account for the observed difference in truth-value. One way of implementing this proposal would be to say that the utterer's standard of taste, or preferences, are the contextual factor in question. Thus, on this proposal, the sentence "Matisse is better than Picasso." would exhibit the same sort of context-sensitivity as the sentence "According to standards of taste like my own, Matisse is better than Picasso.". Both sentences will express different propositions in the mouths of suitably different speakers.

Secondly, the dependence of the truth of utterances of the sentence on a standard of taste could be construed in analogy with tensed propositions. On a temporalist view, the sentence "MK is hungry." expresses the same proposition

³ For example, Lewis 1979a and Perry 1979. Tensed propositions seem now to be out of fashion—a different, quantificational treatment of the tenses has become more customary, apparently without any conclusive reason for this development (see King 2003). Perspectival propositions, on the other hand, are alive and well, in the guise of propositions that are identical to sets of "centred possible worlds" (see for example Chalmers 2002 and Stalnaker 2001).

whenever it is used. However, the truth-value of that proposition varies with time: on most days it is true before lunch and (fortunately) not true after lunch. Similarly, on the second proposal, the sentence “Matisse is better than Picasso.” expresses the same proposition whenever used, but that proposition varies in truth-value with a standard of taste, so that the proposition can be true relative to one standard of taste and false relative to another. Just like temporalism, this view will need to spell out further under what conditions uttering this sentence counts as correct—the envisaged motivating datum was that the correctness of an utterance of the sentence depended on the utterer’s standard of taste. One possible answer would follow the paradigm of temporalism and construe the truth of an utterance as the truth of the proposition expressed relative to the standard of taste of the person who made the utterance. Another, apparently more radical answer would maintain that even the truth of utterances remains relative to a standard of taste. I shall return to this difference.

Let me summarize briefly. The most familiar picture of linguistic meaning has it that the meaning of a sentence centrally consists in a rule or function that determines the content or proposition the sentence would express for each context in which the sentence may be used. This proposition in turn determines a truth-value for each way the world may happen to be (each possible world), and in particular it determines a truth-value given the way the world actually is. An older challenge to the traditional picture concerns whether the truth-values of propositions depend on a time or an agent. The focus of this book is whether there are novel truth-determining factors, such as standards of taste and states of knowledge, and how exactly such a determination relation should be construed. The two basic rival options are as follows. First, the view that the sentences in question merely exhibits a hitherto unnoticed contextual dependence analogous to indexicality: the same sentence expresses different standard propositions in different contexts. Secondly the view that the sentences in question express non-standard propositions that exhibit a relativity of truth analogous to that postulated, for example, by temporalists: the sentences in question express non-standard propositions whose truth-values are relative to an extra factor.

I shall call views of the second kind “relativist”. Relativism is therefore the view that some propositions vary in their truth-value with some parameter(s) over and above the possible world parameter.⁴

⁴ My definition of “relativism” is thus fairly wide and differs from MacFarlane’s. MacFarlane wants to reserve the label “relativist” for those who claim that the truth of *utterances* is relative (2005, 325). Saying that propositions vary in truth-value with a non-standard parameter is not sufficient for relativism in MacFarlane’s sense, for even with non-standardly relativized propositions one can still define utterance truth in such a way as to be absolute (this corresponds to the first of the two possible answers mentioned two paragraphs back).

1.2 STANDARD SEMANTICS

In order to examine how any of these relativizations can be motivated, I propose to look at the foundations of natural language semantics in a little more detail. I shall begin by looking at the form of semantic theories, and the role of truth-predicates in them. It will turn out that semantic theories for natural languages define a three-place truth-predicate applicable to sentences, and that some extra-semantic principles are needed in order to relate this semantic truth-predicate to truth in any pre-theoretic sense.

We can provide a semantics for a *formal* language by assigning *semantic values*, or *extensions* to the basic expressions of that language together with some rules as to how the semantic values of complex expressions are determined by those of the basic expressions from which they are built up. In the case of sentences (one kind of expression), the semantic values are truth-values. Given that a language consists of a body of sentences, a semantics for that language can take the form of a recursive definition of a one-place truth-predicate $T(s)$ ranging over the sentences s of the language in question.

In extending this approach to natural languages, the assignment of semantic values is initially complicated by two fundamental problems: non-extensional phenomena and context-sensitivity.⁵ Let us look at each of these in turn.

First, natural languages contain complex expressions whose extension is not fully determined by the extensions of its component expressions. For example, the extensions (truth-values) of “Joe believes that Hesperus is Phosphorus.” and of “Bush might not have won the election.” seem to depend not just on the extensions of the component words. The standard solution originates with Carnap (1947): we take expressions to have their extensions not absolutely but relative to *possible worlds*. In other words, expressions have *intensions*, and for each possible world an intension determines an extension. This allows one to treat non-extensional contexts as intensional operators. For example, we can say that the sentence resulting from prefixing a sentence s with the operator “possibly” is true at a possible world just if there is a possible world (accessible from it) at which s is true. The familiar semantic clause runs something like

(Poss) For all s, w :

$T(\text{“Possibly”} \wedge s, w)$ iff there is a w^* accessible from w such that $T(s, w^*)$.⁶

Similarly, we can say that the result of prefixing a sentence s with “Joe believes that” is true at a world just if s is true at all possible worlds that Joe does not

⁵ There will of course be further and detailed problems, for example problems of formalization.

⁶ “ $\wedge$ ” is a concatenation sign in the metalanguage.

rule out at that world.⁷ Whatever the details, the semantics now defines not a one-place truth-predicate $T(s)$, but a two-place truth-predicate $T(s, w)$ relating sentences to the possible worlds at which they are true.

Secondly, semantic theories for natural languages need to make room for context-sensitive expressions or *indexicals*. Some expressions do not have the same extension whenever they are used. For example, the very same expression “my uncle” will have a certain extension if used by me and a different extension if used by you (except if you are my sibling). Thus, it seems that extensions also depend on the *context of use*. This motivates adding an argument place to the sentential truth-predicate: sentences are true relative to contexts of utterance.

For reasons explained by Kaplan (1977: 507–10), these two relativizations of the sentential truth-predicate—to a possible world and to a context of use—must be treated as separate. We need “double-indexing”. The reason, briefly, is that we want to distinguish between necessity (truth at all possible worlds) and logical truth (truth at all contexts of utterance). As a result, a standard semantic theory for a natural language will define not a two-place, but a three-place predicate $T(s, c, w)$, relating a sentence with a context in which it may be used and a possible world. This is called a “double-index semantics”.

In the previous section, I have already introduced one very familiar way of thinking about double-indexing, which involves portraying it as a two-stage process, in which a sentence in a context of use determines a proposition, and a proposition in a possible world determines a truth-value. The best-known exposition of this type of approach is probably Kaplan’s “On Demonstratives” (but see also Lewis 1980). Kaplan thinks of a semantic theory as assigning to each expression a *character*, where a character is just a function from *contexts of use* as arguments to *contents* (i.e. intensions) as values. Contents in turn are functions which take *circumstances of evaluation* as argument and have extensions as values. On the standard picture, circumstances of evaluation are merely possible worlds. However, as we shall see, some theorists, including Kaplan himself, add further factors to the circumstances, such as a time or a person.

There are different ways of construing contexts of use. Some think of a context of use as an n -tuple $\langle a, l, t, w, \dots \rangle$ of an agent, a location, a time, a possible world etc. (e.g. Kaplan 1977, Predelli 2005). Others construe contexts of use as a concrete situation in which a sentence is—or can be—uttered (Lewis 1980). But a context does a certain job in the semantics, and it has to be rich enough to do that job. A context’s job will be to determine for each indexical the intension that indexical has with respect to that context. If the semantics is for a natural language like English, then each context must at least determine a speaker, a location, a time, a world, etc, because of indexicals like “I”, “here”, “now”, “actual”. Thus, in whichever way we construe contexts, we must be able to make

⁷ For one account along these lines, see Stalnaker 1987.

sense of the idea of the speaker, location, time, world etc. *determined by* a context (or the speaker, location, time etc. *of* a context).⁸

A standard semantics will define the sentential truth-predicate $\text{True}(s, c, w)$ indirectly via another, more basic semantic value, namely the truth, or satisfaction, of a *formula* at a context, at a world and with respect to an assignment of values to the individual variables ($\text{Sat}(s, c, w, a)$). This is required for a compositional treatment of the quantifiers. I will ignore this complication in what follows and concentrate on closed sentences.

The job of a standard semantic theory, then, is to deliver a definition of a three-place sentential truth-predicate $\text{True}(s, c, w)$. The definition will be recursive and involve compositional clauses like (Poss) above.⁹ Such a theory is, however, a purely abstract theory, that does not, *by itself* have any observable consequences. This would be so even if the sentences for which we are defining $\text{True}(s, c, w)$ were exactly the sentences (or phonetic types) of the English language. That this is so becomes clear if one considers the fact that we have no pre-theoretical data concerning a three-place sentential truth-predicate that relates sentences with contexts and possible worlds. Such a predicate is a technical term introduced by semanticists to express a theoretical notion. Pre-theoretical observations are of a different kind. In general, they are data concerning the correct use of language.¹⁰ They might be data concerning the conditions under which utterances are *correct*, perhaps in various different respects, or which utterances say something true. Or they might be data concerning which sentences are synonymous, or utterances of which sentences are incompatible with the utterance of which other sentence. So what, then, are the testable predictions that a semantic theory for a natural language makes?

Semanticists regularly, and often silently, assume a principle of application that relates the claims made in a theory of semantic content for a natural language to pre-theoretical claims about that language.¹¹ One simple such principle is that adopted by Kaplan when he defines a notion of truth for “occurrences” of sentences, which is supposed to correspond to a pre-theoretical

⁸ On the n -tuple approach to context, it makes sense to restrict the set of contexts $\langle a, l, t, w, \dots \rangle$ to those where a is located in l at t and at w .

⁹ The compositional clauses for “Possibly” will normally concern *formulae* in general, so (Poss) is only a special case of the sort of compositional clause one would find regarding a possibility operator in a semantic theory.

¹⁰ Some purists (e.g. some corpus linguists) might say that all we can observe is the time and place at which tokens of sentences are produced, not the correctness of such performances. However, natural language semanticists usually operate with data concerning which utterances (real or imagined) are correct, felicitous, or true. I am assuming that this practice is legitimate because (a) semantics models competence, i.e. correct usage, not actual performance and (b) ordinary competent speakers’ judgements are good indicators (*ceteris paribus*) as to which utterances under what conditions are correct or true.

¹¹ I myself have been making such an assumption in my exposition so far, for example when explaining why semantic theories for natural languages should be double-indexed. For a sophisticated discussion of the relationship between semantics and language use see Predelli 2005.

notion of truth we have for occurrences or utterances of sentences. On Kaplan's definition, an *occurrence* of a sentence in a context is true just if the proposition expressed by the sentence in that context is true with respect to the possible world of that context.¹² In other words, an occurrence of a sentence in a context is true just if the sentence is true (in the sense defined by the semantics) in the context of use of that occurrence and at the world of that occurrence:

- (P) For all sentences s and contexts c :
 an occurrence of s in c is true just if $T(s, c, W(c))$.
 (where $W(c)$ is the world of c)

Given such a bridging principle, we can use what the semantics says about the truth of sentences relative to contexts and worlds to generate predictions as to which occurrences will be true under what kinds of condition.

To give a concrete example, let's consider a semantics that delivers a clause like this:

- (1) For all c, w : $T(\text{"I am hungry."}, c, w)$ iff in w the speaker of c is hungry at the time of c .

In what sense, if any, would clause (1) be intuitively correct? Once we assume (P), we can generate predictions as to the truth of particular utterances of the sentence. For example, let's consider two utterances of the sentence, both by Oscar, one at noon, one at 2 p.m. Suppose Oscar is hungry at noon but not hungry at 2 p.m. Then intuitively, we would say that the first utterance is true, while the second utterance is not. This intuitive result is also predicted by the semantics and (P), for from the assumptions we made:

- (a) u is an occurrence of "I am hungry." in a context $c1$.
- (b) v is an occurrence of "I am hungry." in a context $c2$.
- (c) $S(c1) = S(c2) = \text{Oscar}$, $T(c1) = 12 \text{ noon}$, $T(c2) = 2 \text{ p.m.}$
- (d) Oscar is hungry at 12 noon. Oscar is not hungry at 2 p.m.

we can generate the prediction that u is true and v is not.¹³ To the extent to which the theory's predictions conform to the intuitive judgements as to the truth

¹² Cf. Kaplan 1989: 522. Similar definitions of truth for utterances or tokens are offered by Lewis 1975: 172, and Chalmers (forthcoming, principle (T_3)). This principle is discussed in detail under the label "the Classic Reduction" by Predelli and Stojanovic in their contribution to this volume.

¹³ (1*) $\forall_w [T(\text{"I am hungry."}, c1, w) \text{ iff in } w S(c1) \text{ is hungry at } T(c1)]$ instance of (1)

(1**) $\forall_w [T(\text{"I am hungry."}, c2, w) \text{ iff in } w S(c2) \text{ is hungry at } T(c2)]$ instance of (1)

(c*) $\forall_w [T(\text{"I am hungry."}, c1, w) \text{ iff in } w \text{ Oscar is hungry at } 12 \text{ noon}]$ from (1*) and (c)

of utterances, the theory is confirmed. Without some bridging principle like (P), a semantic theory is not an empirically testable hypothesis, and any impression that some formal semantics for a natural language is intuitively correct will rely on the assumption of some such principle.

Evidently, application principle (P) represents a considerable simplification. To begin with, we don't ordinarily call *utterances* or *occurrences* of sentences true. When we do apply some notion of truth to utterances or occurrences of sentences, then it is usually in an at least semi-technical context.¹⁴ In ordinary, non-theoretical contexts, we might make judgements about whether for example what someone said, asserted, claimed, believed, or supposed is true. Our intuitive judgements about truth therefore seem to be about the truth of the *objects* of assertion, belief, etc., i.e. of propositions. Another difficulty is that linguistic utterances can convey a host of different propositions, each of them triggering intuitions as to truth. For example, a hyperbolic utterance of "I could eat an ox." might carry a false literal content and a true communicated content. Which of these is relevant to adjudicating the output of a semantic theory?

One response to this sort of difficulty would be properly to take into account the *illocutionary* aspects of natural language. In addition to the contents or propositions that sentences express in contexts of use, sentences also have illocutionary force markers, i.e. they can be declarative, interrogative, or imperative. When a natural language sentence is uttered, a proposition is not just "expressed", it is also, for example, asserted. If we add a theory of illocutionary force to a pure semantics of the sort already described, we get predictions of the form: the utterance of *s* in *c* is an assertion that *p*. Or: the utterance of *s* in *c* is a request that *p*.¹⁵ If we have a pre-theoretical grasp of asserting and requesting, we would then be in a position to confirm or

(c**) $\forall_w [T(\text{"I am hungry."}, c2, w) \text{ iff in } w \text{ Oscar is hungry at } 2 \text{ pm}]$ from (1**) and (c)

(P*) *u* is true iff Oscar is hungry at 12 noon in $W(c1)$. From (c*), (a) and (P)

(P**) *v* is true iff Oscar is hungry at 2 p.m. in $W(c2)$. From (c**), (b) and (P)

The result that *u* is true and *v* is not now follows from (d), (P*) and (P**).

¹⁴ Cf. § 7 of MacFarlane (this volume).

¹⁵ A new application principle might look like this:

(P*) An utterance of a sentence *s* in context *c* is an assertion that *p* if

- (i) *s* is assertoric (and *c* meets certain normality conditions)
- (ii) For all *w*: $T(s, c, w)$ iff *p* is true in *w*.

The general approach corresponds to that outlined in McDowell 1980. McDowell envisages a reductive account of assertion which will enable us to test the bipartite theory, consisting of a semantics and a theory of illocutionary force, against a general psychological theory. No such psychological reduction of assertion and of other propositional acts is in sight. However, I believe an account of assertion and other illocutionary acts that identifies the norms characteristically associated with each act type can fill this gap.

disconfirm the combination of the semantics with the modified application principle.

In my view, this is the most promising approach to natural language semantics. However, it would lead us too far astray if we entered a detailed discussion of assertion and other illocutionary acts at this point. Fortunately, a shortcut is available. We are currently interested in a debate about whether truth as it occurs in semantics ought to be relativized to a novel parameter, and if so, how exactly this relativization should be implemented. For these purposes, it will be sufficient to concentrate on assertoric utterances only, and on the normative role of truth in theories of assertion.

Any theory of assertion is committed to the view that truth is *in some sense* a norm for assertion. It is relatively uncontroversial to say that the norms of assertion require asserters to assert only true propositions. Different theorists of assertion might have different views of the status of this requirement, some may say that it is merely a derived norm, others that it is central or basic. For example, some may think that this norm derives from a basic assertoric norm that requires sincerity (“Assert only what you believe!”) and a general norm requiring the truth of beliefs (“Believe only what is true!”). However, most will agree that it is *a* norm in the sense that if a proposition lacks truth then the norms of assertion require in some sense that it not be asserted (even though, of course, there are also other requirements which may override this requirement). We can articulate the norm as follows:

- (A1) An assertion of a proposition p in a context c is correct (violates no norms of assertion) only if p is true in the possible world of c .

The idea is that, given the norm requires avoidance of untrue assertions, an assertion is correct or fault free (in the sense of violating no norms at all) only if the proposition asserted is true. Assuming that an utterance of a sentence s in a context c expresses the proposition that is true in exactly the worlds w in which $T(s, c, w)$, and also restricting ourselves to assertoric, literal utterances, we thus get a principle that resembles the left-to-right direction of Kaplan’s principle (P):

- (A2) For all sentences s and contexts c :
 an utterance of s in c is correct only if $T(s, c, W(c))$.
 (where $W(c)$ is the world of c)

In other words, as long as we are speaking about assertoric utterances only, an utterance of a sentence violates no norms only if the sentence is true relative to the context in which the utterance was made and to the world in which the utterance was made.

(A2) avoids some of the more obvious difficulties with the very simple principle (P) familiar from Kaplan and others. However, it is still sufficiently strong to allow intuitions about the correctness of utterances to constrain semantic theorizing somewhat. In particular, it is strong enough to serve the purposes of those who

wish to argue that standard semantics is in need of modification (as we shall see in a moment).

1.3 HIDDEN INDEXICALS IN SENTENCES EXPRESSING JUDGEMENTS OF TASTE

(A2) has the following useful consequence: suppose there is one (literal, assertoric) utterance of a sentence s and another (literal, assertoric) utterance of that sentence's negation $\text{not-}s$. Suppose further that the contexts in which these utterances are made are similar in all the respects that matter for the assignments of content to any indexicals contained in s . Then we can conclude that at most one of the utterances is free from faults. The reasoning is as follows: given that the contexts are similar in any relevant respects, and given that one of the sentences uttered is the negation of the other, the two utterances must express contradictory propositions. Thus any world in which the proposition expressed by the first utterance is true is a world in which the proposition expressed by the second is not true, and vice versa. So, if one utterance is an assertion of a true proposition, the other is an assertion of an untrue proposition. Thus, we can infer (iii) from (i) and (ii):

- (i) u is an utterance of s in C1, v is an utterance of $\text{not-}s$ in C2.
- (ii) s expresses the same proposition in C1 and C2.
- (iii) u and v are not both fault free.¹⁶

By contraposition, if two utterances of contradicting sentences s and $\text{not-}s$ are both fault free, then the contexts C1 and C2 in which the utterance were made must be relevantly different, i.e. different in such a way that s would express different propositions in C1 and C2.¹⁷

The contrapositive principle can be applied in the following case. Suppose Anna asserts that Depp is more handsome than Pitt and Barbara asserts that he is not. Intuitively, I believe, many would be inclined to say that under certain conditions it is possible that neither Anna nor Barbara has committed any fault. These conditions might include that Anna and Barbara both believe what they

¹⁶ For a more careful exposition of this type of argument, and one that is more sensitive to the position of intuitionists, see my 2003.

¹⁷ This pattern of argument has certain similarities with what Cappelen and Lepore 2004 call "context-shifting arguments". Context shifting arguments start from a premiss that says that two utterances of the *same* sentence have opposing truth-values, and conclude (via a principle like (P)), that the sentence must be context-sensitive and for that reason express different contents in the two utterances. The style of argument here discussed is more powerful because (a) it does not need to rely on a simplification like (P), and (b) its starting premiss claims an *absence* of norm violations, i.e. a kind of claim that is less liable to reinterpretation.

say, and they believe so for good reasons. Now, let's imagine exactly such a case and consider the following two utterances:

(1a) Anna (in C1): "Depp is more handsome than Pitt."

(1b) Barbara (in C2): "Depp is not more handsome than Pitt."

Now, if we want to maintain (and ex hypothesi we do) that neither Anna nor Barbara is at fault, we must conclude that the contexts C1 and C2 in which utterances (1a) and (1b) are made are relevantly dissimilar. And "relevantly dissimilar" here means: the sentence "Depp is more handsome than Pitt." contains an indexical element which is sensitive to a feature of context with respect to which C1 and C2 differ. In other words, standard semantics forces us to conclude that the sentence "Depp is more handsome than Pitt." expresses different propositions in C1 and C2.¹⁸ Consequently, (1a) and (1b) do not express contradictory propositions, which explains how both utterances can be correct—despite the fact that syntactically the sentence used in (1a) is the negation of the sentence used in (1b). By analogy, an utterance of "I am the queen of England." and an utterance of "I am not the queen of England." will not express contradictory propositions if uttered by different people, even though the second sentence is syntactically the negation of the first.

What kind of indexicality could be responsible for (1a) and (1b) *not* expressing contradictory propositions? On the face of it, the sentences uttered do not seem to contain any indexical elements. So, if we are to accept the conclusion that there is indexicality here, we will have to say that it is somehow beneath the surface. Now, the names "Depp" and "Pitt" do not seem to be context-sensitive in any way that could help here. So the assumption must be that the apparent two-place predicate "is more handsome than" has a silent third argument place, one that is assigned a referent contextually. One way (but not the only way) in which all this might be true would be if the sentence used in (1a) were similar to the sentence "On my standard, Depp is prettier than Pitt.". For each utterer with a different standard, such a sentence would express a different proposition.

An indexical proposal along these lines is the only way in which a standard semantic framework can save the possibility that neither Anna nor Barbara has violated any norms, or perhaps more precisely that neither of them has violated any norms that they are subject to. But such a proposal faces a number of problems which may motivate a relativist modification of the standard framework. Let us examine some of these problems.

A fairly superficial complaint would criticize the postulation of implicit syntactic structure. However, it would be hard to argue against implicit syntactic

¹⁸ Strictly speaking, there are at least two ways in which a sentence can semantically express different propositions in different contexts. It might be that the sentence is ambiguous, i.e. has several meanings, and different meanings are relevant to the different contexts. Or the sentence might have context-sensitive meaning which causes it to express different contents in the different contexts. For simplicity I mostly ignore the first possibility.

structures in general. But arguing that in this particular case postulating implicit structure is wrong would presumably amount to showing that there are acceptable alternatives. So this line of criticism would rest on a detailed examination of the alternatives—including relativism.

A more substantial set of criticisms revolves around the impression that any indexical proposal distorts the content of the utterances, for it claims that Anna and Barbara assert (and believe) propositions that concern their own standards, even though it seems that they are merely comparing Depp and Pitt without their assertions and thoughts having any reflective content. By itself, this is a fairly vague worry. But there are a number of concrete problems that crystallize the issue.

The first problem concerns attitude and speech reports. If the sentence “Depp is prettier than Pitt.” were indexical in such a way that it expresses different propositions in (1a) and (1b), then we should expect speech-reports of such utterances to be sensitive to this difference. Speech-reports are subject to the constraint that correct speech-reports must adjust indexical elements in the utterance reported to any relevant changes between the context of the report and the context of the original utterance. For example, take our earlier example: if Oscar utters the words “I am hungry.” at 12 noon, then a report by Alistair: “Oscar said that I am hungry.” would not correctly report Oscar’s utterance, because “I” in the report would refer to Alistair. Similarly, if at 2 p.m. Oscar reports his own 12 noon utterance with the words “I said that I am hungry.”, then the report would be incorrect, or at the very least odd: the present tense of “am” in the report suggests that the reported utterance concerned the time of the report (“I said then that I would be hungry now” rather than “I said then that I was hungry then.”). The following general rule articulates some of these principles of speech reporting:

- (SR) If a sentence *s* is indexical in such a way that an utterance of *s* in context *c*₁ expresses a different proposition from an utterance of *s* in a context *c*₂, then an utterance by someone *A* of *s* in *c*₁ cannot be correctly reported in *c*₂ by using the form of words ‘*a* said that *s*.’ (where “*a*” is some term referring to *A*).¹⁹

Now consider the indexical hypothesis that (1a) and (1b) do not express contradictory propositions because “Depp is better than Pitt.” contains an indexical which is sensitive to a change in context that occurs from (1a) to (1b) (for this is our proposed explanation of how both utterances are fault free). On this hypothesis, Barbara cannot, in the context of (1b) (or in a context relevantly like it) correctly report (1a) by saying “Anna said that Depp is prettier than Pitt.” For the sentence “Depp is prettier than Pitt.” contains an indexical element which changes its content with the change from (1a) to (1b). However,

¹⁹ Single quotes are here used like corner quotes.

this prediction is incorrect, for such a report by Barbara would clearly be correct. Thus the hypothesis is wrong.

It may be claimed that expressions like “local” or “enemy” are context-sensitive without meeting constraint (SR), and thus represent a relevant counterexample. Consider two utterances (2a) and (2b), one reporting the other:

(2a) John (in el Raval, Barcelona): “I met her in a local bar.”

(2b) Ben (in Digbeth, Birmingham): “John said he had met her in a local bar.”

The bar mentioned in (2a) is not in the surroundings of the place at which the report, (2b), is made, thus it seems that if (2b) is a correct report of (2a), then (SR) is false.²⁰

However, such cases are not conclusive counterexamples to (SR). It seems to me that there are several construals of these examples on which (SR) is not affected. One construal would view report (2b) (in the envisaged context) as involving unmarked mixed quotation: the reporter is using the exact words of the reportee. Another construal, which I prefer, denies an underlying assumption about the way in which “local” is context-sensitive. Let me explain. (2b) is a counterexample only if the context-sensitivity of “local” is such that it refers to different areas in the contexts of (2a) and (2b). For example, if we assume that “local” always picks out the surroundings of the place of utterance, then clearly “local” picks out disjoint areas in the two contexts. However, it is not clear that this is how the context-sensitivity of “local” works. A more plausible suggestion is that “local” picks out the surroundings of the place salient in the context of utterance. Thus, if prior to (2b), el Raval in Barcelona has been raised to salience, then that’s the area “local” will pick out in it. If not, then not, but then (2b) would not seem a felicitous report of (2a). Thus on this construal of “local”, the case does not present a counterexample to (SR).

Nevertheless, the objection to the indexical hypothesis that is based on (SR) is not conclusive. For there might be a class of context-sensitive expressions that do not conform to (SR), i.e. expressions that do vary in content from context to context like ordinary indexicals, but whose content in attitude report contexts is determined not by the context of the report but by some other context (possibly the context of the attitude being reported).²¹ Claiming that sentences like the one used in (1a) are exceptions to (SR) would avoid the charge of being ad hoc only if there were further, independent cases. As we shall see below however, there is a surprisingly large range of cases that give rise to analogous problems. So a determined indexicalist could use this to defend herself against the charge of making ad hoc exceptions.

²⁰ This type of counterexample is based on François Recanati’s comments on my paper at the relativism workshop in Oslo in 2005. Thanks also to Darragh Byrne for discussion.

²¹ Thus, in terms used by Kaplan (1977), attitude report contexts are selective “monsters”, i.e. they are monstrous with regard to some indexicals and not with regard to others. See Schlenker 2003 for extensive discussion.

The second objection to the indexical hypothesis is this. The indexical hypothesis says that the propositions expressed by (1a) and (1b) respectively are not contradictory—they could both be true at once. As a consequence the indexical hypothesis predicts that Barbara could come to accept what Anna has asserted (and vice versa) without changing her mind. However, in reality it is clear that Barbara cannot come to accept what Anna has said without changing her mind. So the indexical hypothesis is wrong.

Obviously, more needs to be said to make this argument work.²² To begin with, one would need to consider various proposals on the exact nature of the indexicality of “Depp is prettier than Pitt.”. However, there are ways in which the indexical proposal can be saved. Let me sketch what I regard as the best move an indexicalist can make at this point.²³ Suppose the precise hypothesis is that in any context of use, the sentence “Depp is more handsome than Pitt.” expresses the proposition that the speaker of the context has a standard of taste that ranks the looks of Depp higher than those of Pitt. In other words, the sentence

(3) Depp is more handsome than Pitt.

is equivalent to the sentence

(3I) On my standard Depp is more handsome than Pitt.

in the sense that both sentences express the same proposition in any context of use. Now, how could one explain that Barbara cannot accept Anna’s assertion? For if Anna were to utter (3I), then Barbara could perfectly happily say: “What Anna said is true. But still Depp is not more handsome.”. This answer would be incoherent as an answer to an utterance of (3) instead, i.e. (1a). That this is so is not a matter of the contingent features of this situation, but rather seems to have to do with some permanent difference between (3) and (3I).

The thing to say here for the indexicalist is that (3) and (3I) do indeed differ in their meaning, even though they express the same proposition in any context. This is because the difference between them is not a difference in semantic content, but rather a difference in the presuppositions conventionally triggered by them. (3) triggers the presupposition that speaker and audience have standards that converge (at least in the case under discussion). (3I) does not trigger such a presupposition. That explains why Barbara cannot coherently accept Anna’s assertion without changing her mind, even though—according to the indexicalist—she is perfectly happy to believe the proposition Anna asserted. The propositions that are according to the indexicalist asserted by Anna and Barbara in (1a) and (1b) are compatible, however, they are not compatible on the

²² See my 2004 for some more detail on this argument with respect to the moral case.

²³ See Lopez de Sa (this volume) and Kölbel 2007 for some more detail.

presupposition that Anna's and Barbara's standards converge, a presupposition that is conventionally carried by (3).²⁴

The problems with the indexical proposal show that it is not easy to account for cases like (1a) and (1b) within the standard semantic framework. It requires admitting exceptions to (SR) as well as making complex claims about non-truth-conditional aspects of the meaning of sentences like (3). I believe that this provides some motivation for looking into possible modifications of the standard semantic framework. In the next section I shall look at relativist proposals for modifying the standard framework in the taste case. After that I shall review other areas besides matters of taste, including epistemic modals, knowledge ascriptions, and future contingents.

1.4 RELATIVISM ABOUT TASTE

It was the assumption that utterances (1a) and (1b) are both correct that led to the conclusion that the sentences uttered must be context-sensitive in some way. One way to avoid the conclusion would be to oppose the starting assumption. One might point out that it is incoherent to accept both claims as true, and that therefore it is incoherent to deny that one of them is not true. And that would seem to show that it is incoherent to deny that one of them is at fault, given constraint (A1), i.e. the principle that asserting something untrue constitutes a fault. This is an important observation. However, the intuition that the two utterances may be correct and without fault survives it. It is coherent to maintain that both (1a) and (1b) are correct in the sense that neither violates any norm *to which it's utterer is subject*. While no one can coherently accept what (1a) says and also accept what (1b) says, it remains coherent to say that it is correct for one person to accept the one, and correct for another to accept the other.

The relativist takes this idea quite seriously: there are propositions that it is correct for some people to accept (believe, assert etc.) and not correct for others. Something closely analogous is already familiar from the standard framework: for actual people it is correct to believe the proposition that Carnap emigrated to the US. But it is not correct to believe this for non-actual people from a world in which Carnap never emigrated to the US. The proposition in question is contingent and thus varies in its truth-value from one possible world to another. The novelty of the relativist proposal is simply that it aims to allow that *amongst actual people* there may be differences as to what it is correct to believe.

²⁴ Can a similar account be given claiming that the presupposition is triggered not conventionally but conversationally? I doubt this, given that the impression of incompatibility between (1a) and (1b) is quite robust and not easily cancelled by providing further information about the situation.

Thus, instead of saying that the semantic content of the sentences in question varies with the context of use, the proposal is to say that while the content remains stable the content's truth-value varies not just with a possible world but also with an additional factor. The proposal is thus (at least initially) analogous to older and better known departures from the standard picture, such as temporalism. In Kaplanian terms, the proposal is that the circumstances of evaluation consist of a possible world and another parameter, in this case a standard of taste. So, the semantics still defines a three-place sentential truth-predicate $T(s, c, e)$, but now the third relatum of the truth-relation is a complex circumstance of evaluation $e = \langle w, p \rangle$ consisting of a possible world w and a standard of taste p .²⁵

Just as one may legitimately ask "What is a possible world?", one can ask "What is a standard of taste?". My preferred answer to that question postpones the substantial issues until later: a standard of taste is a function from pairs of propositions and possible worlds to truth-values.²⁶ This answer is very abstract. However, as I observed above in § 2, semantic theories by themselves are a purely abstract affair in any case. The definition of a semantic truth-predicate for sentences is empirically irrelevant until we fill it with empirical content via an application principle on the model of (P) or (A1) and (A2). Similarly a revised semantics, which construes circumstances of evaluation as consisting of a possible world and a standard of taste, will be filled with empirical content only once we introduce some application principles.

To recall, the most simple application principle for standard semantics was (P), the principle that an utterance is true just if the sentence uttered is true in the context of the utterance and the possible world of the utterance. Or in more picturesque terms: an utterance is true just if the proposition expressed by the sentence in the context of the utterance is true at the world of the utterance. The more refined and partial constraint (A2) said that an utterance is correct (fault free) only if the sentence uttered is true in the context of the utterance and the world of the utterance. Or in picturesque terms: an utterance is correct only if the proposition it expresses is true at the world of the utterance. What would be analogous principles for a relativist semantics?

In standard semantics, we select the world at which an utterance takes place as the value of the world parameter that is relevant for evaluating the utterance as correct. One possibility is to proceed analogously in the case of the standard of taste parameter. Thus, we might say that the context of an utterance determines

²⁵ Some relativists, for example MacFarlane (e.g. this volume), increase the adicity of the truth-predicate to four, making truth a relation that relates a sentence with a context of use, a possible world, and a "context of assessment". This is merely a notational variation. The discussion that follows can easily be transposed to MacFarlane's key.

²⁶ One might restrict the functions that qualify as standards of taste by saying that standards of taste are constant with respect to propositions that do not concern matters of taste. The substantial issues would remain to be answered.

one particular standard of taste which will be relevant for evaluating utterances as correct. We would then get analogues to (P), (A1) and (A2) along these lines:

- (P^R) For all sentences s and contexts c :
 an occurrence of s in c is true just if $T(s, c, <W(c), S(c) >)$.
 (where $W(c)$ is the world of c and $S(c)$ is the standard of taste determined by c).
- (A1^R) An assertion of a proposition p in a context c is correct only if p is true at $<W(c), S(c) >$.
- (A2^R) For all sentences s and contexts c :
 an (assertoric, literal) utterance of s in c is correct only if $T(s, c, <W(c), S(c) >)$.

Now, a definition of utterance truth along the lines of (P^R), and the corresponding normative principles about assertion, are completely analogous with standard semantics in that utterance truth comes out as absolute. A sentence and an utterance context together determine whether, if an utterance in that context were assertoric and literal, it would be true, and thus can meet the norms of assertion. The situation is exactly analogous to temporalism, where the time at which a sentence is uttered is the one relevant for the evaluation of an utterance. Following Recanati's (this volume) distinction between moderate and radical relativism, let me call this kind of approach to pragmatics "moderate". (MacFarlane (2005, forthcoming) calls a relativist semantics paired with a moderate pragmatics "non-indexical contextualism".)²⁷

Before considering approaches that are not moderate, let me examine in more detail how a moderate approach could be implemented. The most important detail will be the way in which the function S , as it occurs in (P^R), (A1^R) and (A2^R) is to be cashed out. We are familiar with the W function; i.e. with the idea that the context of an utterance determines a possible world, the world at which the utterance takes place. What is new is the idea that an utterance context determines a standard of taste. There are many possibilities:

- (S1) $S(c)$ = the standard of taste of the speaker of c .
- (S2) $S(c)$ = the standard of taste of the audience of c .
- (S3) $S(c)$ = the standard of taste most relevant to the purposes of the conversation at c .

²⁷ (P^R), (A1^R) and (A2^R) work by singling out, in each context of use, one privileged standard of taste. Alternatively, one could single out a range of standards (in each context) and say that correct assertion requires truth in relation to all or some of the standards in the range, in analogy with super- or subvaluationism (compare MacFarlane 2005a: 330). I would count these variations also in the "moderate" category, given that again a sentence and a context of use together determine a truth-value of sorts—though in this case we get truth-value gaps (at the pragmatic level). I myself have proposed this kind of approach for vague sentences (Kölbel forthcoming).

There are many more possibilities.²⁸ (S1) seems to be the most popular hypothesis amongst proponents of relativism.²⁹ So I shall briefly explore one way in which (P^R), (A1^R) and (A2^R) might be applied to Anna's and Barbara's utterances (1a) and (1b), if the S function is interpreted according to (S1).

First, it is necessary to explain the idea that speakers can "have" a standard of taste. My own idea here is that if I have a standard, then I should accept only propositions that are true in the actual world and on my standard. In the case of judgements about matters of taste, which standard I have will in part depend on my personal preferences and aesthetic responses to things. Thus, for example, only someone who prefers Depp to Pitt will have a standard of taste on which the proposition that Depp is more handsome than Pitt is true in the actual world. Standard possession is therefore a theoretical notion that models our idea that it depends on certain individual speaker features (e.g. their preferences) which propositions it is correct for them to believe, and ultimately to assert.³⁰

Now suppose that Anna and Barbara have a different standard of taste due to differences in their personal preferences and dispositions to respond aesthetically. Then we can explain why both utterances ((1a) and (1b)) may be correct, even though the propositions asserted by them contradict one another. The propositions asserted contradict one another in the sense that for each circumstance of evaluation *e*, if the proposition asserted by Anna is true at *e*, then the proposition asserted by Barbara is not true at *e*, and vice versa. However, when evaluating the assertions (and beliefs for that matter) for truth, a different circumstance of evaluation is relevant to each assertion. The context of (1a) determines Anna's standard of taste, while the context of (1b) determines Barbara's standard of taste as relevant. So neither of them fails to be true in the sense of (P^R); neither of them violates the norm expressed by (A2^R). That's why, in the absence of violations of any other norms, both utterances may be without fault.

The relativist with a moderate pragmatics can therefore claim to make sense of the idea that (1a) and (1b) are both correct (fault free) without giving up on the idea that the propositions expressed by the two utterances are contradictory.

It remains for me to discuss the possibility of an "immoderate" pragmatics, i.e. a pragmatics that does not use an absolute notion of utterance truth along the lines of (P^R), (A1^R) and (A2^R) as a central assertoric norm. Departure from an absolute truth norm in pragmatics could have various motives. For example, we

²⁸ Some possibilities will not be very interesting. For example any function S with a constant value would be useless with regard to the problem of understanding utterances (1a) and (1b).

²⁹ See for example Kölbel 2002, 2003, Lasersohn 2005, Recanati this volume, MacFarlane 2008.

³⁰ It is important to be aware that on this model of standard possession, people can make mistakes: they sometimes believe or assert propositions that are not true according to their own standard. For example, Barbara may come to believe, as a result of listening to Anna's utterance (1a), that Depp is more handsome than Pitt. She might later realize that this belief is a mistake, because she prefers Pitt.

might be unable to find a suitable function *S* that determines the value of the new parameter for each context of use with respect to which assertions are to be evaluated. Or there may be a number of distinct and complementary pragmatic norms, each of which picks out a different value of the parameter as relevant. For example, utterers might be subject to the norm (a) to assert a proposition only if it is true on the audience's standard, but *also* subject to the norm (b) that they shouldn't assert a proposition unless it is true on any standard shared by all members of the audience and the speaker, and also subject to the norm (c) that they shouldn't assert a proposition unless it is true on either Paul Bocuse's standard or Ferran Adrià's standard. Perhaps these norms are weighted: some are more important, and can overrule others. Under those circumstances it might be appropriate to say that there is no single value of the standard of taste parameter that is pragmatically relevant. In that case no interpretation of *S* will yield a correct principle (P^R), or ($A2^R$).³¹

There is no reason, *prima facie*, to suspect that this sort of pragmatics cannot work.³² However, any such pragmatics will have to offer an alternative to the model of communication that moderate approaches offer. Those advocating moderate approaches can say that audiences know under what conditions an assertion is correct, and thus can, if they have reason to assume that an assertion is correct, come to believe that those conditions obtain. Any immoderate pragmatics will have to find a replacement for this story. MacFarlane (2003, 2005, this volume) has on several occasions begun to tell such a story. It would depend on the details of such a story as to how the apparent faultlessness of utterances like (1a) and (1b) can be explained. I will briefly return to immoderate approaches in the next section when discussing the case of future contingents.

This concludes my exposition of the alternative, relativist account of Anna's and Barbara's utterances, and correspondingly the relativist account of the semantics of sentences concerning matters of taste. In the next section, I shall review a number of further cases in which similar issues arise, or in which similar relativist solutions seem at least initially promising.

1.5 HIDDEN INDEXICALS OR RELATIVISM IN OTHER CASES

The problem of accounting for the possibility of superficially contradictory utterances that are nevertheless fault free, and the range of possible

³¹ Very abstractly speaking, all theories that distinguish between semantic content and content otherwise conveyed (Grice being the classic example), can be seen as immoderate pragmatics. On Grice's view, clearly various norms compete, namely the various conversational maxims, and a delicate process weighs these norms against one another.

³² Evans 1985 seems to think that only a moderate pragmatics is feasible. Recanati's and García-Carpintero's contribution to this volume discuss Evans's arguments and come to a similar conclusion.

responses—indexical or relativist—is not unique to discourse about matters of taste, even though perhaps this is the case that is most intuitive. The same problem and responses recur with remarkable similarity in a wide range of further cases, some of which I shall now briefly review. I shall only expound the initial motivation and remark on some specific differences. It will be left to the reader to restore the detail provided in the taste case to the other cases.

1.5.1 Epistemic Modals

Imagine Anna is trying to find out who emptied her bottle of whisky. At the beginning of her investigations, at time t_1 , she thinks that it might have been Barbara, because at that time she has no evidence that rules out the possibility of Barbara having emptied the bottle. Much later, at t_2 , she finds out that Barbara could not have emptied the bottle because she was not in town at the relevant time. Anna makes two utterances:

(4a) Anna at t_1 : “It might have been Barbara.”

(4b) Anna at t_2 : “It might not have been Barbara.”

Let us assume that the contexts of (4a) and (4b) are such that in both cases “it” picks out the same emptying of the same whisky bottle. Thus, on the face of it the two utterances look like the assertion and denial, respectively, of the same epistemic possibility. The utterances *seem* to express contradictory propositions: the proposition that it might be the case that Barbara emptied the whisky bottle and the proposition that it might not be the case that Barbara emptied the whisky bottle. Nevertheless, there is a clear sense in which Anna is right both times. Given that her knowledge at t_1 does not rule out the possibility in question,³³ (4a) is the right thing to say (and think) at t_1 . Given that her state of knowledge at t_2 does rule out the possibility in question, (4b) is the right thing to say (and think) at t_2 . Both utterances comply perfectly with the norms of assertion.

As in the taste case, there are two main ways in which one can accommodate the impression that both utterances are correct (without giving up on the idea that truth is a norm of assertion). First, one could make a diagnosis of hidden indexicality: the propositions expressed by (4a) and (4b) respectively are not contradictories, for the sentences used exhibit a form of indexicality that causes some change of reference from t_1 to t_2 . For example, the first sentence used might be equivalent to

(4a*) Anna at t_1 : “My current state of knowledge is compatible with it having been Barbara.”

³³ And perhaps given there is also no easy way in which Anna could have obtained such knowledge. See DeRose 1991.

As in the taste case, the indexical hypothesis generates various seemingly problematic predictions. One prediction concerns speech reports. We would expect, given (SR), that Anna cannot, at t_2 , correctly report her utterance (4a) homophonically. (4a*) clearly could not be reported homophonically at t_2 . However, contrary to this prediction (generated from (SR) and the indexical hypothesis), a homophonic report at t_2 is perfectly correct: “I said that it might have been Barbara.” Explaining this would require dropping (SR) and saying that the hidden indexicals differ from ordinary indexicals. A similar complaint is that Anna might add at t_2 : “But I was wrong. For it could not have been Barbara.”.³⁴ If the indexical hypothesis were correct, one would not expect this. One would expect her to continue to endorse the proposition asserted at t_1 : “What I said was true: my knowledge at t_1 did not rule out Barbara being the perpetrator.”. But it might be argued that this is not how we would retrospectively assess such an utterance.³⁵

The problem is that the way we report utterances like (4a) and (4b) and the way we evaluate them sometimes suggests that they have contradictory contents, and an indexical hypothesis cannot easily accommodate that. It seems that after the new information has come in Anna re-assesses the *same* content from a different perspective, namely that of the knowledge that she now has. Thus again a relativist construal is a natural alternative: the proposition in question varies in truth-value with an extra parameter in the circumstances of evaluation. The proposition that it might have been Barbara is true relative to Anna’s earlier state of knowledge, and it is false relative to her later state of knowledge.³⁶

A moderate pragmatic account could easily deliver the result we need, namely that both (4a) and (4b) are in compliance with the norms of assertion. The most plausible account, it seems, would say that an assertion of the proposition that it might be that p is only correct if *the speaker’s* knowledge is compatible with p .³⁷

1.5.2 Knowledge Attributions

One of the areas of most active debate in recent epistemology has been a position called “contextualism about knowledge”.³⁸ According to the most common form of this view, sentences used for knowledge attributions, such as “Charles

³⁴ See MacFarlane forthcoming a, for a version of this argument against an indexical view.

³⁵ But compare the objections Wright raises against this reasoning in the article that appears in this volume.

³⁶ Relativism about epistemic modals is defended by Egan, Hawthorne and Weatherson 2005 and MacFarlane forthcoming a. See also Egan 2008 and Gillies and von Fintel 2006.

³⁷ Some uses, however, would suggest that it is not always the speaker’s knowledge that is relevant. For example, consider the game show host who says: “The prize might be behind this door.” while knowing it is not. Compare Gillies and von Fintel 2006, Bach 2006, and Egan 2008.

³⁸ Cohen 1986, DeRose 1992, Lewis 1996 are classic articulations. For a recent overview of the vast literature on this topic see Rysiew 2007.

knows he has hands.”, are context-sensitive in unobvious ways. The sentence is obviously indexical because in different contexts of use, it will concern different times in Charles’ life. The contextualist’s claim is that the sentence “Charles knows he has hands.” is in addition sensitive to the standard of knowledge associated with the context of use. Thus, in a context C1 (perhaps when waking up in hospital after a horrific accident), the standards required for knowledge are relatively low: Charles’ feeling his right hand with his left and vice versa may be enough for knowledge that he has hands. But in the context C2 of discussing scepticism, when a sceptic has just raised a sceptical possibility—perhaps that of Charles being a handless brain in a vat tricked by a computer into thinking he has got feelings in his hands—the standard sufficient for knowledge is much higher. Thus the following two of Anna’s utterances may both be correct:

(5a) Anna at C1: Charles knows (at noon 1/1/2007) that he has hands.

(5b) Anna at C2: Charles does not know (at noon 1/1/2007) that he has hands.³⁹

Being able to say that knowledge attributions vary in this way with the contextual standards for knowledge enables contextualists to give an answer of sorts to sceptical arguments. It allows them to concede sceptical conclusions such as that expressed by (5b) when faced with a sceptical argument, without having to give up all or even most claims to knowledge. For (5a) and (5b) are both correct.

This result is achieved, in ordinary forms of contextualism, through the type of indexical hypothesis with which we are by now familiar. The sentence

(6) Charles knows (at noon 1/1/2007) that he has hands.

is claimed to be more or less equivalent to the sentence

(6I) By the currently salient standards for knowledge, Charles knows (at noon 1/1/2007) that he has hands.

Thus, the two utterances (5a) and (5b) do not express contradictory propositions. Rather, (5a) expresses the proposition that Charles meets C1’s standard for knowing (at noon 1/1/2007) that he has hands, while (5b) expresses the proposition that Charles fails to meet C2’s standards for knowing the very same thing. Postulating a hidden form of indexicality enables the contextualist to claim that both (5a) and (5b) are true.

The parallelism with the previous cases is not difficult to detect. Similarly, it is not difficult to see that the indexical hypothesis will be subject to some of the

³⁹ I added the time specification in order to highlight that the contextualists are not talking about ordinary time-sensitivity of knowledge attributions. In other words, they are not claiming that Charles “forgets” that he has hands from C1 to C2 (or that he “learns” that he has hands, in case C2 precedes C1).

same problems that arose for the indexical solutions in the previous cases. Thus, for example, it would seem to be correct to report (5a) in C2 by saying “Anna said that Charles knows that he has hands.”. Indeed, if one pointed out to Anna after her utterance (5b) that she said earlier that Charles does know that he has hands, then it would be quite normal for her to say: “Yes, I did say that. But I was wrong.”. These observations seem to be in conflict with epistemic contextualism in the form currently discussed.⁴⁰

Again, a form of relativism has been offered as an alternative.⁴¹ Instead of saying that sentences like (6) express different propositions in different contexts of use, due to variations in the relevant standards for knowledge, we could say that the proposition expressed remains constant but that the constant proposition expressed is unusual in that it varies in its truth-value with an extra circumstantial parameter. Rather than complicating the function from context of use to content (i.e. the character) by adding parameters to the context, we instead complicate the function from circumstances of evaluation to extensions by adding a circumstantial parameter.⁴²

Again, relativists about knowledge attributions will have to say something about how the changes in the semantics affect the pragmatics, and settle on a moderate or immoderate approach. I believe that the mainstream of contextualism is a form of *attributor* contextualism, thus the most obvious starting point would be a moderate pragmatics in which the context of use determines a relevant standard for knowledge against which that utterance is evaluated. However, MacFarlane (2005b) argues that an adequate pragmatics must allow us to evaluate the same utterance by recourse to varying standards.

1.5.3 Moral Values

The case of moral value is largely similar to that of aesthetic value already discussed at length. However, in the moral case, the driving assumption that there is some extra factor on which value depends is much more controversial. To give a very brief sketch, let us consider two utterances, one made in an Indian moral context by Arvind, one made in a Western European moral context by Barbara

(7a) Arvind in C1: One ought not to marry outside one’s own caste.

(7b) Barbara in C2: It’s not the case that one ought not to marry outside one’s own caste.

⁴⁰ See Yourgrau 1983 and Kompa 2002 for similar observations about contextualism.

⁴¹ See MacFarlane 2005b.

⁴² As pointed out earlier, MacFarlane (2005b) prefers to introduce in addition to context of use and circumstance of evaluation, a “context of assessment”. However, as he has confirmed in correspondence, the choice between adding the new parameter to the circumstances of evaluation or introducing a new context of assessment, is a matter of presentation, not of substance.

Now, to say that both of these utterances may be correct is highly controversial, by far more controversial than in the taste case. Nevertheless, some philosophers and some anthropologists have been driven by a variety of different considerations to the conclusion that moral judgements depend for their correctness on some implicit parameter such as a moral code. Usually positions of this sort are called “moral relativism”. However, the best known position that goes under that name, namely Harman’s moral relativism,⁴³ is actually a thesis of hidden indexicality. Again, (7a) and (7b) do not express contradictory propositions and this explains the (controversial) starting datum that neither (7a) nor (7b) were mistaken. Much the same problems have been pointed out with this indexicalist view, and a relativist version (in the current sense of “relativist”) has been suggested.⁴⁴ So far, the associated pragmatic approach that has been proposed is moderate.

1.5.4 Painted Leaves

There is a seemingly unrelated debate in theoretical linguistics and the philosophy of language about the general role of context in linguistic communication. Contextualists, in this debate, challenge a certain traditional model of linguistic communication associated with formal semantics. The formal semanticists believe that the linguistic meaning of a natural language sentence will determine for each context of use a proposition that it semantically expresses at that context. The route from context to proposition expressed is, according to the traditionalists, a purely mechanical one completely anticipated by the semantics of the sentence used. The paradigm for this role of context is given by typical indexical expressions such as “I” or “here” or “now”. However, according to contextualists, the route from context to content is not one that can be anticipated in this way in a semantic theory. Contextualists put much more emphasis on pragmatic processes of determining which proposition is expressed by any particular utterance. They often claim that a sentence by itself together with the context of use only determines an incomplete or truncated proposition, one that needs completion via an interpretive process of “free enrichment”. Now, the cases presented by contextualists in support of their view bear an uncanny resemblance to the sort of cases we have been discussing.

One of the best known examples is from Charles Travis (1997). Suppose Pia utters the sentence “The leaves are green.” twice, once talking to a photographer who is looking for a green motif, and once talking to a botanist looking to classify the plant in question. The twist is that the leaves that each of the utterances refers to are the leaves of a Japanese maple. As is normal with Japanese maples, the leaves were originally russet in colour. But subsequently they have been painted green. According to Travis (1997), the first utterance

⁴³ Harman 1975. See also Dreier 1990 and 1999.

⁴⁴ See Kölbel 2002, 2004, and 2007.

is true, while the second utterance is false. The lesson is supposed to be that a semantic rule alone cannot anticipate which propositions might be expressed by a sentence in which contexts, and more examples on the same model can be easily constructed.

The typical defensive move of a traditional semanticist would be to insist that even though the utterances may *convey* or *communicate* different propositions which then differ in truth-value, the proposition semantically expressed by the two utterances is the same. Thus, the formal semanticist might say that both utterances are true, but the second is misleading in that it also pragmatically conveys a false proposition, namely that the natural colour of the leaves is green. Or she might say that they are both false, but that the first one also pragmatically communicates a true proposition, namely that the leaves are superficially green. The problem Travis points out is that the formal semanticist cannot offer a reason to prefer one of these two options to the other. Now, if this problem were restricted to only one case, it might have been feasible to provide some rationale for preferring one of the options. However, given that many more examples can be constructed, this will not remove the problem.

Now, in order to preserve the continuity with the earlier discussions, let me adjust Travis's example. Consider two utterances of superficially contradictory sentences:

- (8a) Pia in C1: The leaves are green. [directed to a photographer who is looking for something green as background for his photo]
- (8b) Pia in C2: The leaves are not green. [directed to a botanist who is trying to classify the tree]

There is a strong intuition that both utterances are correct. How can we accommodate that intuition in a traditional semantic framework? As before, we have an indexical and a relativist option. According to the indexical option, the predicate "is green" is indexical in such a way that in C1, it expresses a property that applies to the painted leaves, and in C2 it expresses a property that does not apply to the leaves. Thus the propositions expressed in (8a) and (8b) do not contradict one another. By contrast to the previous cases, it is not easy to generate intuitions about speech reports or retrospective evaluations that would create a problem for an indexical proposal. However, there are other reasons why it seems ill-advised to postulate hidden indexicality in this case. If the current case provides evidence for "is green" being indexical in the way suggested, then it is likely that we will be able to find evidence for the indexicality of any predicate whatsoever. For Travis-style examples abound.⁴⁵

⁴⁵ This is the slippery slope of which Cappelen and Lepore warn us in their 2004.

The conclusion contextualists want to draw is that a formal semantic theory for a language just cannot spell out in advance the semantic contents that will be expressed by the language's sentences in all possible contexts. A form of relativism may help avoid this conclusion: instead of concluding that the propositions expressed by sentences vary in unpredictable ways from context to context, we can say that the propositions predictably expressed by sentences in contexts are evaluated, and used in communication, in unpredictable ways. The proposition expressed by the sentence "The leaves are green." remains invariant as long as we are talking about the same leaves at the same time. However, whether that proposition counts as true, whether asserting it counts as correct, depends on the specific purpose against which we are evaluating that proposition. We cannot in advance predict for what purposes people may venture to assert the proposition that those leaves are green at that time.⁴⁶ Thus again, we need to add a new parameter to the circumstances of evaluation: propositions have their truth-value relative to pairs $\langle w, p \rangle$ of a possible world w and a purpose p .

As before, once the circumstances of evaluation are enriched, the relativist also needs to say something about the normative significance of truth in pragmatics. The impression that (8a) and (8b) are fault-free will need to be explained by reference to some pragmatic account. If the purpose (or purposes) relative to which we evaluate an utterance are completely determined by the sentence used and the context of use, then a moderate pragmatics will do. However, if there is a possibility that the very same utterance may be evaluated now against this, and now against that purpose, then a more sophisticated pragmatics will need to be construed.

1.5.5 Future Contingents

The case of claims about the contingent future differs considerably from the previous ones.⁴⁷ While this is again a case where certain considerations support the view that certain propositions vary in truth-value with a non-standard parameter, the variation is never from false to true, or from true to false. Rather, the variation occurs either between having no truth-value and being true, or between having no truth-value and being false. For this reason, I shall in this

⁴⁶ Stefano Predelli's (2005) answer to Travis's problem seems to follow roughly these lines. MacFarlane has also offered a similar answer to context-shifting arguments in his forthcoming b. In MacFarlane's terminology, an invariant proposition can have a varying intension, rather than the proposition *being* the intension and varying in truth-value not just with a possible world, but also with a purpose.

⁴⁷ The main arguments presented in this subsection are drawn from MacFarlane 2003 and from MacFarlane's contribution to this volume.

case consider two utterances of the same sentence, one made in 2005, and one made in 2007:

(9a) Italy are the winners of the 2006 Football World Cup. [uttered in 2005]

(9b) Italy are the winners of the 2006 Football World Cup. [uttered in 2007]⁴⁸

The sentence used appears to be eternal (i.e. it seems to express the same proposition no matter when it is used). Since (9b) is in fact true, this would mean that (9a) must also be true. However, indeterminism requires that utterance (9a), at least when made, was not true, for it was not yet determined in 2005 that Italy would win in 2006. Mere semantics should not settle controversial metaphysical issues, so the indeterminacy in truth-value of (9a) should at least be an option within our semantic framework. So the truth of the sentence used in (9a) and (9b) should be made, somehow, to depend on a new parameter of time: the sentence is not true (and not false) in relation to times earlier than the 2006 World Cup final, but true relative to all other times.

Might we allow a variation in truth-value between (9a) and (9b) by postulating extra indexicality and saying that after all the two utterances express *different* propositions? No simple indexical solution will work, for the problem can be raised just as easily with just *one* utterance. Consider utterance (9a). It seems that before the 2006 World Cup final, (9a) is not true, because the future is open. But after the 2006 World Cup final, when considering whether utterance (9a) was true, we'll have to say that it was true. Thus *the very same utterance* is true when assessed after the crucial time and not true when assessed before. But if the variation in truth-value concerns the very same utterance, then the variation cannot be due to a variation in the context of use, for we are talking about the same utterance. What has changed is simply the perspective from which we assess the utterance.

Is this another case where we need to introduce a new truth-determining factor in the semantics? A little reflection shows that not. The way to think about the open future is to take any event, e.g. an utterance, to be "located in" or "compatible with" a range of possible worlds, namely the range of worlds that coincide with the past up to that point. Given that the past does not completely determine the future, there will be more than one possible world that shares an event's past, each one corresponding to one way the future might turn out. On this view, the semantics can remain completely standard: a definition of sentential truth $T(s, c, w)$, where s is a sentence, c is a context of use and w is a

⁴⁸ In order to maintain the parallelism with the previous cases, I have deliberately chosen a sentence where the tense does not mark the relationship between the events described and the time of utterance.

possible world. The interesting questions begin when we ask whether a particular utterance is true. If we maintain that the future is open, then a context of use will not determine one particular world of a context, but at best a range of worlds. Thus we cannot follow the model of the simple application principle (P) above, which said that an utterance is true just if the sentence used is true at the context in which the utterance was made and at the world of that context. For there is no “world of the context”, but rather a range of worlds that is compatible with the context.

MacFarlane proposes to define a notion of utterance truth that varies even if the sentence uttered and the context of use are held fixed. (In other words, such a notion of utterance truth must be immoderate, because a sentence together with a context of use do not yet determine a truth-value.) Utterance truth will vary, then, with the “context of assessment”, i.e. a concrete situation at which a proposition can be assessed. A context of assessment, just like a context of use, determines a range of worlds that are compatible with the past preceding that context. We can then define the truth of an utterance at a context of assessment in such a way that if a context of assessment is compatible only with worlds in which the utterance is true, then the utterance is true at that context of assessment.⁴⁹ Thus, (9a) as assessed now is true, but it was not true as assessed before the 2006 World Cup final (because any context of assessment prior to that time is compatible with some worlds where (9a) is false). Any proponent of such a framework will have to offer some further considerations as to how an immoderate pragmatics explains communication, for example by showing how an utterance made in one context can be subject to the norm that that utterance be true in relation to certain future contexts of assessment.

There are some alternatives to this relativist framework. One alternative to relativism is a sophisticated kind of indexicalist proposal. We could say that utterances like (9a) are indexically about one particular possible future, however, which future that is is still indeterminate at the time the utterance is made. So, no determinate proposition is expressed by (9a), and this explains why in 2005, (9a) is neither true nor false. But why is (9a) true as assessed in 2007? Because all the propositions (9a) might express in 2007 are in fact true. This is an elegant solution. However, the problem with it is that Anna’s utterance, like other utterances about the future, already seems to have a context in 2005. Anna manages to express a proposition already at the time of utterance, and, if the utterance was sincere, Anna believes this proposition in 2005, and those in Anna’s audience who believe what she told them also believe this proposition already in 2005. It would be complicated to maintain all these claims

⁴⁹ In the case in which the worlds compatible with the context of use is a proper subset of the set of worlds compatible with the context of assessment. This complication is explained in detail in MacFarlane’s contribution.

if we wanted to say that Anna's utterance only begins to express a proposition much later.⁵⁰

Another alternative is supervaluationism: the view that an utterance is true if the proposition expressed by it is true in all worlds that are compatible with the context of use in which the utterance was made, that it is false if it is not true in any of these worlds, and that it receives neither truth-value if it is true in some but not in others. This means that utterance truth is gappy: some utterances, like (9a), are neither true nor false. This does justice to our assessment of (9a) in 2005. But does it make room for our evaluation of (9a) as true when we assess it in 2007? MacFarlane's contribution to this volume will provide an answer.

1.5.6 Further Cases: Probability, Fiction, Vagueness

There are at least three further potential areas of application for the relativist strategy. First, on a subjectivist construal of probability, probability ascriptions are plausible candidates for a relativist treatment similar to that of epistemic modals.⁵¹ Whether it is right for some thinker to call some outcome probable will depend on the evidence available to that thinker or even on the probability function with which that thinker started out.

Another case of application is that of fictional utterances. The very same sentence, for example "Mozart's requiem was commissioned by Salieri," can be used to make a correct remark about Forman's movie *Amadeus*, and an incorrect remark about the real world. The standard treatment involves the postulation of contextually implicit operators.⁵² However, implicit operators can be avoided if one takes seriously the idea that the truth of a proposition can vary with an extra parameter.⁵³ The extra parameter is presumably something like a possible world, so that the world of Forman's *Amadeus* and the actual world can figure as values of such a parameter.⁵⁴ If such a parameter is allowed, we can be much more differentiated in the pragmatic evaluation of utterances as true, and thus do justice to the impression that the very same proposition—that Salieri commissioned Mozart's requiem—can be asserted correctly in one context and incorrectly in another.

Finally, a "sharpening" parameter has been proposed in the semantics of vague predicates and gradable adjectives. Supervaluationists have long used this to deal

⁵⁰ We would have to introduce new entities that can figure as the objects of Anna's assertion and belief, entities that can vary in their truth-value with the time of assessment, i.e. entities that are quite similar to the propositions of the relativist.

⁵¹ See Price 1983 and Kölbel 2002 for some discussion.

⁵² For example, Field 1973 and Lewis 1978.

⁵³ The basic idea for this is contained in Predelli 1997, 2005. However, Predelli seems to have changed his mind—see his 2006.

⁵⁴ However, the job of the new parameter cannot be done by the standard possible world parameter. For there are possible worlds in which Salieri does commission the requiem and Forman's fiction portrays him as not commissioning it.

with the sorites paradox.⁵⁵ A supervaluationist semantics starts by defining an auxiliary relativized truth-predicate $T^A(s, c, \langle w, p \rangle)$ which relates a sentence s to a context c and a pair of a possible world w and a sharpening p . They then proceed to define truth proper as truth in all sharpenings: For all s, c, w : $T(s, c, w)$ iff for all p : $T^A(s, c, \langle w, p \rangle)$. The result is a three-valued semantics with some associated problems. Now some of these problems can be avoided if we accept the auxiliary truth-predicate T^A as the ordinary semantic truth-predicate. It will then be a task for the pragmatics to deal with the extra sharpening parameter, for example by specifying that it is correct to assert a proposition only if it is true in all sharpenings, thus in effect giving a “supervaluational pragmatics”.⁵⁶

1.5.7 Outlook and Interaction of the Various Parameters

Given that I have just described nine different candidate areas for relativist treatment, the question naturally arises how the different parameters of relativization would interact. Suppose we were articulating a semantics for a language that contains, for example, both epistemic modals and predicates of personal taste. Should our semantics construe the circumstances of evaluation as incorporating two separate new factors, namely a state of knowledge and a standard of taste? The semantic truth-predicate would relate a sentence with a context of use, a possible world, a state of knowledge and a standard of taste. Or can we rationalize and add just one multi-purpose parameter?

I cannot treat this question properly here. However, it is clear that the answer will in part depend on whether there are operators (in analogy with modal operators) that shift the new parameters, and how they interact with one another. Thus, for example, suppose we have a complex operator “For N, p ” which shifts the taste parameter to the standard of taste of N , so that “For Anna, Depp is more handsome than Pitt.” is true relative to a context, world, and standard of taste just if “Depp is more handsome than Pitt.” is true relative to that context, that world, and Anna’s standard of taste.⁵⁷ Now, given that such an operator shifts the taste parameter *without* shifting the world parameter, we can’t rationalize and fuse the world and taste parameters together.⁵⁸ If there is a similar operator shifting the state of knowledge parameter *separately*, then that would presumably constitute a reason to maintain a separate state of knowledge

⁵⁵ The classic source is Fine 1975.

⁵⁶ This sort of approach was proposed by MacFarlane at the ESAP conference 2002 in Lund. One version of the approach is defended in more detail in Kölbel forthcoming. See also Richard 2004.

⁵⁷ See Kölbel 2008a. If natural languages contain constructions that are best treated as operators that shift a novel circumstantial parameter, then this would of course constitute a powerful argument in favour of introducing such a parameter. However, a case will have to be made that such an operator treatment is better than a corresponding quantifier treatment, which would naturally go with the corresponding indexicalist position.

⁵⁸ See Einheuser’s contribution to this volume.

parameter. More interesting questions arise concerning the mutual embedding of such operators and resulting scope issues.⁵⁹

Whatever the outcome of future research in this area, I believe that the number and breadth of areas in which a relativist approach can be motivated shows that there are general underlying issues here that require principled research. There may be a lot wrong with recent proposals of relativism, but it is not that they are ad hoc or lack motivation.

1.6 OVERVIEW OF THE CONTRIBUTIONS TO THIS VOLUME

For better orientation, the contributions to this volume have been ordered according to their focus. There are four parts. Part I contains essays that elaborate relativism in one way or another. Part II contains two papers considering the metaphysical consequences of relativism. Part III consists of papers that offer detailed objections to relativism. And finally Part IV comprises three papers that offer alternatives to relativism.

François Recanati opens the first part with an essay that explains the parallelism between recent forms of relativism (for example about epistemic modals and gradable adjectives) and temporalism, the view that there are tensed propositions, propositions whose truth-value varies with time. He defends these forms of relativism against two classic objections that have been levelled against temporalism, namely those by Gareth Evans and Mark Richard. However, his defense extends to moderate forms of relativism only.

Stefano Predelli and Isidora Stojanovic elaborate relativistic semantics in a different respect. They offer a detailed account of the effects on Kaplan's (1977) semantics if it is modified in a relativist way, and in particular if what they call Kaplan's "Classical Reduction" (none other than principle (P) in § 3 above) is abandoned. They examine in detail how the interaction between indexicality and intensionality is affected by relativism.

John MacFarlane has been one of the pioneers of the recent wave of interest in relativism, as well as arguably its strongest proponent, having defended a relativistic view in a wider range of cases and in more detail than anyone else. His prize-winning 2003 paper "Future Contingents and Relative Truth" was the starting point of much further research, by others and by MacFarlane himself. Here he returns to the case of future contingents, revising the case he originally made. He shows how his own original argument against the main competitor to relativism in this area, supervaluationism, can be resisted. However, he also

⁵⁹ Bonnay and Egré, in this volume, undertake a detailed examination of the interaction of a possible world parameter and a perspectival parameter.

presents a new argument to the effect that supervaluationism, unlike relativism, cannot account for our use of the “actually” operator.

Denis Bonnay and Paul Egré propose a novel and specialized application of relativism. A margin for error principle for knowledge says that one knows that p only if in a situation that is indistinguishably different from one’s own, p is still the case. Epistemicists about vagueness use this principle to explain why we cannot have knowledge of borderline cases. A KK-principle says, roughly, that when one knows that p one also knows that one knows that p . These principles are well known to be incompatible because, roughly, each time we iterate knowledge more, the margin for error required for knowing becomes wider. So if knowledge implies knowledge that one knows, then the margin for error must be maximally wide. Thus no contingent knowledge would be possible, if both the margin for error principle and the KK-principle were true. Williamson consequently rejects the KK-principle. However Bonnay and Egré propose an ingenious way of making versions of the two principles compatible. Their proposal involves adding a new, perspectival parameter in the circumstances of evaluation, a parameter that captures what a subject of knowledge knows, thus making use of the relativist strategy.⁶⁰

Manuel García-Carpintero develops a differentiated view about relativism. After assessing what is at stake in discussions about relativism, he argues that radical versions of relativism are untenable. In particular, he argues against such a view in the case of gradable adjectives. On the other hand, he defends a version of moderate relativism for the same case.

Part II contains two essays exploring the metaphysical consequences of relativism. In the first, Crispin Wright distinguishes two versions of relativism. One version says that propositional truth is a ternary relation that relates a proposition with a world and a perspective of sorts. On this view, propositions can no longer be viewed as properly representational of the world. The other version says that propositional truth continues to be a binary relation relating propositions with worlds, however, on this view some worlds are accessible only from certain perspectives. He goes on to draw out a number of important consequences of the version that claims propositions to be non-representational.

Iris Einheuser also explores the metaphysical consequences of relativism in the cases of matters of taste as well as epistemic modals. Sticking for the moment to matters of taste, relativists will claim that propositions about matters of taste should be evaluated relative to a possible world *and* a standard of taste. Einheuser proposes a form of relativism—“fact-relativism”—that is akin to Wright’s second version, i.e. one according to which propositional truth is

⁶⁰ Bonnay and Egré model their idea by modifying the standard Hintikka epistemic semantics which involves an impersonal “It is known” epistemic operator, thus it is not straightforward to compare their project with that proposed by, for example, MacFarlane in his 2005b. However, I think it is safe to say that Bonnay and Egré’s perspectival parameter has a different function from MacFarlane’s and thus should count as a separate application of the relativist strategy.

binary, relating a proposition with a possible world. She shows that this requires a novel conception of possible world, according to which each possible world consists of an objective “substratum” and a standard of taste. This adjustment has the radical-sounding consequence that people with different standards of taste inhabit different worlds. Einheuser prefers such an account because it maintains the close link between what a propositional content represents and the facts. If we want to say that propositions have their truth-values perspectivally, we should therefore conclude that the facts represented by these propositions are likewise perspectival.

Part III contains three essays that mount detailed objections to relativism. Sebastiano Moruzzi tries out versions of the classic self-refutation objection against MacFarlane’s relativism as applied to matters of taste. He concludes that a fourth version of the objection provides a significant challenge because it shows that relativists cannot explain how a dispute on a matter of taste can be rational.

In his paper, Sven Rosenkranz objects to relativism about matters of taste primarily on the grounds that it does not respect what he regards as an unassailable Fregean truth, namely that to assert a proposition is to present it as true simpliciter, and similarly that one can derive the correctness of an assertion that p from p . He also points out that if assertion is not regarded as involving these norms of correctness, for example if to assert is merely to present as true on one’s own standard, then a case where one person asserts that p and another asserts that not- p is not a proper disagreement.

Richard Dietz offers a detailed criticism of relativism concerning epistemic modals. Relativists claim that when assessing a claim of epistemic possibility (e.g. “There might be a counterexample.”) one will assess it against the state of knowledge of the assessor at the time of the assessment. Thus, when I first claim that there might be a counterexample to Goldbach’s conjecture, and later find a proof of it, then at the later point I will say that my earlier claim was wrong. Dietz points out that this may be the case in the example just given, but that in cases where the assessor has less information than the one who made the claim assessed, this is not so: if I first claim that there cannot be a counterexample because I am in possession of a proof, and then later lose memory of, and access to, that proof, then it seems that I cannot correctly say that my earlier claim was wrong.

Finally Part IV comprises three papers that offer alternatives to relativism. Herman Cappelen approaches the issue from a particular point of view regarding the relationship between language use and semantics, namely speech-act pluralism, a view according to which the propositions asserted and communicated by an utterance are only quite loosely related with the proposition semantically expressed (“said”) by that utterance. He argues that the data pattern usually presented in support of a relativistic semantics does not support such a semantics given this pluralistic view about the relationship between semantic content and speech-act content. He thus underlines the point I made in § 2 (p. 7) above, that

a semantic theory alone is not testable against any data pattern without some pragmatic assumptions, and shows how making certain pragmatic assumptions can render semantic relativism unnecessary.

In a similar way, Andrea Iacona argues that an alternative way to make sense of alleged cases of faultless disagreement about matters of taste is to distinguish between subjective and objective “readings” of the uses of predicates of personal taste. Where the sense is subjective, there is no disagreement, where the sense is objective, there is no faultlessness. The claim is not that taste predicates are systematically ambiguous, but rather that while the semantic content of these predicates is stable, there is variation in what we communicate with them.

Finally, Dan López de Sa explores a different alternative to relativism, namely indexical relativism. According to indexical relativism about taste, taste predicates are implicitly indexical. Accordingly, when you say “Homer is funny.” and I say “He is not.”, we may be expressing propositions that do not contradict one another, for example on one proposal you express the proposition that you find Homer funny, while I express the proposition that I find Homer funny. Thus there is no problem in understanding why neither of us need be wrong. However, the difficulty for indexical relativists is to explain the appearance of disagreement. López de Sa shows how we can explain this on the basis of a presupposition that we have a common sense of humour, thus answering one of the major problems for indexical relativism.

References

- Bach, Kent (2006) “Perspective on Possibilities: Contextualism, Relativism or What?”. Paper presented at the Semantics & Pragmatics Workshop Sheffield, 2 December 2006.
- Borg, Emma (2004) *Minimal Semantics*, Oxford: Oxford University Press.
- Cappelen, Herman and Ernie Lepore (2004) *Insensitive Semantics*, Oxford: Blackwell.
- Carnap, Rudolf (1947) *Meaning and Necessity*, Chicago: Chicago University Press.
- Carston, Robyn (2002) *Thoughts and Utterances*, Oxford: Blackwell.
- Chalmers, David (2002) “On Sense and Intension”, *Philosophical Perspectives* 16: 135–82.
- (forthcoming) “Two-Dimensional Semantics”, in Lepore and Smith (eds.), *The Oxford Handbook to the Philosophy of Language*, Oxford: Oxford University Press.
- Cohen, Stewart (1986) “Knowledge and Context”, *Journal of Philosophy* 83: 574–83.
- DeRose, Keith (1991) “Epistemic possibilities”, *Philosophical Review* 100: 581–605.
- (1992) “Contextualism and Knowledge Attributions”, *Philosophy and Phenomenological Research* 52: 913–29.
- Dreier, James (1990) “Internalism and Speaker Relativism”, *Ethics* 101: 6–26.
- (1999) “Transforming Expressivism”, *Noûs* 33: 558–72.
- Egan, Andy (forthcoming) “Billboards, Bonds & Shotgun Weddings”, *Synthese*.

- Egan, A., J. Hawthorne and B. Weatherson (2005) "Epistemic Modals in Context", in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*, Oxford: Oxford University Press.
- Einheuser, Iris (this volume) "Three Forms of Truth-Relativism".
- Evans, Gareth (1985) "Does Tense Logic Rest on a Mistake?", in his *Collected Papers*, Oxford: Clarendon Press, 341–63.
- Field, Hartry (1973) "Theory Change and the Indeterminacy of Reference", *Journal of Philosophy* 70: 462–81.
- Fine, Kit (1975) "Vagueness, Truth and Logic", *Synthese* 30: 265–300.
- Frege, Gottlob (1892) "Über Sinn und Bedeutung", *Zeitschrift für Philosophie und philosophische Kritik* 100: 25–50. Reprinted in Günther Patzig (ed.), *Funktion, Begriff, Bedeutung*, Göttingen: Vandenhoeck and Ruprecht 1986. Eng. trans. in McGuinness 1984.
- (1918) "Der Gedanke. Eine Logische Untersuchung". *Beiträge zur Philosophie des deutschen Idealismus* 1, 2: 58–77. Reprinted in Günther Patzig (ed.), *Logische Untersuchungen*, Göttingen: Vandenhoeck and Ruprecht 1986. Eng. trans. in McGuinness 1984.
- [1915] "Meine grundlegenden logischen Einsichten". Note published posthumously in H. Hermes, F. Kambartel and F. Kaulbach (eds.), *Nachgelassene Schriften*, Hamburg: Felix Meiner (1983). Eng. trans. in *Posthumous Writings*, trans. P. Long and R. White, Oxford: Blackwell (1979), 251–2.
- Gibbard, Allan (1990) *Wise Choices, Apt Feelings*, Cambridge: Harvard University Press.
- Gillies, Anthony and Kai von Fintel (2006) "CIA leaks". Paper presented at Pacific APA meeting in Portland, 2006.
- Harman, Gilbert (1975) "Moral Relativism Defended", *Philosophical Review* 84: 3–22.
- Kaplan, David (1977/1989) "Demonstratives", in Almog *et al.* (eds.), *Themes from Kaplan*, Oxford: Clarendon Press (1989).
- Kölbel, Max (2002) *Truth Without Objectivity*, London: Routledge.
- (2003) "Faultless Disagreement", *Proceedings of the Aristotelian Society* 104: 53–73.
- (2004) "Indexical Relativism vs Genuine Relativism", *International Journal of Philosophical Studies* 12 (Oct. 2004), 297–313.
- (2007) "How to Spell Out Genuine Relativism and How to Defend Indexical Relativism", *International Journal of Philosophical Studies* 15: 281–88.
- (2008a) "The Evidence for Relativism", in *Synthese*.
- (2008b) "'True' as Ambiguous", *Philosophy and Phenomenological Research*.
- (forthcoming) "Vagueness as Semantic", in R. Dietz and S. Moruzzi (eds.), *Cuts and Clouds: Issues in the Philosophy of Vagueness*, Oxford: Oxford University Press.
- Kompa, Nicola (2002) "The Context Sensitivity of Knowledge Ascriptions", *Grazer Philosophische Studien* 64: 11–18.
- King, Geoffrey (2003) "Tense, Modality and Semantic Value", *Philosophical Perspectives* 17: 195–245.
- Lasersohn, Peter (2005) "Context Dependence, Disagreement, and Predicates of Personal Taste", *Linguistics and Philosophy* 28: 643–86.
- Lewis, David (1975) "Languages and Language", in *Minnesota Studies in the Philosophy of Language* 7: 3–35. Rep. in Lewis 1983.

- (1978) “Truth in Fiction”, *American Philosophical Quarterly* 15: 37–46. Rep. in Lewis 1983.
- (1979a) “Attitudes *De Dicto* and *De Se*”, *Philosophical Review* 88: 513–43. Rep. in Lewis 1983.
- (1979b) “Scorekeeping in a Language Game”, *Journal of Philosophical Logic* 8: 339–59. Rep. in Lewis 1983.
- (1980) “Index, Context, and Content”, in Stig Kanger and Sven Öhman (eds.), *Philosophy and Grammar*, Dordrecht: Reidel. Rep. in Lewis 1998.
- (1983) *Philosophical Papers*, vol. 1, Oxford: Oxford University Press.
- (1996) “Elusive Knowledge”, *Australasian Journal of Philosophy* 74: 549–67.
- (1998) *Papers in Philosophical Logic*, Cambridge: Cambridge University Press.
- McDowell, John (1980) “Meaning, Communication and Knowledge”, in Zak Van Straaten (ed.), *Philosophical Subjects: Essays presented to P. F. Strawson*, Oxford: Clarendon Press.
- MacFarlane, John (2003) “Future Contingents and Relative Truth”, *Philosophical Quarterly* 53: 321–36.
- (2005a) “Making Sense of Relative Truth”, *Proceedings of the Aristotelian Society* 105: 321–39.
- (2005b) “The Assessment Sensitivity of Knowledge Attributions”, forthcoming in Tamar Szabó Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology* 1, Oxford: Oxford University Press, 197–233.
- (2008) “Nonindexical Contextualism”, *Synthese*.
- (this volume) “Truth in the Garden of Forking Paths”.
- (forthcoming a) “Epistemic Modals are Assessment-Sensitive”, in Brian Weatherson and Andy Egan (eds.), *Epistemic Modals*, Oxford: Oxford University Press.
- (forthcoming b) “Semantic Minimalism and Non-indexical Contextualism”, in G. Preyer and G. Peter (eds.), *Context-Sensitivity and Semantic Minimalism: Essays on Semantics and Pragmatics*.
- McGuinness, Brian (1984) *Frege: Collected Papers on Mathematics, Logic and Philosophy*, Oxford: Basil Blackwell.
- Perry, John (1979) “The Problem of the Essential Indexical”, *Noûs* 13: 3–21.
- Predelli, Stefano (1997) “Talk about Fiction”, *Erkenntnis* 46: 69–77.
- (2005) *Context: Meaning, Truth and the Use of Language*, Oxford: Oxford University Press.
- (2006) “Operators on Contexts in the Semantics of Talk about Fiction”, paper presented at *Logic & Language 2006 Birmingham*.
- Price, Huw (1983) “Does ‘probably’ modify Sense?”, *Australasian Journal of Philosophy* 61: 396–408.
- Prior, Arthur (1967) *Past, Present and Future*, Oxford: Clarendon Press.
- Recanati, François (2004) *Literal Meaning*, Oxford: Clarendon Press.
- (this volume) “Moderate Relativism”.
- Richard, Mark (2004) “Contextualism and Relativism”, *Philosophical Studies* 119: 215–42.
- Rysiew, Patrick (2007) “Epistemic Contextualism”, *The Stanford Encyclopedia of Philosophy* (Fall 2007 edn.), Edward N. Zalta (ed.), forthcoming URL = <<http://plato.stanford.edu/archives/fall2007/entries/contextualism-epistemology/>>.

- Schlenker, Philippe (2003) "A Plea for Monsters", *Linguistics & Philosophy* 26: 29–120.
- Stalnaker, Robert (1987) *Inquiry*, Cambridge, Mass.: MIT Press.
- (2001) "On Considering a Possible World as Actual", *Proceeding of the Aristotelian Society* Supp. Vol. 75: 141–56.
- Stanley, Jason (2000) "Context and Logical Form", *Linguistics & Philosophy* 23: 391–434.
- Travis, Charles (1997) "Pragmatics", in Bob Hale and Crispin Wright (eds.), *A Companion to the Philosophy of Language*, Oxford: Blackwell.
- Wilson, Deirdre and Dan Sperber (2001) "Truthfulness and Relevance", *Mind* 111: 583–632.
- Wright, Crispin (2001) "On Being in a Quandary", *Mind* 110: 45–98. Rep. in C. Wright, *Saving the Differences*, Cambridge, Mass.: Harvard University Press (2003).
- (2006) "Realism, Relativism and Rhubarb", forthcoming in Greenough and Lynch (eds.), *Truth and Realism*, Oxford: Oxford University Press.
- Yourgrau, Palle (1983) "Knowledge and Relevant Alternatives", *Synthese* 55: 175–90.

PART I

RELATIVISM ELABORATED

This page intentionally left blank

2

Moderate Relativism

François Recanati

In modal logic, propositions are evaluated relative to possible worlds. A proposition may be true relative to a world w , and false relative to another world w' . A proposition whose truth-value varies across worlds is said to be *contingent* (as opposed to *necessary*). Relativism is the view that the relativization idea extends beyond possible worlds and modalities. Thus, in tense logic, propositions are evaluated relative to times. A proposition (e.g. the proposition that Socrates is sitting) may be true relative to a time t , and false relative to another time t' . A proposition that has this property is said to be *temporal* (as opposed to *eternal*). The view that there are such propositions may be called ‘Temporal Relativism’, or ‘Temporalism’ for short.

Further applications of the relativization idea easily come to mind. The proposition that it is raining (at a given time, in a given world) is true relative to some places, and false relative to others. The proposition that one is a philosopher is true relative to some persons, and false relative to others. The proposition that spinach is delicious is true relative to some standards of taste, and false relative to others. The proposition that the treasure might be under the palm tree is true relative to some epistemic situations, and false relative to others. The proposition that John is tall is true relative to some standards of height, and false relative to others.

In this paper I will discuss, and attempt to rebut, two classical objections to Relativism. Both objections are concerned specifically with Temporalism, but the issues they raise are quite general, as we shall see. Likewise, my responses are intended as a general defense of Relativism—not merely Temporalism.

The first objection, due to Frege, is the objection from incompleteness. I will distinguish two possible relativist responses to that objection, one of which corresponds to the view I actually defend: Moderate Relativism. Responding to that objection will therefore enable me to expound my view in some detail. The

I am indebted to Manuel García-Carpintero and John MacFarlane for comments on an earlier draft, and to the participants in the Barcelona workshop on ‘Relativizing Utterance Truth’ (especially Kit Fine) for remarks which inspired me. Thanks are also due to Julien Murzi for helping me update the Dummett reference, and to CSMN for financial support.

second objection is due to Mark Richard, who argued that the objects of belief cannot be relativistic (specifically, they cannot be ‘temporal propositions’). I will show that that objection can be met within the Moderate Relativist framework.¹ In the last section, I will deal with special forms of disagreement that have loomed large in recent discussions of Relativism.

2.1 THE OBJECTION FROM INCOMPLETENESS

2.1.1 Content and circumstance

As I understand it, the relativization idea has two component sub-ideas, which I will call ‘Duality’ and ‘Distribution’. Distribution presupposes Duality, but it is possible to accept Duality while rejecting Distribution.

[*Duality*] To get a truth-value, we need a circumstance of evaluation as well as a content to evaluate. (As Austin puts it, ‘It takes two to make a truth’.)

[*Distribution*] The determinants of truth-value distribute over the two basic components truth-evaluation involves: content and circumstance. That is, a determinant of truth-value, e.g. a time, is *either* given as an ingredient of content *or* as an aspect of the circumstance of evaluation.

The distribution idea is apparent in the literature that stems from John Perry’s work on unarticulated constituents. According to Perry (1986), if something is given as part of the situation which an utterance (or, for that matter, a mental representation) concerns, and against which it is evaluated, it does not have to be articulated in that representation. Thus Perry draws a distinction between ‘It’s raining here’, which explicitly mentions a place, and ‘It’s raining’, which leaves the place out of the picture. Perry describes the content of the latter not as a complete proposition but as a *propositional function*, true of some places and false of others. The place which actually determines the truth-value of the utterance is fixed not by the content of the utterance but by the situation which that utterance concerns (the situation the speaker manifestly intends to characterize). When I say ‘It’s raining here’ the situation which my utterance concerns is typically more complex since it involves several places at once, between which a contrast is drawn (Recanati 1997, 2000). In this case the place must be articulated and cannot be left out of the picture, since it is not independently fixed by the situation talked about.

The distribution idea also comes up in Kaplan’s well-known argument for temporal propositions (Kaplan 1989: 502–4). Kaplan’s argument is based on the existence of temporal operators. The contents temporal operators operate

¹ For reasons of space, I cannot discuss the objections to Temporalism raised by King (2003), as I had planned. I do so elsewhere.

on must be temporally neutral, Kaplan argued, for if they are not—if they are temporally-specific—the temporal operators will be vacuous. A temporal operator specifies the time(s) with respect to which the proposition it operates on is to be evaluated. If the proposition itself specified a time, embedding the proposition under the temporal operator would have no effect whatsoever. Being already specified by the content to be evaluated, the time of evaluation would be fixed once for all and could no longer be shifted. So temporal operators must operate on temporally-neutral propositions—propositions which are true with respect to a time, and false with respect to another time, but which do not specify the time relative to which they are supposed to be evaluated.

The general principle which emerges is a principle of economy or optimality according to which a determinant of truth-value is either given as an ingredient of content or as an aspect of the circumstance of evaluation, but not both. The richer the circumstance, the poorer the content evaluated with respect to that circumstance; and the richer the content, the poorer the circumstance. In particular:

- If the circumstance consists of a possible world only, the content must be a complete proposition (something that determines a function from possible worlds to truth-values).
- If the circumstance is richer and involves a time and a place in addition to a world, then the content can be less than fully propositional: it can be place- and time-neutral and determine only a propositional function (a function from place-time pairs to functions from possible worlds to truth-values, or equivalently, a function from centered worlds to truth-values).

2.1.2 The objection

Frege rejected the very idea of a temporal proposition, i.e. a proposition that is true at some times and false at other times. Such a proposition is not a genuine proposition, he held, because it is not evaluable as true or false, or at least, it is not evaluable *unless* we are given a particular time. In the absence of a time specification, the alleged proposition is only ‘true-at’ certain times and ‘false-at’ others. It is, therefore, *semantically incomplete* by Frege’s lights:

A thought is not true at one time and false at another, but it is either true or false, *tertium non datur*. The false appearance that a thought can be true at one time and false at another arises from an incomplete expression. A complete proposition or expression of a thought must also contain a time datum. (Frege 1967: 338, quoted in Evans 1985: 350)

As Evans points out, the problem of semantic incompleteness does not arise in the modal case. Even if a thought is said to be ‘true-at’ one world and ‘false-at’ another, as in modal logic, this does not prevent it from being true (or false) *tout court*. It is true *tout court* iff it is true-at the actual world. But the ‘thought’

that it is hot cannot be evaluated as true or false *tout court*. In the absence of a contextually supplied time it can *only* be ascribed relative, 'truth-at'-conditions. Only a particular, dated utterance of such a sentence can be endowed with genuine truth-conditions. What this shows is that the time of utterance is part of the (complete) content of the utterance, or, in a Fregean framework, part of the expression of such a content; hence it cannot be deemed external to content and treated like the world of evaluation. So the objection goes. And the same objection applies to the place-neutral content of 'It's raining': such a content is not complete, since the utterance cannot be evaluated unless a place is contextually provided.

There are two possible responses which a Relativist can make to that objection.² A Radical Relativist will insist that the nonclassical contents that we are led to encompass if we accept Distribution *are* complete. Thus the Stoics posited '*lekta*' that were "in many respects reminiscent of the 'propositions' that many modern philosophers postulate as meanings of eternal assertoric sentences", save for the fact that they were "temporally indefinite in the same way as occasion sentences" (Hintikka 1973: 70). Such *lekta* were thought by them to be complete, despite their temporal neutrality. In his review of Mates 1953, which brought Stoic logic (and temporal propositions) to the forefront of attention, Geach wrote that for the Stoics, "though the truth-value of 'Dion is alive' changes at Dion's death, the sentence still expresses the same complete meaning (*lekton*)" (Geach 1955: 144). This idea, which aroused Prior's interest, Evans later found incomprehensible and even incoherent (Prior 1967: 17; Evans 1985: 348–50). If the *lekton* is complete, Evans argued, it can be evaluated as correct or incorrect; but if the *lekton* is temporally neutral, its evaluation as correct or incorrect will vary with time, hence it will not be evaluated as correct or incorrect once for all. Does this not entail that the *lekton cannot* be evaluated as correct or incorrect (*tout court*) after all? Evans writes:

To say that the sentence type 'Socrates is sitting' . . . expresses a complete meaning seems to imply that . . . to know what assertion is being made by an utterance of a tensed sentence all you need to know is which tensed sentence was uttered; you do not need further information to tie the sentence down to a particular time. . . . It would follow that such an 'assertion' would not admit of a stable evaluation as correct or incorrect; if we are to speak of correctness or incorrectness at all, we must say that the assertion is correct at some times and not at others. (Evans 1985: 349)

For Evans, this consequence (the instability of evaluation) is a *reductio* of the whole position. At this point, however, it is useful to consider MacFarlane's discussion of future contingents, for MacFarlane seems to bite Evans's bullet. MacFarlane (2003) argues that a sentence like 'There will be a sea-battle tomorrow' is neither true nor false when it is uttered (since the future is indeterminate)

² See García-Carpintero (this volume) for a similar distinction between two versions of Relativism.

but turns out to be true or false, as the case may be, when it is evaluated the next day. So MacFarlane gives up the constraint that the evaluation of a thought as correct or incorrect must be temporally stable: the truth-value of an utterance may well depend upon the context of evaluation (e.g. the time at which it is evaluated), so an utterance or thought that is evaluated in a certain way at a certain time may be evaluated differently at a different time. In such a framework, reminiscent of Aristotle, we could maintain that the tensed sentence 'Dion is alive' expresses a complete content, and is (therefore) evaluable (at any given time), since we reject the constraint that the evaluation process itself must be 'eternal', hence stable, rather than context-sensitive and unstable.³

Whatever we think of this line of argument, I will not be concerned with the radical forms of Relativism in this paper, but only with a moderate form which I myself advocate. In response to the Fregean objection, a Moderate Relativist will *concede* that the complete content of the utterance/thought 'Dion is alive' involves more than the temporally neutral *lekton* it expresses; it additionally involves the time of utterance, which is tacitly referred to and against which the utterance is meant to be evaluated. Distribution can be construed as saying that the complete content, *in the sense of Frege and Evans*, distributes over the two components which Duality posits, namely the circumstance of evaluation (which may include more than a world) and the content to be evaluated, in the narrow sense of content. Once it is admitted that we need these two components, we can tolerate contents that are not 'semantically complete' in Frege's sense, i.e. endowed with absolute truth-conditions. We can, because the circumstance is there which enables the content to be suitably completed. Thus the content of tensed sentences is semantically incomplete, yet the circumstance (the time) relative to which such a sentence is evaluated is sufficient to complete it. It follows that we must distinguish *two* levels of content. The content we evaluate with respect to the circumstance is the content in the narrow sense; it may, but need not, be semantically complete by Frege's lights. What is semantically complete in any case is the content in the broad sense. It consists of the (narrow) content *and* the circumstance with respect to which that content is meant to be evaluated. Distribution only induces us to *analyse* the complete content of an utterance into two components, corresponding to those distinguished in Duality.

³ Note, however, that this is not what MacFarlane himself would say. MacFarlane's brand of Radical Relativism consists in making room for a new form of context-sensitivity: sensitivity to the context of evaluation and not (or not merely) to the context of utterance. On MacFarlane's view, some expressions are 'assessment sensitive', and others are not—just as some expressions are utterance sensitive, and others are not. In this regard, future contingents are a *special* case. (Evans himself seems to accept that there is something special about future contingents, and for that reason, he says, he confines his discussion to sentences in the past. See Evans 1985: 350, fn. 9.) As far as 'Dion is alive' is concerned, MacFarlane holds that its truth depends on the time of utterance, as Frege points out, but *not* on the time of assessment. This shows that one may be a Radical Relativist with respect to some sentences—those whose truth-value is assumed to depend upon the context of assessment and whose evaluation is therefore unstable—and not with respect to others.

2.1.3 Two levels of content

The position I have sketched has been argued for by several authors, more or less explicitly. Thus Hintikka, in an interesting article on 'Time, Truth and Knowledge in Aristotle and Other Greek Philosophers', says the following:

It is obvious that the sentence, 'It is raining', as uttered by me today, is made true or false by a set of facts different from those that verified or falsified my utterance yesterday, 'It is raining'. But it is very natural to say that in some sense the state of mind or attitude toward my environment that is expressed by the two utterances is the same. The facts to which yesterday's utterance refers are referred to today by the sentence, 'It was raining yesterday'. But the 'state of mind' that this utterance appears to express seems to be entirely different from that expressed by yesterday's present-tense utterance, 'It is raining'. (. . .) Hence the idea that spoken words are symbols for unspoken thoughts encourages the idea that one and the same temporally indefinite form of words expresses one and the same belief or opinion at the different times when it is uttered. (Hintikka 1973: 85)

To me at least, this suggests that the complete content of an utterance (that which determines its truth-conditions) involves two factors: the thought that is expressed, and the time at which it is expressed. The sentence 'It is raining' expresses the same thought whenever it is uttered, and that thought is evaluated with respect to the time of utterance. Since the latter changes, the truth-value is liable to change even though, in the narrow sense of content, the content is the same. The truth-conditions also change: An utterance of 'It's raining' at t is true iff the thought expressed by the sentence is true at t ; an utterance of the same sentence at t' is true iff the same thought is true at t' . On this view the complete content of two successive utterances of 'It is raining' need not be the same, since the speaker does not merely express a certain content, but also tacitly refers to a certain time (the time of utterance) as relevant for the evaluation of that content. The complete content corresponds to the utterance's truth-conditions which, according to Hintikka, depend upon an external factor, namely the actual time at which the utterance is made or the thought entertained.

In a similar vein, Dummett attempts to make sense of Prior's position (in response to Evans's critique) by distinguishing two levels. He points out that temporal propositions are, for Prior, the contents of *sentence-types*. The content of a sentence-type is a function from times to truth-values, hence a sentence-type only has relative truth-conditions: it is true at some times and false at other times. This does not prevent us from introducing a notion of absolute truth, by shifting to the level of *utterance* content. According to Dummett, when a sentence is uttered the function which is its content is applied to some contextually provided time (typically, the time of utterance). The time in question serves as circumstance of evaluation for the utterance: the utterance is true *tout court* iff the sentence is 'true-at' the contextually provided time. As Dummett emphasizes,

The variable truth-value and the absolute truth-value attach to different things; it is the type sentence that is true at one time, false at another, but the utterance that is true or false simpliciter. (Dummett 1997 55n.)

Since there are two distinct levels, corresponding to the sentence-type and the utterance, there is no harm in taking the utterance to possess a ‘content’ also (content_u), distinct from that of the sentence (content_s). For example, we can treat the utterance as expressing a structured proposition consisting of (i) the contextually provided time as subject, and (ii) the content of the sentence-type, predicated of that time. But if we do so, we must acknowledge the unarticulated nature of the ‘subject’ in the content_u of tensed utterances. As Prior says, “tensed propositions are understood as directly or indirectly characterising the *unmentioned* time of utterance” (Prior 1977: 30). Hence there is a trade-off: if we want to restrict ourselves to what is linguistically articulated, we must focus on the content_s, which is ‘semantically incomplete’ by Frege’s lights—it corresponds to the content of a predicate rather than to that of a complete sentence in a logically perfect language. If, following Frege, we want to focus on the complete content of the utterance, that which makes it truth-evaluable in absolute terms, we must acknowledge the role played in that content (content_u) by unarticulated constituents corresponding to the circumstances in which the content_s is evaluated.

Another author who ought to be mentioned in connection with Moderate Relativism is Jon Barwise. Barwise also put forward a semantic theory with two levels of content: the ‘infor’ or ‘state of affairs’ and the ‘Austinian proposition’ (Barwise 1989, Barwise and Etchemendy 1987). The infor is the content to be evaluated with respect to a given situation, and the Austinian proposition is the proposition to the effect that that situation supports that infor. In what follows I will use the notion of Austinian proposition, corresponding to the complete content of an utterance/thought. But I will use the Stoic term ‘*lekton*’, rather than Barwise’s theoretically-loaded term ‘infor’, to refer to the content in the narrow sense.⁴ So ‘It is raining’ expresses a constant *lekton* whenever and wherever it is used, a content that can be modelled as a function from situations to truth-values or as a set of situations (viz. the set {*s*: it is raining in *s*}); but the complete content of an utterance of ‘It is raining’ is the Austinian proposition that a certain situation (that which the utterance/thought ‘concerns’) fits that *lekton*, i.e., belongs to the set of situations in question.

In Dummett’s framework the partial content was the content of the sentence-type. Now the infor, according to Barwise, is the content of the sentence *with respect to context*: if the sentence contains indexicals, the contextual values of the indexicals contribute to the infor. I will retain that feature of Barwise’s account. In my framework, the *lekton* is the content of the sentence in context,

⁴ Evans also has coined a term for that entity. He calls it the ‘Stoic-proposition’ (Evans 1985: 350).

so an indexical sentence will express different *lekta* in different contexts. But the context comes into the picture a second time: it not only provides values for the indexicals, which values contribute to the *lekton*, but it also determines the situation against which the *lekton* is to be evaluated. The complete content of the utterance involves the *lekton* together with the situation of evaluation.⁵

2.1.4 Moderate Relativism: two versions

Following various authors, I have suggested that we need the *lekton* as a level of content even though it is not the complete content (that which determines the utterance's possible-worlds truth-conditions). The *lekton* is the content of the sentence (with respect to context, if the sentence is indexical), but the complete content of the utterance involves something more: it involves a situation with respect to which the utterance is meant to be evaluated. Change the situation of evaluation, you change the complete content of the utterance, even though the content of the sentence (with respect to context) remains constant.

The debate between classical theorists and Moderate Relativists bears upon the indispensability of the *lekton* as a level of content. According to the classical theorist, the only thing we need is the complete, truth-conditional content on the one hand and the meaning of the sentence-type on the other. One reason for positing an extra level of content, viz. the *lekton*, is that it enables us to represent what the sentence (or possibly the thought) explicitly articulates (in a possibly indexical manner). Again, 'It's raining here' says something different from what 'It's raining' says, even in a context in which they are both true iff it is raining at the place of utterance. The difference lies in the fact that the place in question is (indexically) articulated in the former case while it is left unarticulated in the latter. It follows that the *lekton* differs, even though the truth-conditions are the same.

There is another debate, concerning the special case in which what the sentence explicitly articulates is a classical proposition. Suppose the speaker says 'It is raining here and now', without leaving anything unarticulated (save the world of evaluation). Both the time and the place are explicitly articulated, hence they are both part of the *lekton*. It follows that we don't need a rich circumstance to evaluate that content: the *lekton* already determines a function from possible-worlds to truth-values, hence the only thing we need to determine a truth-value is a possible world. No further relativization is needed. So it seems that, *with sentences whose content is not semantically incomplete*, there is no need to invoke a double layer of content. The content of the sentence-in-context,

⁵ Moderate Relativism, thus understood, is what MacFarlane calls 'Non-indexical Contextualism' (MacFarlane forthcoming). It also corresponds to Kaplan's own position, since Kaplanian 'contents' do not determine a classical proposition unless a circumstance (involving a time and, possibly, a place in addition to a world) is contextually provided.

insofar as it has an absolute truth-value, is the only thing we need, as in the classical theory. Or, to put it in slightly different terms: in such cases the *lekton* is the complete content. This position defines one version of Moderate Relativism, namely the weak version ('WMR', for 'weak moderate relativism'). But there is another, strong version, which has been argued for by Barwise and which I also advocate.

On the strong version ('SMR'), the content of a sentence (*whatever* the sentence) is a function from situations to truth-values. Hence the relativity of truth, construed as a property of sentences: the same sentence may be true relative to a situation and false relative to another one. That is so *even if the sentence itself is not semantically incomplete*. Even when the sentence is truth-evaluable in the absolute sense—when it is 'semantically complete' by Frege's lights—SMR says there is a principled distinction between the content of the sentence (the *lekton*) and the content of the utterance (the Austinian proposition). In such a case, the *lekton* will be a 'classical' proposition (a function from possible worlds to truth-values), but the Austinian proposition will still contain a situation in addition to that proposition. What the utterance 'says' is that *the situation in question supports the proposition in question*. It follows that two distinct evaluations are possible, in such cases. We can evaluate the sentence itself (i.e. evaluate the proposition with respect to the actual world), or we can evaluate the utterance, that is, evaluate the proposition *with respect to the situation figuring in the Austinian proposition*.

To illustrate this point I usually quote my favourite example, from Barwise and Etchemendy 1987. Commenting upon a poker game I am watching, I say: 'Claire has a good hand now'. What I say is true, iff Claire has a good hand in the poker game I am watching at the moment of utterance. But suppose I made a mistake and Claire is not among the players in that game. Suppose further that, by coincidence, she happens to be playing bridge in some other part of town and has a good hand there. Still, my utterance is not intuitively true, because the situation it concerns (the poker game I am watching) is not one in which Claire has a good hand at the time of utterance. But we can say that the *sentence* (or the sentence-in-context) is true: for it says that Claire has a good hand at the time of utterance, and Claire *has* a good hand (somewhere) at the time of utterance. The unarticulated constituent which distinguishes the *lekton* from the Austinian proposition makes all the difference here, and it accounts for our intuitive classification of the utterance as non-true.

This sort of approach can easily be extended to deal with standard problems such as that of quantifier domain restriction. It is natural to hold that 'All *F*s are *G*' expresses a proposition that is true (in a world, at a time) if and only if all the *F*s are *G* (in that world, at that time). Thus 'All students are French' expresses the proposition that all students are French. Many theorists feel compelled to give up this natural view, and claim that the sentence is semantically incomplete or

covertly indexical, so that it expresses no proposition (independent of context).⁶ They say so because they are impressed by the fact that the truth-conditions of an utterance of that sentence typically involve a contextually restricted domain of quantification. In the SMR framework, however, we can stick to the simple and straightforward view regarding the proposition expressed by ‘All *F*s are *G*’, while fully acknowledging contextual domain restriction. The two layers of content enable us to do just that. The sentence is said to express a proposition that is evaluable with respect to an arbitrary world (or, perhaps, an arbitrary world-time pair)—the proposition that all students are French—but that proposition can *also* be evaluated with respect to the specific situation that features in the Austinian proposition. That is what happens when we evaluate an utterance of this sentence, instead of evaluating the sentence itself.⁷

2.2 THE OBJECTION FROM BELIEF REPORTS

2.2.1 Richard 1981

In ‘Temporalism and Eternalism’, Mark Richard put forward what many take to be a knock-down argument against Temporalism (the view that there are temporal propositions). Since Temporalism is a particular form of Relativism, we must consider his argument to see whether or not it threatens SMR.

According to Richard’s argument, “the temporalist is unable to give an adequate treatment of attributions of belief” (Richard 1981: 3). Richard asks us to consider the following piece of reasoning:

- [1] Mary believed that Nixon was president
- [2] Mary still believes everything she once believed

Ergo

- [3] Mary believes that Nixon is president

As Richard points out, “this argument is not a valid argument in English” and “we ought to reject any position which is committed to [its] validity” (Richard 1981: 4). Temporalism, Richard claims, is one such position. For the temporalist holds that ‘Nixon is president’ expresses a temporal proposition p_1 , true at any time t iff Nixon is president at t . Let us assume, plausibly enough, that a belief report ‘ x believes that S ’ states that the individual referred to by the subject term is belief-related to the proposition expressed by the embedded

⁶ See e.g. Stanley and Szabó 2000.

⁷ A well-known difficulty for the situation-theoretic approach to contextual domain restriction comes from the fact that distinct quantifiers in a single sentence may involve distinct restrictions. The answer to that difficulty consists in associating sub-sentential expressions with (local) circumstances of evaluation. See e.g. Recanati 1996.

sentence. It follows that ‘Mary believes that Nixon is president’ expresses the proposition that Mary believes p_1 . On the equally plausible assumption that the past tense in the embedded clause of ‘Mary believed that Nixon was president’ is semantically vacuous, it follows that [1], ‘Mary believed that Nixon was president’, expresses the proposition that Mary believed p_1 . Now this, together with [2] (the proposition that Mary still believes everything she once believed), entails that she still believes p_1 , i.e., that she still believes that Nixon is president! Since that conclusion does not actually follow, there is something wrong with Temporalism.

Richard takes his argument to show that “the objects of belief expressed by sentences are all eternal” (Richard 1981: 10), i.e., they are classical propositions, not temporal propositions. Temporalism can be rescued, Richard points out, if we give up the assumption that “a sentence expresses at most one thing (a proposition) at a time” (1981: 9). Moderate Relativism as I have described it precisely rejects that claim, since it posits two levels of content for every utterance. I will return to Moderate Relativism shortly. Richard himself describes a view which he calls ‘Moderate Temporalism’, which rejects the ‘single content’ assumption:

We distinguish two different relations of expressing (say, *expresses*₁ and *expresses*₂) and two distinct classes of objects, which we may call contents and propositions. Expression₁ is a relation between sentences and contents; expression₂ is a relation between sentences and propositions. Contents may be either eternal or temporal; propositions are all eternal.

We now take contents to be the bearers of truth and falsity expressed by sentences, propositions to be the objects of belief so expressed. A sentence S is true, relative to time t , iff there is a content c such that S expresses₁ c at t and c is true at t . A sentence S expresses, relative to t , a belief of a person u iff there is a proposition p such that S expresses₂ p at t and u believes p at t . (Richard 1981: 10)

In terms of temporally neutral content, we can make sense of the claim that, in a certain sense, two persons who say that it is raining (at different times) ‘say the same thing’: that it is raining. Their respective utterances are true iff and only if that constant *lekton* is true at the times of their respective utterances. But what Richard’s argument about belief reports is supposed to establish is that the content of *belief* is not such a temporally neutral *lekton*: the content of belief is a classical (eternal) proposition. For a Moderate Temporalist of the sort Richard describes, what is said is a temporal proposition, but what is believed, or what the utterance presents the speaker as believing, is a classical proposition.

Richard does not find Moderate Temporalism particularly attractive, because there are utterances like ‘What you say is true and I believe it, too’ which show that the object of assertion is, or at least can be, the same as the object of belief. At this point, Richard argues, the Moderate Temporalist will have to distinguish the object of assertion thus understood (a classical proposition, like the object of belief) from ‘what the speaker says’ in the temporally neutral sense (the *lekton*). But that temporalist notion of ‘what the speaker says’, distinct both from what

the speaker asserts and from what she believes, becomes suspicious, and it is unclear that we need it. “Until some clarification of this notion of ‘what is said’ by an utterance is given”, Richard concludes, “we should remain sceptical” (Richard 1981: 12).

2.2.2 Richard 2003

I have spelled out Richard’s argument in some detail, in order to make clear where the Moderate Relativist differs from the hypothetical Moderate Temporalist described by Richard. Like Richard’s Moderate Temporalist, the Moderate Relativist distinguishes two types of content and two relations of expressing; but he would deny that one type of content is what is said, and the other what is believed. The distinction between the two types of content cuts across the distinction between saying and believing. That means that, whether we consider the speaker’s assertion or the speaker’s belief, we can distinguish two things: the *lekton* (content in the narrow sense) and the complete content or Austinian proposition. Richard himself comes close to that conclusion when, on behalf of the Temporal Relativist, he draws a tentative distinction between what is asserted (a classical proposition) and what is ‘said’ (a temporal proposition). In a later paper, he gives example like

- (1) When Susan saw Kate two winters ago, she swore that Kate was pregnant, and when Mindy saw her this spring, that’s what she said too.

and he comments as follows:

- (1) seems to report Susan and Mindy as literally saying the same thing; if they do, presumably they each say something temporally neuter. But . . . suppose that last spring Susan saw Kate and said to herself, ‘(I guess that) she wasn’t pregnant two winters ago, but she is now’. Then we can surely go to Kate and say,
- (2) When Susan saw you two winters ago, she said that you were pregnant, but now she takes that back/denies that/denies what she said.

All this, it might be said, suggests that when someone utters a tensed, but temporally unspecific, sentence, two distinct reports of what she said will be possible: one reporting her as having said something temporally specific, and one reporting her as having said something temporally unspecific. And this suggests that utterances of temporally unspecific sentences express, or at least typically express, two things, one temporally unspecific, the other specific. (Richard 2003: 39–40)

That is exactly what a Moderate Relativist will say; and the Moderate Relativist will point out, as Richard himself does in the later paper, that the same point can be made with respect to belief. Richard gives the following example of a belief ascription where the object of belief seems to be a temporal proposition:

- (3) Bob went to the monkey house, and now he thinks that he’s been infected with the Ebola virus. Every time he goes there he thinks that; he’s convinced one of the monkeys is a carrier.

The word ‘that’ in ‘Every time he goes there he thinks that’ seems to refer to a temporal proposition (since the eternal proposition believed by Bob after visiting the monkey house changes from one visit to the next). Richard, however, thinks the evidence is misleading. He has a story to tell regarding examples like (3), a story which does not appeal to temporal propositions as objects of belief.⁸ His reason for resisting the view that there are two possible objects of belief, corresponding to the two levels of content distinguished by the Moderate Relativist, is that “diachronic agreement or disagreement seems to be, of necessity, a matter of agreement or disagreement about something temporally specific” (Richard 2003: 40). So we are back to Richard’s original argument: whether one changes one’s mind or retains one’s belief is a matter of still believing (or ceasing to believe) the same classical propositions. *When it comes to assessing inter- or intra-individual (dis-)agreement, only classical contents count.* So if we know that Mary retained all of her previous beliefs, we will not conclude that she still believes that Nixon is president even though we know that, twenty years ago, she believed that Nixon was president. On this issue, Richard has not changed his mind from 1981 to 2003, and his objection to temporal propositions as the objects of belief still stands.

2.2.3 Reply to Richard

In response, the Moderate Relativist can point out that belief reports have (at least) two distinct functions. First, one may report someone’s beliefs in order to assess their (dis)agreement with either (i) the facts, or (ii) the beliefs of other people, or the beliefs held by the same person at different times, about the same facts. Such belief reports will typically focus on the truth-conditional properties of the belief, hence on its complete content. Second, one may report someone’s beliefs in order to link those beliefs to other states or acts of the same person, for example her sensory experiences, her actions, or other beliefs potentially or necessarily held by her. Such belief reports focus not on the truth-conditional properties of the belief but on what McGinn calls its ‘intra-individual causal-explanatory role’. This distinction between two functions of belief reports is well known and it has been extensively documented in the late seventies. In his classical paper on these issues, McGinn writes:

Our concept of belief combines two separate elements, serving separate concerns: we view beliefs as causally explanatory states of the head whose semantic properties are, from that point of view, as may be; and we view beliefs as relations to propositions that

⁸ According to Richard (2003: 41–2), in ‘Every time he goes there he thinks that’, ‘that’ does not refer to a temporal proposition denoted by the antecedent ‘that’-clause ‘that he has been infected with the Ebola virus’. Rather than construe ‘that’ as a device of cross-reference, we may, “with a fair amount of plausibility”, construe it as a device of ellipsis, Richard says (p. 42). On that analysis ‘Every time he goes there he thinks that’ is *short for* ‘Every time he goes there he thinks that he’s been infected with the Ebola virus’, and in *that* sentence the belief that is ascribed to Bob is temporally specific (eternal) rather than temporally neutral.

can be assigned referential truth-conditions, and so point outward to the world. This bifurcation of content can be seen as stemming from the point that beliefs involve internal representations, and these inherently present a dual aspect. (McGinn 1982: 216)

Since there is this duality in our notion of belief, it is not surprising that there is an ambiguity in a belief report like 'Susan believes that Kate is pregnant'. This may ascribe to Susan either the internal state of believing Kate pregnant, a state one may be in at different times ('relativist' interpretation); or it may ascribe to her a belief with a certain truth-conditional content, which content depends, as we have seen, upon external factors such as the time at which the belief is held ('classical' interpretation). On the latter interpretation, Susan's belief can change from one occurrence of the internal state to the next, even though the internal state itself does not change. At t , Susan is in the state of believing Kate pregnant, and she thereby believes the classical proposition that Kate is pregnant at t ; at t' Susan is in the same state, but the classical proposition she now believes is the (distinct) proposition that Kate is pregnant at t' . If, on the classical interpretation, we say that someone's beliefs have not changed, then it follows that she believes all the *classical* propositions she formerly believed; but it does not follow that her internal doxastic state has not changed. On that interpretation the argument Richard presents as invalid is indeed invalid. From the fact that, at a certain time t , Mary was in the state of believing Nixon president, and thereby believed the classical proposition that Nixon is president at t , plus the fact that she still believes all the classical propositions she once believed, it does not follow that she still is in the state of believing Nixon president and thereby believes the classical proposition that Nixon is president now.

So, on the classical interpretation of a belief report, the object of belief is indeed the complete content (which we can represent either as a classical proposition, or as an Austinian proposition). But that is not the only possible reading of a belief report. There is another reading, where the ascriber is interested in the intra-individual causal-explanatory role of the ascribed belief. In such cases what matters is the *lekton*, not the complete truth-conditional content.

That we need the *lekton* in such cases has been forcefully argued by Barwise, who gives the following example. Suppose Holmes and Watson face each other. In between stand the salt and the pepper. Holmes says 'The salt is left of the pepper', because the salt is left of the pepper from Holmes's perspective. From Watson's perspective, the pepper is left of the salt; however, Watson is mistaken as to which shaker is which, and he wrongly says 'The salt is left of the pepper'. Holmes and Watson apparently 'say the same thing' (so they express the same *lekton*) but Holmes is right and Watson wrong (so they believe different classical propositions, or different Austinian propositions, because they each relativize the *lekton* to their own perspective). In the classical framework, Barwise points out,

... we have nothing in the theory that classifies the similarity in attitudes of Holmes and Watson in cases like these. And it is this similarity that leads them to make the same

bodily movements, reaching in the same direction, though toward different objects, when they want the salt. (Barwise 1989: 240).⁹

In other words, if what we are interested in is the state Holmes and Watson are both in, and the causal-explanatory role of that state, then we should accept that (in the relevant sense) they believe the same thing: they both have a belief with a certain *lekton* as content, which *lekton* determines different truth-conditions when evaluated with respect to their distinct perspectives. So there is a sense in which Holmes and Watson believe the same thing in that situation, and there is also a sense in which they do not believe the same thing. When arguing that the content of belief must be eternal (classical), Richard simply focuses upon the sense which is relevant to belief reports whose function is to assess (dis)agreement with the facts; but everybody knows that that is not the sole function of belief reports.

If I am right, shouldn't there be an interpretation in which the argument Richard discusses *is* valid? Let us reconsider that argument:

- [1] Mary believed that Nixon was president
- [2] Mary still believes everything she once believed
- [3] *Ergo*: Mary believes that Nixon is president

Since [3] does not follow, Richard argues that [1] cannot be interpreted as saying that Mary stood in a certain relation to the temporally indefinite *lekton* <Nixon, being president>; for if it could be so interpreted, [3] would follow. Now I hold that [1] *can* be interpreted in this 'relativist' manner. On that interpretation, indeed, [3] ought to follow—but it does not. Does this not show that Richard is right, and the Moderate Relativist wrong?

I do not think so. The reason why [3] is hard to accept in this context may be due to the fact that the universal quantification over beliefs in premiss [2] has to be understood either as quantifying over complete contents or as quantifying over *lekta*, since those are different sorts of things. Let's assume that a choice has indeed to be made from the outset when interpreting [2]. Then, arguably, the 'classical proposition' interpretation of [2] is more salient, and that is why [3] does not follow from [1] and [2] even though [1] *could* be interpreted as saying that a certain relation obtains between the subject and a temporally indefinite *lekton*. To check that this is correct, we have only to rephrase [2] so as to make the *lekton* interpretation more salient.

To show this, let us go back to the Susan/Kate example and run the Richard argument in that context, while suitably modifying (the counterpart of) premiss [2]. We get:

- [1'] Susan believed that Kate was pregnant

⁹ A classical theorist could respond that the linguistic meaning of the sentence type (Kaplan's 'character') can be invoked to account for the relevant similarities. See Perry 1979 and Richard 1982. For reasons of space, I have to leave that issue aside.

[2'] Susan is exactly in the same doxastic state she was in. (Imagine a Rip Van Winkle context in which, unbeknown to her, she has just awoken from a two-year sleep.)

[3'] *Ergo*: she still believes that Kate is pregnant.

This argument seems valid. Now the only substantial thing I have done (besides transposing the example) is change the second premiss in order to force the causal-explanatory reading of the first premiss. That is precisely the reading which Richard wrongly treats as nonexistent. So my example shows that, *pace* Richard, *there is* a reading in which the embedded clause in a belief report like [1] or [1'] is used to pick out the *lekton* rather than the complete content of the ascribed belief.

2.3 RELATIVISTIC DISAGREEMENT

2.3.1 Shared circumstances

In responding to Richard's argument, I have conceded that when the point of a belief report is to assess (dis)agreement with the facts or with other people or with oneself at different times, it is the complete content of the belief that matters, since the complete content is what determines possible-worlds truth-conditions. But this too can be disputed.

There are two types of example which cast doubt on the idea that agreement or disagreement can only be about complete contents. The first type of example involves cases in which all the parties to the conversation are in the same situation—for example, they share their location. In such cases, there is no objection to one of them asserting a place-neutral proposition such as 'Sydney is nearby', true at any place *l* iff Sydney is near *l*. (I borrow the example, and the argument, from Egan 2007.) Let us assume the belief-transfer model of assertion, according to which the function of assertion is to transfer belief in the asserted proposition (Stalnaker 1978). In the present case the transfer can be described as follows:

Speaker A, at place *l*, accepts the place-neutral proposition that Sydney is nearby, and thereby believes the classical proposition that Sydney is near *l*.

Speaker A asserts the place-neutral proposition that Sydney is nearby.

As a result, audience B, also at place *l*, comes to accept the place-neutral proposition that Sydney is nearby, and thereby believes the classical proposition that Sydney is near *l*.

Even though what is asserted, and what is transferred, is a place-neutral proposition, the audience comes to believe the same classical propositions as the speaker. Since their situations are the same, the *lekta* cannot determine different

truth-conditions with respect to their respective situations. That is why the belief transfer can take place directly at the *lekton* level, in this type of case. So it is not true that the content of assertion can only be the complete content. In shared-situation cases, the *lekta* go proxy for the complete contents (Barwise 1989: 253), and we can assert them (Egan 2007).

Just as we can assert *lekta* in shared-situation cases, we can agree or disagree about them. The function of assertion is to transfer belief: the speaker says something, and the audience is supposed to accept what the speaker has said. Sometimes, however, the speaker says something and the audience does *not* accept what the speaker has said: they disagree. Like assertion, disagreement, in shared-situation cases, can be over *lekta*. I say ‘Sydney is nearby’, and you respond ‘No, it isn’t’. If our locations were different, there would be no point in so disagreeing about *lekta*: our respectively accepting and denying the place-relative proposition that Sydney is nearby would not entail any truth-conditional incompatibility between our beliefs. But genuine (dis)agreement over *lekta* is possible whenever the situation of evaluation is shared.

This straightforwardly applies to the temporal case: there can be genuine agreement or disagreement over temporally neutral propositions whenever the time of evaluation is shared. If it is not—as in cases of diachronic disagreement—then the content in dispute must be temporally specific. Here, Richard seems to be right: “diachronic agreement or disagreement seems to be, of necessity, a matter of agreement or disagreement about something temporally specific” (Richard 2003: 40). Even that has been disputed, however. This brings us to the second type of example which casts doubt on the view that, in matters of agreement or disagreement, only complete contents count.

2.3.2 Faultless disagreement

Sometimes, it seems that we genuinely disagree about a certain *lekton*, even though we are not in the same situation. Thus, looking at a painting, I say: ‘This is beautiful’. You disagree: ‘No, it’s ugly’. In a sense, we are both right, since for me it is beautiful, while for you it is ugly; but we disagree nonetheless. Or consider epistemic modals. I say ‘The treasure might be under the palm tree’. I am right since, for all I know, the treasure might be there—nothing in my epistemic state rules out the treasure’s being there. Later, however, I learn that the treasure is not on the island (where the palm tree is). This rules out the treasure’s being under the palm tree, and in my new epistemic situation, I assert: ‘The treasure cannot be under the palm tree’. Again, I am right since, in my new epistemic situation, there is something that rules out the treasure’s being under the palm tree. What is strange, however, is that I can now disagree with my former self. I can say: ‘I was wrong—the treasure cannot be under the palm tree’. How can that be? If I was right, given my epistemic situation then, how can I later judge that I was wrong?

Such cases of ‘faultless disagreement’ suggest that sometimes at least, agreement or disagreement is about the *lekton*, even though the disagreeing parties evaluate the *lekton* with respect to distinct situations. This makes sense if one is a Radical Relativist. For a Radical Relativist, the *lekton* is complete. It is the content—that which one asserts, believes, and over which one agrees or disagrees with others.¹⁰ The situation of evaluation is not an aspect of content (broadly understood) but something external to content. As Prior puts it,

Aristotle . . . says that ‘statements and opinions’ vary in their truth and falsehood with the times at which they are made or held, just as concrete things have different qualities at different times; though the cases are different, because *the changes of truth-value of statements and opinions are not properly speaking changes in these statements and opinions themselves, but reflexions of changes in the objects to which they refer* (a statement being true when what it says is so, and ceasing to be true when that ceases to be so). (Prior 1967: 16; emphasis mine)

So a Radical Relativist has a story to tell about faultless disagreement, and that involves giving up the claim that agreement or disagreement in non-shared situations is of necessity a matter of agreement or disagreement about classical content. (Strangely enough, Mark Richard is, or seems to be, among the recent advocates of Radical Relativism—see Richard 2004.)

Some authors have claimed that faultless disagreement merely calls for a distinction between the content over which we agree or disagree (the *lekton*) and the utterance’s possible-world truth-conditions (Kölbel, this volume; Lasersohn 2005¹¹). The latter depend upon, and covary with, the situation of evaluation. On this view the Moderate Relativist framework, with its distinction between two levels of content, is sufficient to account for faultless disagreement. I think these authors are mistaken, however. There are plenty of cases in which we must distinguish between the *lekton* and the complete content, but in which there can be no genuine (dis)agreement about the *lekton*. Thus I call you on the phone, and commenting upon my situation I say ‘It is raining’. If you say ‘No, it isn’t’, meaning that there is no rain in *your* situation, there is misunderstanding rather than genuine disagreement. Or, adapting Barwise’s example, suppose that Watson says ‘The salt is left of the pepper’, and Holmes,

¹⁰ MacFarlane protests that in many cases, *he* does not take the *lekton* to be complete, even though he is a Radical Relativist: “So, for example, I could say (with the temporalist) that the time of utterance is not part of the *lekton*, but rather part of what a use of the *lekton* concerns, and still tell my story about the assessment-sensitivity of future contingents” (MacFarlane, p.c.). But MacFarlane’s Radical Relativism is not absolute, as we have seen (footnote 3): MacFarlane is a Radical Relativist with respect to some examples but not with respect to others. My claim is conditional: *if* one is a Radical Relativist with respect to a given type of sentence, e.g. ‘The treasure may be under the palm tree’, whose truth-value is relative to something in addition to a possible-world (in this case, an epistemic state), *then* one holds that the *lekton*, that is, the content of the sentence independent of that thing, is complete and can be the object or assertion, belief, or (dis)agreement.

¹¹ My comments on Lasersohn’s paper are based upon a draft dated 2004.

speaking from his own perspective, replies ‘No it is not’. Clearly, there is no substantive disagreement here. If Watson and Holmes are each talking about their own perspective (distinct from that of the other), there is misunderstanding rather than genuine disagreement. The same considerations apply to the temporal case. At time *t*, you say ‘It is raining’. Later, when the sun is shining again, you say ‘It is not raining’. You cannot conclude ‘so I was wrong’. Here, as Richard points out, genuine disagreement can only be about temporally specific contents.

So Moderate Relativism by itself, with its two levels of content, does not provide a solution to the problem of faultless disagreement, contrary to what Kölbel and Lasersohn believe. This does not mean that no solution is available within the Moderate Relativist framework, however. I think a solution is available, along the following lines.¹²

2.3.3 Speaking for the community

The speaker says ‘The painting is beautiful’, and her audience disagrees. In the Moderate Relativist framework, the statement that the painting is beautiful must be evaluated with respect to some aesthetic standards, but those standards cannot be *the speaker’s* aesthetic standards. If that were the case, then the speaker’s ‘It is beautiful’ and the audience’s reply ‘No, it’s ugly’ would *not* contradict each other. When the speaker says ‘It’s beautiful’, she means something intuitively stronger than merely ‘It is beautiful by my standards’ or ‘I find it beautiful’. To be sure, when realizing that the audience does not share her taste, the speaker may retreat to the weaker statement: ‘I find it beautiful’. This gives us a hint as to what the proper analysis of the stronger, unqualified statement should be. I suggest the following: ‘It is beautiful’ means something like *It is beautiful for us*, that is, for the community to which the speaker and his audience belong. When the audience says ‘No it isn’t’, the speaker realizes that the stronger statement is incorrect (since the audience does not actually find the painting beautiful), and she retreats to the weaker statement.

Lasersohn has an objection to that analysis. Were it correct, he says, it would not be possible for the speaker to *maintain* the stronger statement and keep disagreeing with his audience. But this is clearly possible:

- A: This painting is beautiful.
- B: No, it’s ugly.
- A: I tell you it is beautiful!
- B: Absolutely ugly!

¹² I will discuss only a case involving standards of taste. Whether the type of solution I put forward extends to other cases (such as epistemic modals) is an issue I will leave for further research. (See von Fintel and Gillies 2004 and 2006 for a step in that direction.)

To handle this type of counter-example, we must introduce a certain flexibility as to what counts as 'the community' from whose point of view the *lekton* is meant to be evaluated. I may judge that my audience deviates, by her bad taste, from the aesthetic standards of the community to which we both belong, those standards being fixed by e.g. the community's experts. This enables me to disagree with my audience, even though what I claim is that the painting is beautiful *for us*, i.e., for our community.

Whether, at the second step, the speaker retreats to the weaker statement ('I find it beautiful') or not ('I tell you it's beautiful!'), the right conclusion to draw is that there is *no* genuine faultless disagreement in those cases (Stojanovic forthcoming). The speaker who retreats to the weaker statement realizes that he was *mistaken* when he assumed that the painting was beautiful for both him and his audience; so he was at fault. And when the speaker refuses to retreat and maintains his position, because he "judges that his audience deviates, by her bad taste, from the aesthetic standards of the community to which they both belong, those standards being fixed by e.g. the community's experts", he finds *her* at fault. In both cases, there is genuine disagreement but it is not faultless. Nor are those cases cases in which the speaker and his audience evaluate the *lekton* with respect to different standards: they both appeal to the standards of the community to which they belong.

Those, however, are the simple cases. More difficult to handle is the following case (inspired from an example by Lasersohn). Suppose the speaker finds the painting beautiful, *and no one else (and in particular, no expert) does*. Suppose, moreover, that the speaker knows that he is alone in finding the painting beautiful. Still, as the opinionated person he is, he maintains: 'This is beautiful!' Clearly, he does not mean simply that *he* finds the painting beautiful. He refuses to retreat, and maintains the stronger claim. Can we handle that example consistently with the view that 'This is beautiful' means *This is beautiful for us*? Can we maintain that the speaker, even in such a case, appeals to the standards of the community?

I think we can. As Johan Brännmark writes (in connection with moral judgments),

When speaking for a collective, I cannot deviate from its present view without ceasing to speak for the collective. But communities are multi-generational and by their very nature they always have one foot in the future . . . So I think we can distinguish between two ways in which we can speak for the community: first, there is the purely representative one; second, there is a progressive one. In passing progressive moral judgment *I am deviating from the present community in the direction towards which I find that the community would move if the people in it thought things through really well*. (Brännmark forthcoming emphasis mine).

In this sort of case—when no appeal to experts is relevant and each of the two disputants is expressing his or her own taste and trying to impose it—it makes

sense to say that they are “both right”, even though they disagree: for they are both entitled to interpret the community’s standards the way they want in passing progressive judgment. This freedom explains how they can disagree about the truth-value of ‘This is beautiful’ even though, in evaluating the statement, *they both appeal to the standards of the community*. This is possible because the standards in question are not fixed once for all—they are up to us. Each of us can contribute to shaping them, and this is what we do in this type of case when we say ‘It is beautiful’ or ‘It is ugly’. Austin would perhaps make this point by saying that there is a performative element in such statements.

If we accept this explanation of the phenomenon, we can maintain that agreement or disagreement is about complete contents, except when the situation of evaluation is shared (in which case the *lekton* can go proxy for the complete content). Alleged ‘faultless disagreement’ cases such as those we have discussed are no exception. In such cases, the disagreement is about the Austinian proposition consisting of the *lekton* together with the standards of the community.¹³ The speaker says that the painting is beautiful (for the community), and the audience denies that it is beautiful (for the community). The disagreement ultimately bears upon what the community standards are, or should be. If the disagreement bears upon what the community standards *are*, it is not faultless. If, as in the last case we considered, it bears upon what the standards *should be*, then it is, arguably, faultless, but even in that case the disagreement is over the complete content.

References

- Barwise, J. (1989). ‘Situations, Facts, and True Propositions’. In his collection *The Situation in Logic*, Stanford: CSLI Publications, 221–54.
- and J. Etchemendy (1987). *The Liar: An Essay on Truth and Circularity*. New York: Oxford University Press.
- Brännmark, J. (forthcoming). ‘Speaking for the Community: On Morality and We-Judgments’. *MS*.
- Dummett, M. (1997). ‘Existence, Possibility and Time’. In Meggle, G., *Analyomen* 2/1, Berlin: de Gruyter, 43–67.
- Egan, A. (2007). ‘Epistemic Modals, Relativism and Assertion’. *Philosophical Studies* 133: 1–22.
- Evans, G. (1985). ‘Does Tense Logic Rest on a Mistake?’ In his *Collected Papers*, Oxford: Clarendon Press, 343–63.
- Fintel, K.v. and A. Gillies (2004). ‘“Might” Made Right’. *MS*.
- (2006). ‘CIA leaks’. *MS*.
- Frege, G. (1967). *Kleine Schriften*. Hildesheim: Olms.

¹³ Since both disputants implicitly appeal to the standards of the community, we can also treat that case as a case in which the situation of evaluation is shared.

- García-Carpintero, M. (this volume). 'Relativism, Vagueness and What is Said'.
- Geach, P. (1955). Review of Mates 1953. *Philosophical Review* 64: 143–5.
- Hintikka, J. (1973). *Time and Necessity: Studies in Aristotle's Theory of Modality*. Oxford: Clarendon Press.
- Kaplan, D. (1989). 'Demonstratives'. In Almog, J., Perry, J. and Wettstein, H. (eds.), *Themes from Kaplan*, New York: Oxford University Press, 481–563.
- King, J. (2003). 'Tense, Modality, and Semantic Values'. *Philosophical Perspectives* 17: 195–245.
- Kölbel, M. (this volume). 'Motivations for Relativism'.
- Laserson, P. (2005). 'Context Dependence, Disagreement, and Predicates of Personal Taste'. *Linguistics and Philosophy* 28: 643–86.
- MacFarlane, J. (2003). 'Future Contingents and Relative Truth'. *The Philosophical Quarterly* 53: 321–36.
- (forthcoming). 'Nonindexical Contextualism'. *MS*.
- McGinn, C. (1982). 'The Structure of Content'. In Woodfield, A. (ed.), *Thought and Object*, Oxford: Clarendon Press, 207–58.
- Mates, B. (1953). *Stoic Logic*. University of California Press.
- Perry, J. (1979). 'The Problem of the Essential Indexical'. Repr. in Perry 1993, 33–52.
- (1986). 'Thought without Representation'. Repr. in Perry 1993, 205–19.
- (1993). *The Problem of the Essential Indexical and Other Essays*. New York: Oxford University Press.
- Prior, A. (1967). *Past, Present, and Future*. Oxford: Clarendon Press.
- and K. Fine (1977). *Worlds, Times and Selves*. London: Duckworth.
- Recanati, F. (1996). 'Domains of Discourse'. *Linguistics and Philosophy* 19: 445–75.
- (1997). 'The Dynamics of Situations'. *European Review of Philosophy* 2: 41–75.
- (2000). *Oratio Obliqua, Oratio Recta. An Essay on Metarepresentation*. Cambridge, Mass.: MIT Press/Bradford Books.
- Richard, M. (1981). 'Temporalism and Eternalism'. *Philosophical Studies* 39: 1–13.
- (1982). 'Tense, Propositions, and Meanings'. *Philosophical Studies* 41: 337–51.
- (2003). 'Introduction to Part I: The Philosophy of Tensed Language', In Jokic, A. and Smith, Q. (eds.), *Time, Tense and Reference*, Cambridge, Mass.: MIT Press/Bradford Books, 25–45.
- (2004). 'Contextualism and Relativism'. *Philosophical Studies* 119: 215–42.
- Stalnaker, R. (1978). 'Assertion'. *Syntax and Semantics* 9: 315–32.
- Stanley, J. and Z. Szabó (2000). 'On Quantifier Domain Restriction'. *Mind and Language* 15: 219–61.
- Stojanovic, I. (forthcoming). 'Talking about Taste: Disagreement, Implicit Arguments and Relativized Contents'. *Linguistics and Philosophy*.

3

Semantic Relativism and the Logic of Indexicals

Stefano Predelli and Isidora Stojanovic

Traditional relativist approaches to matters of philosophical interest have recently been re-assessed from a genuinely semantic viewpoint. The ensuing semantically oriented relativist apparatus has been applied not only to issues generally considered as at least *prima facie* amenable to relativist treatments, such as questions pertaining to the palatability of liquorice or the inelegance of blue suede shoes, but also to areas not customarily falling within the province of the relativists' interests, such as the treatment of vague expressions, the interpretation of epistemic modals, and the analysis of future contingents.¹ The philosophical relevance and adequacy of Semantic Relativism (SR) has thus been assessed on a two-fold level. What is at issue is, on the one hand, whether the resources provided by SR are adequate for the solution of the puzzles of vagueness or future contingency, and, on the other, whether the resulting proposals are indeed interestingly related to traditional forms of philosophical relativism.

But SR is arguably of considerable philosophical interest independently of either of the issues to which we just alluded. As we explain below, its semantic apparatus is closely related to, but importantly different from, the framework for semantic systems of philosophical interest, in particular the theoretical constructs developed for the analysis of indexical intensional languages. It is thus at least initially legitimate to suspect that the discrepancies between SR and the traditional treatment of indexicality also result in contrasting conclusions having to do with questions of general semantic import, such as conclusions pertaining to the relationships between truth and meaning, and to the significance and scope of semantic validity.

In this essay, we thus momentarily eschew the independently interesting questions pertaining to the applicability of SR, and to the extent to which SR echoes genuinely relativistic viewpoints. Instead, we focus on the consequences of SR pertaining to the logic of indexical intensional languages, and, as a consequence,

¹ See for instance MacFarlane 2003 and Egan, Hawthorne, and Weatherson 2005.

to the relationships between meaning and truth. In particular, we aim at explaining the differences between the classic and relativistic definitions of truth and validity, and at unveiling certain fundamental philosophical discrepancies between their approaches to logic and semantics.

For a variety of pedagogical and methodological reasons, our discussion often focuses on the semantic interpretation of artificial languages equipped for the treatment of indexicality and intensionality—more precisely, it focuses on a model-theoretic approach to certain simple indexical intensional fragments. The quasi-formal semantic analysis of some expressions within the simple language we elect as paradigmatic has the methodological advantage of bringing to the foreground the properties of their natural-language counterparts that are of relevance for our aim, while disregarding a variety of independently interesting but distracting other aspects. Section one provides a brief introduction to some aspects of languages of that type and of their model-theoretic analysis. Section two presents what we take to be the relevant point of discrepancy between the classic model-theoretic treatment for indexical languages, such as Kaplan's interpretive system, and the framework developed from the viewpoint of SR. The traditional view, so we explain, is characterized by what we call the *Classic Reduction*, namely the identification of the elements required for the interpretation of indexicals and the parameters relevant for the definition of (unrelativized) truth.

Section three focuses on the motivations for the traditional commitment to the Classic Reduction, with particular attention to Kaplan's 'Demonstratives'. We tentatively conclude that the Classic Reduction reflects an understanding of logical relations grounded on what we call *warranted utterability*, and that SR provides a more appropriate apparatus for a conception of validity as *truth in virtue of meaning*. Section four discusses some logical consequences of SR, and contrasts them with certain typical results of the classic approach to indexical languages.

3.1 PRELIMINARIES

In this section, we briefly discuss certain general aspects of the analysis of indexicality and intensionality in natural languages, and their relationships to the model-theoretic study of relevant artificial languages. In particular, we present the idea of indexicality and meaning-encoded contextuality (sub-section 1.1), and we rehearse the customary compositional treatment of complex expressions within intensional languages (sub-section 1.2).

3.1.1 Contexts and indexicality

It is by now a relatively commonplace observation that the contextual sensitivity affecting the interpretation of English expressions such as 'I', 'there', or 'now' is *meaning-encoded*. In particular, these and other indexical expressions are

conventionally associated with a particular descriptive material, of such a kind that their semantic behaviour is importantly dependent upon the vicissitudes of the context with respect to which semantic interpretation is being carried out. On the basis of these initial gestures towards an understanding of indexicality, it seems natural to associate an indexical expression such as 'I' with a descriptive condition, such as that of being the contextually relevant agent, and to determine its semantic value on the basis of the interaction between that condition and the properties of the interpretive context. Accordingly, the results of semantic interpretation take a relativized form. For instance, 'I' is assigned a particular individual i as referent with respect to a context c , or, equivalently, it is the expression-context pair $\langle 'I', c \rangle$ rather than the lone expression-type which receives i as its value.²

For a variety of semantic tasks, the exact composition of the descriptive condition in question may be disregarded as irrelevant, and the semanticist's attention may concentrate on the association between contexts and value-determining parameters—say, in a simple approach to 'I' and similar cases, on the association between contexts and referents. At least for certain purposes, then, the *meaning* of an indexical may be reflected within a formal language by a function from contexts to semantic values, that is, by what is customarily called a *character*. By the same token, the multifarious complexity of the contextual setting which supplies the information required for the interpretation of particular uses of an indexical may be reduced to the formally tractable format of an n -tuple containing the sort of items required by the characters of the expressions under study: an agent (for 'I'), a location (for 'here'), etc.³

Let us then temporarily focus on a simple fragment containing no intensional operators and (the counterparts for) 'I' and 'you' as its only indexicals. A structure m for such a fragment is of the form $\langle U^m, C^m, I^m \rangle$, with U^m a class of individuals, I^m an interpretation function, and C^m a class of contexts c of the form $\langle i, j \rangle$, where $i, j \in U^m$ (intuitively, respectively an 'agent' and an 'addressee').⁴ An expression e within this fragment is assigned a semantic value $[[e]]^m_c$ with respect to a structure m and a context $c \in C^m$ in the obvious fashion. (We ignore for simplicity's sake mention of an assignment of values to variables.) A sentence s is evaluated as true in a structure m with respect to a context c iff the compositional apparatus assigns the appropriate semantic value to it with respect to m and c : $\text{True}^m_c(s)$ iff $[[s]]^m_c = 1$. Validity and cognate notions may thus eventually be defined in the customary unrelativized terms: for instance, $\text{Valid}(s)$ iff for all m and c , $\text{True}^m_c(s)$.

² See Kaplan 1977 and Lewis 1980. One of us has complained that this terminological choice may be pedagogically unsuitable, and may in fact be at the root of some of the problems we briefly mention in Section 2 below (see Predelli 2005).

³ Note though that Stojanovic (2005) has recently explored a different kind of formal account of the descriptive conditions associated with indexicals.

⁴ See Kaplan 1977: 543, for a parallel notion of structure.

3.1.2 Circumstances and intensionality

It is a relatively well-established assumption that certain aspects of the semantic behaviour of some English expressions may be explained by considering the semantic value of other expressions across appropriate parameters. So, for instance, at least certain simple features of ‘it is necessary that’ may be rendered by means of clauses which assign a semantic value to sentences of the form ‘necessarily ϕ ’ on the basis of the behaviour of ϕ across modally relevant parameters—according to mainstream approaches, so-called possible worlds. Equivalently, at least an interesting part of the conventional profile for ‘necessarily’ may be reflected by a (constant) character eventually related to a function which yields a semantic result on the basis of what is commonly called the *intension* of the expression upon which it operates.

The customary semantic treatments for languages containing intensional operators are thus characterized by the introduction of appropriate parameters of relativization within the compositional system. Collections of such parameters are typically labelled *circumstances*, in conformity to the usage recommended by Kaplan in Kaplan 1977. As a result of the demands of compositional analysis, the interpretive procedure attributes to the expressions ultimately under analysis the sort of relativization required by the clauses for the intensional operators: such and such a sentence turns out to have this or that semantic value only with respect to this or that circumstance of evaluation. More formally, and focusing for simplicity’s sake on a fragment containing no indexical expressions and a few modal operators as its only intensional operators, a structure m is of the form $\langle U^m, W^m, \alpha^m, I^m \rangle$, with U^m a class of individuals, W^m a class of circumstances (possible worlds), α^m an element of W^m , and I^m an interpretation function.⁵ (For simplicity’s sake we ignore issues of accessibility among elements of W^m , and we assume a single domain of individuals.) An expression e within this fragment is assigned a semantic value $[[e]]^m_w$ with respect to a structure m and a circumstance $w \in W^m$ in the usual manner.

The sort of circumstance-relativity carried over by the assignments of semantic values is eventually removed by the post-compositional clauses responsible for the definition of *truth (in a structure)*. More precisely, a sentence s is evaluated as true in a structure m iff the compositional apparatus assigns the appropriate semantic value to it with respect to m and α^m : $\text{True}^m(s)$ iff $[[s]]^m_{\alpha} = 1$. Validity and related notions may thus eventually be accounted for in the familiar way: for instance, $\text{Valid}(s)$ iff for all m , $\text{True}^m(s)$.

⁵ α^m is what is usually called the ‘designated world’. See for instance Kripke 1963: 63 for a parallel notion of a (model) structure.

3.2 SEMANTIC RELATIVISM

In this section, we sketch the familiar features of the classic, Kaplan-inspired semantics for intensional indexical languages, and we contrast it with an approach with the framework for SR.

3.2.1 SR-structures and the Classic Reduction

Intensionality and indexicality are compatible phenomena. In particular, a language may not only include intensional operators side by side with indexical expressions; it may also allow for indexical intensional expressions, that is for operators contributing a certain circumstance-sensitive function on the basis of contextual factors. Kaplan's *LD* is indeed a language of this type, and some of the most interesting features of its semantic analysis emerge precisely as a result of the interaction between indexicality and intensionality. For simplicity's sake, in this essay we focus on the cases of *modal* intensional operators, and in particular on the behaviour of the indexical modal operator 'actually'. The simple format for contexts envisioned in sub-section 1.1 (and the corresponding composition of a structure) will thus need to be enriched so as to include at least mention of a possible world: any c in C^m is now understood as at least a triple of the form $\langle i, j, w \rangle$, with, as before, $i, j \in U^m$ (respectively the agent and addressee of c) and $w \in W^m$ (the possible world of c).

In Section 1, we briefly rehearsed certain familiar traits of the model-theoretic treatment of indexical languages (1.1) and intensional languages (1.2). The formal demands which the phenomenon of indexicality imposes upon semantic interpretation consist in the inclusion of a class of contexts among the items constitutive of a structure, and in the resultant relativization of the notions of semantic value and truth: given an expression e , its semantic value $[[e]]^m_c$, and hence its truth-value $\mathbf{True}^m_c(e)$, are relative to the choice of an element within C^m . What intensionality requires, on the other hand, is that structures supply a certain collection of circumstances, that the compositional assignment of semantic values keep track of the choice of elements from that collection, and that truth-value be defined on the basis of the structure's choice of a privileged circumstance. These formal requirements suggest an obvious 'cumulative' solution for the semantic analysis of languages that are both intensional *and* indexical: structures are now understood as collections of the form $\langle U^m, W^m, C^m, I^m, \alpha^m \rangle$, semantic values $[[e]]^m_{c,w}$ are relativized both to contexts and circumstances along the familiar lines of 'double index' semantics, and the truth-value of a sentence s is defined as follows: $\mathbf{True}^m_c(s)$ iff $[[s]]^m_{c,\alpha} = 1$.

We take this straightforward cumulative suggestion as a trademark of SR. It is not, however, the traditional proposal for the treatment of intensional indexical languages such as Kaplan's *LD*. In particular, a classic structure m for *LD* is (*modulo* the usual simplifications) a quadruple $\langle U^m, W^m, C^m, I^m \rangle$ that does not explicitly single out a privileged circumstance. The assignment of semantic value within the compositional stage follows the usual double-index format sketched above, but the definition of truth takes advantage of the appropriate contextual item as a substitute for the missing privileged element of W^m : $\text{True}^m_c(s)$ iff $[[s]]^m_{c,k} = 1$, where $k = c_w$, the circumstance 'determined' by context. In Kaplan's terminology, where $\{\phi\}^M_{c,f}$ is the intension (*content* in Kaplan's terminology) for ϕ in structure M with respect to a context c and an assignment f of values to variables:

Definition: ϕ is true in the context c (in the structure M) iff for every assignment f , $\{\phi\}^M_{c,f}(c_T, c_w) = \text{Truth}$. (Kaplan 1977: 547)

We refer to this identification of contextually determined parameters and privileged circumstance as the *Classic Reduction* of the model-theoretic requirements provided by the phenomena of indexicality and intensionality.

As we mentioned, our attention in this essay focuses for simplicity's sake on modal phenomena. Accordingly, the sort of indexicals of particular interest for our purposes are those whose characters make non-trivial reference to the modal dimension, such as the modal indexical operator 'actually'. If the fragment under study is restricted so as to include only indexicals of this sort, a context c may be understood as supplying (in fact, harmlessly, as identical with) merely a possible world, the 'context world' c_w . Within this framework, then, the Classic Reduction may be phrased as the proposal that the role played by the model's 'designated world' α^m within the definition of truth relative to c be filled by c itself: $\text{True}^m_c(s)$ iff $[[s]]^m_{c,c} = 1$. Conversely, SR's rejection of the Classic Reduction amounts to the distinction between the interpretive parameter appropriate for modal indexicals such as 'actually', and the circumstance relevant for truth.

When applied to the study of natural languages, this distinction is typically assumed to reflect a genuinely relativistic standpoint, because it presumably makes room for the possibility that utterances endowed with the same semantic content receive distinct truth-values, that is, it presumably provides the resources needed for a semantic treatment of so-called 'faultless disagreement'.⁶ Whether faultless disagreement is indeed a phenomenon worthy of semantic attention, and whether SR does provide the semantic resources appropriate for its analysis, are issues we do not discuss in this essay.⁷ We thus remain neutral with respect to the diatribe whether SR does indeed deserve the label of 'relativism', and with respect to the applications of SR to questions pertaining to this or that natural-language

⁶ On this idea see for instance Kölbel 2002.

⁷ Stojanovic (forthcoming) has argued that semantic relativism does not explain how there could be such a thing as faultless disagreement.

phenomenon. Instead, in sections 3 and 4, we stress certain general properties and consequences of SR, and some of the divergences between SR and the traditional approach to indexical languages. A few independent considerations having to do with the treatment of contexts within contemporary SR are however in order before we proceed with our discussion of logic from a relativistic point of view.

3.2.2 SR and Context-Shifts

As highlighted in the previous paragraphs, what is of relevance for SR, at least for the purpose of this essay, is its commitment to a ‘double parameter’ approach to matters of semantic interpretation. In the familiar philosophical jargon, this amounts to the idea that different items are devoted to the interpretation of the indexicals (that is, to the identification of the appropriate content, in Kaplan’s sense of the term) and to the choice of a privileged circumstance (that is, to the assignment of an unrelativized truth-value to that content). In this sense, then, what is characteristic of SR is not the sheer presence of a multiplicity of contexts, but rather the insistence on different *roles* for them—in MacFarlane’s parlance (MacFarlane 2003), respectively the roles of ‘interpreter’ and ‘initializer’.

It is worth noting in this respect that mere ‘multiple context’ semantics *may* be incorporated within a classic framework without any commitment to SR, and in a manner compatible with the Classic Reduction. Many, including one of us (see Predelli 1998a, 1998b, and 2005), have argued that certain utterances are appropriately represented by evaluating some occurrences of expressions within one sentence with respect to one context, and others with respect to another. Kaplan himself suggested the simple example of ‘now you see it and now you do not’, an utterance that is non-contradictory precisely insofar as the occurrences of the temporal indexicals refer to distinct instants. Whether the appropriate input for the semantic interpretation of this utterance is one which provides different contexts, suitably anchored to each occurrence of ‘now’, is a question that need not detain us here. What is of greater urgency is the fact that, were this indeed to be the case, the resulting analysis would not *thereby* qualify as a treatment in the spirit of SR: both contexts, in this case, concur in the role as interpreters, that is, they function merely at the stage relevant for the establishment of Kaplanesque content. To put it otherwise: the discussion of the treatment of ‘now you see it and now you do not’ may well raise interesting ‘pre-semantic’ issues having to do with the proper representation of an utterance in a format amenable to a systematic semantic analysis, but has no bearing on the very structure of that analysis.⁸ This claim is confirmed by considerations pertaining to the proper treatment of logical truth within such a harmless ‘multi-context’ apparatus. On any independently plausible account of validity, context variations of this kind

⁸ On this sense of ‘pre-semantic’ see Predelli 2005.

must be kept at bay by fiat: obviously, the true utterability of ‘now you see it now you do not’ should not be taken to entail any surprising result pertaining to what counts as a contradiction or a valid sentence.

The reason for insisting on this rather obvious point has to do with certain independently important, but for our purpose irrelevant features of certain applications of SR, such as MacFarlane’s treatment of future contingents (MacFarlane 2003). MacFarlane’s proposal falls squarely within the SR camp, insofar as the choice of a privileged circumstance of evaluation is theoretically independent of the choice of the parameter which supplies the interpretation of the indexicals. Still, in his application of the SR framework to the problem at hand, MacFarlane occasionally selects the items provided by the context of assessment for interpretive purposes: his context of assessment thus plays a dual role, as initializer and as partial interpreter. Whether this employment of the resources of SR is independently important, and whether it provides a satisfactory solution to the problem of future contingents, are questions of no immediate relevance here. What matters is rather that the mere fact that a context distinct from the context of interpretation may sometimes be invoked for interpretive purposes is not, by itself, what allows MacFarlane to qualify his approach as a version of Semantic Relativism. It is for this reason that the discussion in what follows proceeds on an important but for our purpose harmless simplification, that is, a sharp division of labour for the contexts in question: the interpretive role shall be assumed to be exhausted by the ‘context of interpretation’, and the role for the ‘context of assessment’ shall be confined to initializing purposes.

3.3 THE VAGARIES OF ACTION

Let us take stock. In the first two sections of this essay, we have focused on the distinct semantic demands put forth by the phenomena of indexicality and intensionality. In the absence of arguments to the contrary, then, model-theoretic treatments of indexical intensional languages ought to be amenable to the possibility that the parameter relevant for the interpretation of indexical expressions (and, a fortiori, of modal indexicals such as ‘actually’) be distinct from that assumed as appropriate for the determination of truth. It is this possibility that characterizes systems in the spirit of SR as we understand it, and which is denied by fiat by the traditional commitment to the Classic Reduction.

Our claim that indexicality and intensionality are distinct phenomena is uncontentious. More importantly, such distinction plays a fundamental role in traditional approaches to indexical and intensional languages such as Kaplan’s or Lewis’s (Kaplan 1977, Lewis 1980). Indeed, it is a trademark of their analyses that the parameters required for the interpretation of indexical expressions (Kaplan’s and Lewis’s ‘contexts’) do not coincide with the *relata* needed for the evaluation of intentional operators (Kaplan’s ‘circumstances’ or Lewis’s ‘indices’). If this is

the case, it is by no means obvious that an appropriate definition of truth ought to proceed along the lines suggested by the Classic Reduction: the idea of a privileged circumstance of evaluation required for the semantics of an intensional language may not be assimilated without further ado with the notion of context needed for the interpretation of indexicals possibly occurring in it.

Still, explicit arguments in favour of the Classic Reduction are not easy to come by. One of the best known presentations of the traditional approach to indexical intensional languages, Kaplan's 'Demonstratives', only devotes a few sentences to this issue:

Since the content of an occurrence of a sentence containing indexicals depends on the context, the notion of *truth* must be relativized to a context. If *c* is a context, then an occurrence of φ in *c* is true iff the content expressed by φ in this context is true when evaluated with respect to the circumstance of the context. We see from the notion of truth that among other aspects of a context must be a possible circumstance. Every context occurs in a particular circumstance, and there are demonstratives such as 'actual' which refer to that circumstance. If you try out the notion of truth on a few examples, you will see that it is correct. If I now utter a sentence, I will have uttered a truth just in case *what I said*, the content, is true in *these* circumstances. (Kaplan 1977: 522–3)

Their brevity notwithstanding, Kaplan's considerations are worthy of attention. The reason for the Classic Reduction that transpires from Kaplan's suggestion that we 'try it out on a few examples', so we argue in the following paragraphs, unveils the fundamental (and controversial) metasemantic assumptions that give the traditional definition of truth its misleading air of harmless neutrality.

According to the passage from 'Demonstratives' quoted above, the motivation for the Classic Reduction lies in the alleged fact that, whenever an utterance is being produced, what matters for the assessment of its truth-value are the conditions within which it takes place ('*these* circumstances'), that is, presumably, the state of affairs which also supplies the parameters needed for the interpretation of 'demonstratives such as "actual"'. This idea returns with added emphasis in Kaplan's 'Afterthoughts', where the context's 'dual role' is motivated in terms of the 'underlying idea . . . [that] the actual world is where we actually are . . . now', and where validity is defined as 'truth no matter what the circumstances were *in which the sentence was used*' (Kaplan 1989: 595–6, emphasis in the original). The 'underlying idea' behind the Classic Reduction, in other words, registers a regularity resulting from presumed inevitable properties of *utterances*, that is, of *uses* of a sentence: whenever a sentence is uttered, so we are told, the circumstance relevant for the determination of truth *tout court* is provided by the very conditions within which that utterance takes place.

Whether this claim is indeed correct depends upon the applicability of relativistic strategies to this or that example: if contemporary relativists turn out to be right, the circumstances appropriate for the truth of utterances involving, say, expressions of taste or epistemic modals, may well fail to coincide with the obvious parameters of utterance. Yet, more general considerations also point

towards the conclusion that, independently of the question whether Kaplan's claim about utterances and their circumstances is indeed correct, the sheer decision to ground the *definition* of truth on questions of use reflects a substantial and contentious approach to logic and semantics. We devote the remainder of this section to an explanation of this claim.

It is a general assumption (common to both classic and relativistic approaches) that simple expressions are conventionally assigned certain meanings (Kaplan's 'characters'), and that the syntactic procedures for the formation of complex expressions conventionally achieve semantic results of one type or another. On the basis of this assumption, some sentences end up being endowed with certain semantic properties solely in virtue of the meaning of the expressions they involve, and/or of the semantic effects of their syntactic layout. For instance, as long as '&' is interpreted in terms of a constant character suitably associated with the truth-function of conjunction, sentences such as 'P & Q' and 'Q & P' end up sharing their semantic values with respect to any particular context and circumstance of evaluation: for all c and w , $[[P \& Q]]_{c,w} = [[Q \& P]]_{c,w}$. According to a well-established metasemantic stand, the concern of logic lies precisely with 'verities of meaning' of this sort (Kaplan 1989: 585): logically true sentences such as 'P & Q $\leftrightarrow$ Q & P' are indeed those that are 'true solely in virtue of meaning'.

Rather uncontroversially, regularities affecting our *use* of certain expressions induce effects that are, or at least may well be, independent of their meaning. Some of these effects may be describable in terms of results about the truth-value of the utterances in question: some expressions may be *warrantedly utterable*, in the sense that any use of them would result in a true utterance, and some arguments may be warrantably utterable, in the sense that any scenario in which utterances of their premises turn out true are also scenarios in which their conclusion may be uttered truly. To cite a rather obvious colloquial example, utterances of 'there are at least six objects' presumably inevitably take place in circumstances in which at least six individuals exist, the English word-types 'there', 'are', 'at', 'least', 'six', and 'objects'. In this sense, then, 'there are at least six objects' turns out to be warrantably utterable: any of its possible uses would take place in circumstances with respect to which the content they express turns out true. Yet, again pretty much uncontroversially, existential claims of this sort apparently fail to be grounded on any aspect encoded in the meaning of these expressions, as testified by our reluctance to accept 'there are at least six objects' as true in virtue of meaning. More strongly: even if, surprisingly enough, such reluctance turned out to be misplaced, conclusions of truth ought to result from a discussion of the conventional profile of the expressions under discussion, and should not automatically stem from general assumptions embedded within the semantic apparatus to which our utterance is consigned. The unexpected validity of 'there are at least six objects' would then result from equally unexpected theses about the meaning of 'there', 'are', etc., rather than from antecedent constraints

for the semantic system. In the absence of claims to this effect, evidence of warranted utterability does not suffice for parallel conclusions of logical validity.

Metasemantic considerations along the lines sketched thus far are in fact prominent in Kaplan's approach to indexical languages. For instance, Kaplan warns us that

... it is important to distinguish an *utterance* from a *sentence-in-context*. The former notion is from the theory of speech acts, the latter from semantics. ... Thus the notion of ϕ *being true in c and M* does not require an utterance of ϕ . (Kaplan 1977: 522)

Kaplan's example in this respect hinges on a (controversial) assumption about utterances: 'utterances take time', and sufficiently slow utterances of even trivially valid (in the sense of meaning-guaranteed) arguments such as 'now I see it, therefore I see it now' may well turn out to involve a true utterance of its premise(s) and a false utterance of its conclusion. Parallel considerations hold with respect to the presumed warranted utterability of examples such as 'I am speaking now' (on a wide understanding of speaking, roughly compatible with a variety of methods for producing the utterance in question). Even granting the (debatable) premise that no use of that sentence may result in a false utterance, the sort of guarantee thus obtained remains idle with respect to the relationship between meaning and truth: nothing in the meaning of 'I', 'now', or 'speak' suffices for the conclusion that, in any context c , True_c ('I am speaking now'). In Kaplan's own words:

... there are sentences which express a truth in certain contexts, but not if uttered. For instance, 'I say nothing'. Logic and semantics are concerned not with the vagaries of actions, but with the verities of meanings. (Kaplan 1989: 584–5)

Of course, none of the above entails that no interesting apparatus may be developed for the systematic analysis of warranted utterability—an apparatus devoted to the explanation of why, say, no utterance of 'I am speaking now' may convey a false content, and of why very slow presentations of 'now I see it, therefore I see it now' may involve a true utterance of its premise, and a false utterance of its conclusion. Clearly, explanations of this type would be contingent upon the factual constraints for utterances of this or that type: 'now you see it, now you do not' may well be uttered truly in scenarios of slow enunciation, but not in situations in which utterances are always produced virtually instantaneously, say, by virtue of displaying written notes. Kaplan tentatively spells out the discrepancy between issues of warranted utterability and questions of truth in virtue of meaning in terms of his distinction between 'utterance-validity' and 'occurrence validity':

I am unclear even as to what arguments *ought* to come out as utterance-valid (as opposed to occurrence-valid). There are different notions of utterance-validity corresponding to different assumptions and idealizations. With no idealizations, the rules of repetition and double negation become invalid. (Kaplan 1989: 585, n. 40).

Yet, so Kaplan rightly insists, regardless of the structure and content of systems devoted to the analysis of this or that form of ‘utterance-validity’, none of their conclusions may be incorporated within a semantic apparatus, in a sense of semantics devoted to the study of the relationships between meaning and truth. Or, at least, none of them may be *immediately* incorporated within such an apparatus, independently of arguments for the conclusion that the production of this or that regularity is indeed encoded within the conventional meanings of the expressions in question.

Notwithstanding Kaplan’s occasional explicit distinctions between warranted utterability and the ‘verities of meaning’, some among the (presumed) unavoidable effects of language-use are notoriously surreptitiously allowed to enter his semantic system. So, for instance, the system of ‘Demonstratives’ yields results of validity for (the formal counterparts of) sentences such as ‘I am here now’ and ‘I exist’, without providing any motivation for the needed (and in our view suspicious) assumption that the meanings of ‘I’, ‘here’, or ‘now’ guarantee a result of True_c with respect to any c . More spectacularly, *indexical free* sentences such as ‘something exists’ are also accepted as valid (even ‘in the neotraditional logic that countenances empty worlds’, Kaplan 1977: 549), in spite of the obvious fact that no aspect of the meaning of ‘something’ or ‘exists’ may justify logical results different from those obtainable within the ‘neotraditional’ approach: if the account of the meaning of ‘something’ or ‘exists’ within systems that countenance empty worlds is indeed correct, and if ‘something exists’ does not turn out valid within those systems, it is mysterious how ‘something exists’ may become valid merely in virtue of being incorporated within richer, indexical languages.

At least in the absence of further clarifications and additional arguments, Kaplan’s motivation for the Classic Reduction is apparently the result of a similar disregard for the general metasemantic distinction so forcefully defended in the more lucid junctures of ‘Demonstratives’. As hinted thus far, and as stressed by the discussion of the examples in section 4, the traditional commitment to the Classic Reduction entails logically relevant conclusions: the identification of ‘initializer’ and ‘interpreter’ yields results of logical truth that would not be obtainable in its absence. But Kaplan’s reasons for accepting such conclusions are explicitly grounded on (presumed) regularities of use: ‘If I now utter a sentence’, so we are told, ‘I will have uttered a truth just in case *what I said*, the content, is true in *these* circumstances’. In Kaplan’s parlance, in other words, the motivation for the Classic Reduction rests on (presumed) regularities relevant for ‘utterance-validity’, but not for ‘occurrence-validity’. In our vocabulary, it rests on (presumed) evidence for conclusions of warranted utterability, rather than of meaning-guaranteed truth preservation.⁹

Of course, our tentative claim that Kaplan’s commitment to the Classic Reduction is grounded on an understanding of matters of logic along the lines

⁹ See also Castañeda’s notion of ‘lalic necessity’ in Castañeda 1957.

of warranted utterability does not suffice for the conclusion that SR provides the most appropriate treatment of questions of ‘truth in virtue of meaning’. Still, the considerations in this section support at least a *prima facie* case for the thesis that a conception of logical relations as meaning-warranted truth-preservation had better approach the Classic Reduction with a good deal of initial scepticism. It follows that, at least in the absence of convincing arguments in its favour, the Classic Reduction may at best be incorporated within an account of the (alleged) regularities affecting our use of language, and may well have to be withdrawn from an explanation of the systematic relationships between meaning and truth. On the basis of this cautious conclusion, we proceed to the analysis of some arguably interesting traditional equivalences that are rejected as invalid from the point of view of SR.

3.4 THE LOGIC OF SR

As explained thus far, what is of particular relevance for the discussion of SR are the interactions between indexicality and intensionality, and between the demands they impose upon the structure of semantic interpretation. For simplicity’s sake, in what follows we continue to focus on fragments involving only *modal* indexicals and operators. For this reason, what a suitable context ought to provide is merely a possible world—in fact, we proceed with the convenient simplification of taking contexts themselves to be possible worlds: $C^m \subseteq W^m$. By the same token, in the absence of non-modal intensional expressions such as temporal operators, circumstances of evaluation shall themselves be understood simply as possible worlds. The result in subsection 4.3 appeals to accessibility relations among possible worlds, and the case in sub-section 4.4 extends our consideration from simple propositional fragments to languages with quantifiers, whose semantic analysis admits of different domains of individuals for distinct possible worlds. In these sub-sections we thus tacitly expand our understanding of a structure m as including respectively a relation R^m among members of W^m , and as supplying a function p^m that partitions U^m across elements of W^m . In the latter case, as usual, assignments of semantic value shall include reference to an assignment f of values to variables along the customary lines. An expression that will receive some attention in what follows is the intensional indexical operator ‘Actually’, defined in classic (propositional) semantics as $[[\text{Act } \phi]]^m_{c,w} = 1$ iff $[[\phi]]^m_{c,c} = 1$.

3.4.1 ‘Actually $\phi \leftrightarrow \phi$ ’

The Classic Reduction, coupled with the traditional definition of validity for indexical languages, yields interesting logical outcomes. An important case is provided by

(1) $\text{Act } \phi \leftrightarrow \phi$

a formula that is classically valid but not SR-valid. The classic validity of (1) is easily obtainable, on the basis of the definitions of validity, truth, and ‘Act’: (1) is valid iff for any structure m and context $c \in C^m$ $\text{True}_c^m(1)$, that is,

$$\begin{array}{ll} \text{iff } [[(1)]]_{c,c}^m = 1 & \text{classic definition of truth} \\ \text{iff } [[\text{Act}\phi]]_{c,c}^m = [[\phi]]_{c,c}^m & \text{def. of } \leftrightarrow \\ \text{iff } [[\phi]]_{c,c}^m = [[\phi]]_{c,c}^m & \text{def. of ‘Act’} \end{array}$$

In contrast to a classic structure, an SR-structure includes a specification of a privileged circumstance of evaluation α . Then, (1) is valid iff for any structure m and context $c \in C^m$, $\text{True}_c^m(1)$, that is,

$$\begin{array}{ll} \text{iff } [[(1)]]_{c,\alpha}^m = 1 & \text{SR definition of truth} \\ \text{iff } [[\text{Act}\phi]]_{c,\alpha}^m = [[\phi]]_{c,\alpha}^m & \text{def. of } \leftrightarrow \\ \text{iff } [[\phi]]_{c,c}^m = [[\phi]]_{c,\alpha}^m & \text{def. of ‘Act’} \end{array}$$

Hence, as long as $[[\phi]]_{c,c}^m \neq [[\phi]]_{c,\alpha}^m$, not $\text{True}_c^m(1)$.

This result is of significance because, as Kaplan points out, the validity of (1) would provide a counter-instance to classic logical inference rules, in particular the rule of necessitation characteristic of normal modal logics.¹⁰

3.4.2 Scope Forcing

In the final sub-sections of this essay, we highlight some properties having to do with the relationships between indexical expressions and matters of scope. In this and the next sub-sections, we pause on the interactions between modal intensional operators and the modal indexical operator ‘actually’. In 4.4, we mention a conclusion pertaining to the relationships between quantifiers and ‘actually’.

Within a SR framework, the effect of occurrences of indexical operators such as ‘Act’ within the scope of intensional operators such as ‘Nec’ does not engender a reading in which the expressions under the effect of the former ‘take wide scope’ with respect to the latter. Consider for instance the equivalence

(2) $\text{Nec } (\phi \ \& \ \text{Act } \psi) \leftrightarrow \psi \ \& \ \text{Nec } \phi$

This biconditional is classically valid iff for all m and $c \in C^m$, $\text{True}_c^m(2)$, that is

$$\begin{array}{ll} \text{iff: } [[(2)]]_{c,c}^m = 1 & (\text{classic def. of truth}) \\ \text{iff: for all } w, [[\phi \ \& \ \text{Act } \psi]]_{c,w}^m = 1 & \\ \text{iff } [[\psi \ \& \ \text{Nec } \phi]]_{c,c}^m = 1 & (\text{def. of } \leftrightarrow \text{ and Nec}) \end{array}$$

¹⁰ See Kaplan 1977: 511, n. 35; 522, and 539 n. 65. Intriguingly, Kaplan’s independent alleged counter-examples to necessitation are provided by examples such as ‘I am here now’, whose special status is arguably even less controversially independent of matters of meaning, and derivable from ‘the vagaries of action’ (see Predelli 2005).

$$\begin{aligned}
 &\text{iff for all } w, [[\varphi]]^m_{c,w} = [[\psi]]^m_{c,c} = 1 \\
 &\text{iff } [[\psi]]^m_{c,c} = [[\text{Nec } \varphi]]^m_{c,c} = 1 \quad (\text{def. of } \& \text{ and Act}) \\
 &\text{iff for all } w, [[\varphi]]^m_{c,w} = [[\psi]]^m_{c,c} = 1 \\
 &\text{iff for all } w [[\psi]]^m_{c,c} = [[\varphi]]^m_{c,w} = 1 \quad (\text{def. of Nec})
 \end{aligned}$$

However, the equivalence in (2) is invalid from the point of view of SR:
 True^m_c (2)

$$\begin{aligned}
 &\text{iff: } [[(2)]]^m_{c,a} = 1 && \text{SR def. of truth} \\
 &\text{iff: for all } w, [[\varphi \& \text{Act } \psi]]^m_{c,w} = 1 \\
 &\text{iff } [[\psi \& \text{Nec } \phi]]^m_{c,a} = 1 && \text{def. of } \leftrightarrow \text{ and Nec} \\
 &\text{iff for all } w, [[\varphi]]^m_{c,w} = [[\psi]]^m_{c,c} = 1 \\
 &\text{iff } [[\psi]]^m_{c,a} = [[\text{Nec } \varphi]]^m_{c,a} = 1 && \text{def. of } \& \text{ and Act} \\
 &\text{iff for all } w, [[\varphi]]^m_{c,w} = [[\psi]]^m_{c,c} = 1 \\
 &\text{iff for all } w [[\psi]]^m_{c,a} = [[\varphi]]^m_{c,w} = 1 && \text{def of Nec}
 \end{aligned}$$

Hence, as long as $[[\psi]]^m_{c,c} \neq [[\psi]]^m_{c,a}$, not True^m_c (2).

3.4.3 Necessity and Actual Necessity

The formula

$$(3) \quad \text{Nec } \varphi \leftrightarrow \text{ActNec } \varphi$$

is classically valid, since True^m_c (3)

$$\begin{aligned}
 &\text{iff } [[(3)]]^m_{c,c} = 1 && \text{classic def. of truth} \\
 &\text{iff } [[\text{Nec } \varphi]]^m_{c,c} = 1 \text{ iff } [[\text{Act Nec } \varphi]]^m_{c,c} = 1 && \text{def. of } \leftrightarrow \\
 &\text{iff: } [[\text{Nec } \varphi]]^m_{c,c} = 1 \text{ iff } [[\text{Nec } \varphi]]^m_{c,c} = 1 && \text{def. of Act.}
 \end{aligned}$$

Yet, (3) is SR-invalid, since True^m_c (3)

$$\begin{aligned}
 &\text{iff } [[(3)]]^m_{c,a} = 1 && \text{SR-def. of truth} \\
 &\text{iff } [[\text{Nec } \varphi]]^m_{c,a} = 1 \text{ iff } [[\text{Act Nec } \varphi]]^m_{c,a} = 1 && \text{def. of } \leftrightarrow \\
 &\text{iff } [[\text{Nec } \varphi]]^m_{c,a} = 1 \text{ iff } [[\text{Nec } \varphi]]^m_{c,c} = 1 && \text{def. of Act.} \\
 &\text{iff for all } w \text{ such that } wR^m_a [[\varphi]]^m_{c,w} = 1 \\
 &\text{iff for all } w \text{ such that } wR^m_c [[\varphi]]^m_{c,w} = 1 && \text{def. of Nec.}
 \end{aligned}$$

Hence, if $[[\varphi]]^m_{c,w} = 1$ for all w accessible from a , but not for all w accessible from c , (4) is not True^m_c .

3.4.4 Universality and Actual Universality

The formula

$$(4) \quad \forall x (\text{Act } Fx) \leftrightarrow \text{Act } \forall x Fx$$

is classically valid. We write $g =_{x,w} f$ for: g is an assignment of values to variables just like f except at most for assigning to x an individual from $p^m(w)$, that is,

except at most for assigning to x an individual in the domain assigned in the structure to w . Then, $\mathbf{True}^m_c$ (4)

iff: $[[(4)]]^m_{f,c,c} = 1$	classic truth
iff: $[[\forall x (\text{Act } Fx)]]^m_{f,c,c} = 1$	
iff: $[[\text{Act } \forall x Fx]]^m_{f,c,c} = 1$	def. of $\leftrightarrow$
iff: for all $g =_{x,c} f$, $[[\text{Act } Fx]]^m_{g,c,c} = 1$	
iff: $[[\forall x Fx]]^m_{f,c,c} = 1$	def. of $\forall$ and Act.
iff: for all $g =_{x,c} f$ $[[Fx]]^m_{g,c,c} = 1$	
iff for all $g =_{x,c} f$ $[[Fx]]^m_{g,c,c} = 1$	def. of Act and $\forall$

But (4) is SR invalid, since $\mathbf{True}^m_c$ (4)

iff: $[[(4)]]^m_{f,c,a} = 1$	SR truth
iff: $[[\forall x (\text{Act } Fx)]]^m_{f,c,a} = 1$	
iff: $[[\text{Act } \forall x Fx]]^m_{f,c,a} = 1$	def. of $\leftrightarrow$
iff: for all $g =_{x,a} f$ $[[\text{Act } Fx]]^m_{g,c,a} = 1$	
iff: $[[\forall x Fx]]^m_{f,c,c} = 1$	def. of $\forall$ and Act.
iff: for all $g =_{x,a} f$ $[[Fx]]^m_{g,c,c} = 1$	
iff for all $g =_{x,c} f$ $[[Fx]]^m_{g,c,c} = 1$	def. of Act and $\forall$

Hence, if for all $i \in p^m(a)$ $i \in [[F]]^m_{g,c,c}$ but for some $j \in p^m(c)$, $j \notin [[F]]^m_{g,c,c}$, not $\mathbf{True}^m_c$ (4).

References

- Castañeda, Hector-Neri (1957) 'Some Nonformal "Logical" Relations', *Philosophical Studies* 8: 89–92.
- Kaplan, David (1977) 'Demonstratives', in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan*. Oxford: Oxford University Press, 1989.
- (1989) 'Afterthoughts', in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan*, Oxford: Oxford University Press.
- Kölbel, Max (2002) *Truth without Objectivity*, London: Routledge.
- Kripke, Saul (1963) 'Semantical Considerations on Modal Logic'. *Acta Philosophica Fennica*, 16: 83–94.
- Lewis, David (1980) 'Index, Context, and Content', in S. Kanger and S. Öhman (eds.), *Philosophy and Grammar*. Dordrecht: Reidel. Repr. in D. Lewis, *Papers in Philosophical Logic*, Cambridge: Cambridge University Press, 1998.
- MacFarlane, John (2003) 'Future Contingents and Relative Truth'. *Philosophical Quarterly* 53: 321–36.
- Egan, Andy, John Hawthorne, and Brian Weatherson (2005) 'Epistemic Modals in Context', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy. Knowledge, Meaning, and Truth*. Oxford: Oxford University Press.

- Predelli, Stefano (2005) *Contexts. Meaning, Truth, and the Use of Language*. Oxford: Oxford University Press.
- (1998a) 'Utterance, Interpretation, and the Logic of Indexicals'. *Mind and Language* 13.3: 400–14.
- (1998b) 'I Am Not Here Now'. *Analysis*, 58: 107–15.
- Stojanovic, Isidora (2005) 'A Different Story About Indexicals'. Amsterdam: ILLC Research Report pp 2005–05
- (forthcoming) 'Talking About Taste: Disagreement, Implicit Arguments and Relative Truth'. *Linguistics and Philosophy*.

This page intentionally left blank

4

Truth in the Garden of Forking Paths

John MacFarlane

This network of times which approached one another, forked, broke off, or were unaware of one another for centuries, embraces *all* possibilities of time. . . . In the present one, which a favorable fate has granted me, you have arrived at my house; in another, while crossing the garden, you found me dead; in still another, I utter these same words, but I am a mistake, a ghost. . . . Time forks perpetually toward innumerable futures.

Jorge Luis Borges, “The Garden of Forking Paths”

4.1 INTRODUCTION

It is very natural to think of time as a *line* running from past to future. Talk about the past concerns the part of the line that lies behind us, while talk of the future concerns the part that lies ahead. In his story “The Garden of Forking Paths,” Borges invites us to reject this linear picture and think of time as *forking* or *branching*. The branches represent possible future continuations of history. At each moment in time there are many possible continuations—many branches—and none of them is marked out as “the” future. The world itself branches.

Unlike the linear picture, the branching picture represents the future as genuinely, objectively “open”—as contingent in a particularly strong sense. The future might be contingent in a number of weaker senses without being genuinely “open.” For example, it might be that even though there is no branching, the evolution of the world is so chaotic that it could not be predicted, even by an agent with full knowledge of its past states and all the laws of nature. Whether the future is open in a stronger sense that requires the branching picture is in part an empirical question. To see “branching worlds” taken seriously as a description of

This paper was first presented at the LOGOS conference on Relativizing Utterance Truth in Barcelona, September 5–7, 2005. I am grateful to Richard Dietz, Iris Einheuser, Daniel López de Sa, and Max Kölbel for comments on an earlier draft, and to my Spring 2005 graduate seminar at Berkeley for helping me work through these ideas. Nuel Belnap deserves credit for getting me to think seriously about semantics in an indeterministic framework.

physical reality, one has only to look at the large literature on the “many worlds” interpretation of quantum mechanics.

But can we really make sense of branching worlds? On the branching picture, it seems, there is no such thing as *the* future. But we make claims about the future all the time. For example, I said ten days ago that it would be sunny today. It is sunny, so, it seems, what I said was true. But how *could* my claim have been true if, as the branching picture has it, there were both rainy and cloudy branches ahead of me when I made it?

This is the puzzle I want to discuss. Faced with the same puzzle, David Lewis concludes that we must either reject branching or accept that our future-directed talk and attitudes are fundamentally confused. (He prefers the first option.¹) I think the dilemma is a false one. In what follows, I will show that it is possible to give a semantic account of tensed talk that makes sense of our talk about the future, whether or not there is branching.

The core idea comes from my paper “Future Contingents and Relative Truth” (MacFarlane, 2003), but I hope here to correct a number of shortcomings in that paper:

1. In FCRT, I followed Prior, Thomason, Kaplan, and Belnap in treating temporal modifiers as sentential operators. This approach has now fallen out of favor among semanticists of natural language (see King (2003)). Here, I will use a more orthodox semantic framework, treating temporal modifiers as referring terms and quantifiers rather than operators.
2. In FCRT, I talked only of sentences and utterances, not of propositions. This made my case against supervaluationism easier than it should have been. In what follows, I’ll show how to deal with propositions and propositional truth in a branching framework, and I’ll explain how the supervaluationist might evade my criticism regarding retrospective assessments of future-tensed assertions by talking of proposition truth rather than utterance truth.
3. In FCRT, I did not explain how to do semantics for a one-place truth predicate in the object language. Nor did I explain how a relativist might handle an actuality operator. Both of these failings are corrected below. Indeed, it will emerge that the real problem with standard supervaluationism is its inability to make good sense of “actually,” not its treatment of retrospective assessments of predictions.

¹ See Lewis (1986: 199–209). Lewis’s main argument against branching worlds is that it “conflicts with our ordinary presupposition that we have a single future.” Lewis makes a distinction, which I do not entirely understand, between branching worlds and branching *within* a world. He concedes the possibility of the latter, though he thinks we can dismiss on “common sense” grounds the suggestion that “we ourselves are involved in branching” (209). Belnap *et al.* (2001: 205) rightly point out that a similar argument from common sense could be used against the claim that there are no reference-frame-independent facts about simultaneity.

Telling this story will require some careful concept-mongering. But with the appropriate concepts in hand, we can give a compelling semantic account of our talk of futurity and actuality that makes sense in both branching and non-branching frameworks. This should suffice to defuse arguments against branching that assume the incompatibility of branching with our ordinary talk and thought about the future.

4.2 A STANDARD FRAMEWORK

Let's start with a basic semantic framework, as "vanilla" as possible. We can then see just what changes are necessary to accommodate branching.

Semantics in the style of Lewis (1980) and Kaplan (1989) requires relativization of truth to

- an *assignment* (to handle quantifiers),
- an *index* (to handle operators), and
- a *context* (to handle indexicals).

An *assignment* is just a function from variables to objects in the domain. (To keep things simple, we will operate with a single, fixed domain of objects.) An *index* is a collection of separately shiftable parameters. Which parameters are required depends on what operators the language contains. If the language contains modal operators, a possible world parameter is needed; if the language contains temporal operators, a time parameter is needed; and so on. We'll work with a language that does not contain temporal operators, so for our purposes, an index can be simply a possible world. A *context* is a possible occasion on which a sentence might be used (or a representation of such an occasion). A context must be capable of supplying semantic values for indexicals and other context-sensitive expressions, but it need not be thought of as a sequence of parameters, as in Kaplan (1989).

We will call a $\langle \text{context, index, assignment} \rangle$ triple a "point of evaluation." A semantic theory for a language consists in a recursive definition of truth at a point of evaluation for arbitrary formulas of the language. Such a theory might include clauses like the following:

- (1) $\ulcorner \forall a \Phi \urcorner$ is true at $\langle C, w, a \rangle$ iff Φ is true at every point $\langle C, w, a' \rangle$ such that a' differs from a at most in the value it assigns to a .
- (2) $\ulcorner \Box \Phi \urcorner$ is true at $\langle C, w, a \rangle$ iff Φ is true at every point $\langle C, w', a \rangle$ such that w' is accessible from w .

When we are considering a sentence (closed formula) used in a particular context, we can strip off these relativizations and get an "absolute" truth value:

- (3) An occurrence of a sentence S in a context C is true iff S is true at $\langle C, w_C, a \rangle$, where w_C is the world of the context C and a is an arbitrary assignment. (Cf. Kaplan (1989: 522, 547).)

I will sometimes write “ S is true relative to C ” or “ S is true at C ” instead of “An occurrence of S in C is true.” All these formulations should be regarded as equivalent.²

Finally, logical truth and consequence are defined, following Kaplan (1989: 522–3), as truth and truth preservation at every context of use:³

- (4) A sentence S is logically true iff for every context C , S is true at C .
- (5) A sentence S is a logical consequence of a set Γ of sentences iff for every context C , if every member of Γ is true at C , then S is true at C .

It is also useful to define a notion of logical necessity that quantifies over points of evaluation rather than contexts, and a corresponding notion of logical implication:⁴

- (6) A formula Φ is logically necessary iff for every point of evaluation π , Φ is true at π .
- (7) A formula Φ is logically implied by a set Γ of formulas iff for every point of evaluation π , if every member of Γ is true at π , then Φ is true at π .

(Note that these concepts, unlike logical truth and consequence, are defined for formulas in general, not just for sentences.) In many familiar semantic frameworks, a sentence is logically true iff it is logically necessary, and a sentence is a logical consequence of a set of sentences iff it is logically implied by this set of sentences. But in languages containing certain kinds of context-sensitive expressions (particularly “now” and “actually”), these notions come apart, and it is possible for a sentence to be logically true without being logically necessary, or to be a logical consequence of a set of sentences without being logically implied by them. We will see a bit later why it is important to have both sets of concepts.

4.3 BRANCHING WORLDS

What changes do we need to make to this framework to make room for worlds that overlap in their initial temporal segments and branch towards the future? Surprisingly few. We can continue to think of indices as worlds, and the semantic clauses for temporal modifiers and indexicals like “tomorrow” can remain exactly as they were.

² Note that if the language contained no indexicals, so that C had no role to play in the definition of truth at a point of evaluation, C could be omitted from the points of evaluation. It would still be needed in (3), however, to “initialize” the world index. Cf. Kaplan (1989: 594), Belnap *et al.* (2001: 148).

³ An additional quantification over models or interpretations is required when—and *only* when—the sentences involved contain schematic “nonlogical constants.” All interpreted expressions are treated here as “logical constants.”

⁴ On the need for these two distinct notions, see Thomason (1970: 273) and Kaplan (1989: 548–50).

We'll need some new operators to express claims about the branching structure of historical possibility from within the object language:

- (8) $\lceil \text{Sett}_t : \Phi \rceil$ is true at $\langle C, w, a \rangle$ iff Φ is true at all points $\langle C, w', a \rangle$ such that w' overlaps with w at t .
- (9) $\lceil \text{Poss}_t : \Phi \rceil$ is true at $\langle C, w, a \rangle$ iff Φ is true at some point $\langle C, w', a \rangle$ such that w' overlaps with w at t .⁵

These operators are called *historical* modalities because they are indexed to a time. What is settled (or inevitable) at one time may not have been settled at an earlier time, and what is still an open possibility at one time might not be later. Thus, speakers of our language can talk about the branching structure of temporal possibility by saying things like

- (10) $\neg \text{Sett}_{\text{yesterday}} : \text{Sunny}(\text{today}) \wedge \neg \text{Sett}_{\text{yesterday}} : \neg \text{Sunny}(\text{today})$.

It was not settled yesterday that it would be sunny today, and it was not settled yesterday that it would not be sunny today.

The only thing that *must* change in our framework when we add branching is the reference to “the world of the context” in the definition of truth for an occurrence of a sentence at a context (3). When two worlds overlap at the time of utterance, an utterance that takes place in one must take place in the other as well. So contexts, construed as possible utterance occasions, will in general be located at multiple worlds that overlap at the time of utterance. There will not, in general, be a unique “world of the context” (Belnap *et al.*, 2001: 231–3).

Understanding this point is crucial to understanding the branching framework. It is tempting to suppose that, although the utterance takes place in many worlds or “possible histories,” they are not all on a par: one of them is marked out as “the actual history,” the one that will really come to pass. But there is no sense to be made of this idea. “Actual” and “actually,” as Lewis (1970a) argued, are indexical notions: their function is to shift the world of evaluation back to the world of the context.⁶ So we cannot appeal to the notion of actuality to single out one of the many worlds to which the context equally belongs.

What gives plausibility to the idea that one of the worlds is privileged, I think, is our tendency to picture the branching tree of histories as a branching network

⁵ Note that one can define epistemic analogues of these modalities in a framework without genuine branching, by appealing to exact qualitative similarity of worlds up through t instead of literal overlap at t . (See Lewis (1986: 206) on “divergence.”) In this sense, to say that it is now possible that there will be a sea battle tomorrow is to say that some other possible world with a past qualitatively identical to our past (and containing “counterparts” of us) holds a sea battle tomorrow. This position amounts to a subtle form of epistemicism: it is only because we cannot now distinguish our world from worlds with qualitatively identical pasts that the future is “open” for us.

⁶ This is a semantic point. It does not depend on Lewis’s modal realism. Whether we are realists or ersatzists about our worlds, we can ask whether a speaker at another world would have spoken truly in saying “Dodos are actually extinct.” It is unreasonable to suppose that the answer depends on whether there are dodos in *our* world.

of roads that we are traveling down (in the back of a pickup truck, perhaps—not in the driver's seat). We reason as follows:

Even though I'm now on *both* Route 66 and Interstate 40—they overlap here—there's a fact of the matter as to which one I'll be going down when they diverge ahead. I may not know at this point which one it will be, but I know I won't be going *both* ways! Similarly, even if I'm now located in many worlds that overlap in the present but diverge in the future, there's a fact of the matter as to which one will be my future. I'll find out when I get there.

But this picture embodies a fundamental confusion. We've already represented time as one of the spatial dimensions of our tree. So what could possibly be represented by the *motion* of a point on this tree? Certainly not a process that takes place in time, since all such processes are already represented spatially on the tree. There is nothing in the branching time model that corresponds to a car moving along the branching road, and nothing that corresponds to the decision the car will have to make to go down one branch or the other. If worlds branch, then we branch too.

4.4 SUPERVALUATIONS

What, then, should we do with our definition of sentence truth at a context, if we can no longer talk of “the world of the context”? The obvious solution is to quantify over all of the worlds to which the possible utterance event belongs:

- (11) An occurrence of a sentence S in a context C is true iff for each of the worlds w that overlap at C and for every assignment a , S is true at $\langle C, w, a \rangle$.

This gives us a supervaluational semantics along the lines of Thomason (1970). An occurrence of

- (12) Sunny(tomorrow)
It will be sunny tomorrow

is counted true at a context just in case it is sunny on the day after the day of the context in every world to which the context belongs. Thus, if it is not settled today that it will be sunny tomorrow, (12) is not true. But it is not false either—that is, its negation is not true. So supervaluationism is committed to truth value gaps for occurrences of sentences at contexts.⁷ This is not obviously an unwelcome result. Given that nothing in the branching picture warrants assigning one truth

⁷ Could the supervaluationist avoid saying this by changing the biconditional in (11) to a conditional? This alternative view would say that S is true at C if it is true at all worlds that overlap at C . But it wouldn't say that S is *not* true if it isn't true at all such worlds. In this case, the theory would just remain silent about the truth of S at C . The problem is that there are no grounds for remaining silent here. It is reasonable to withhold a verdict when the facts are still out, but that is not the case here. The theory will *still* be silent if it is given *all* the facts about the branching worlds

value as opposed to the other in such cases—the situation is completely symmetrical—one might reasonably think that future contingents *must* be “gappy.”

An attractive feature of supervaluational semantics is that it secures this result without giving up the law of excluded middle. Despite the fact that neither (12) nor its negation is logically true on the supervaluational semantics, the disjunction of (12) with its negation, that is,

- (13) $\text{Sunny}(\text{tomorrow}) \vee \neg \text{Sunny}(\text{tomorrow})$

It will be sunny tomorrow or it won't be sunny tomorrow

is logically true, just as it was in the linear framework (Thomason 1970: 270). Thus we can say things like

- (14) $[\text{Poss}_{\text{now}} : \text{Sunny}(\text{tomorrow}) \wedge \text{Poss}_{\text{now}} : \neg \text{Sunny}(\text{tomorrow})]$
 $\wedge [\text{Sunny}(\text{tomorrow}) \vee \neg \text{Sunny}(\text{tomorrow})]$

Although it is now possible that it will be sunny tomorrow and possible that it will not be sunny tomorrow, either it will be sunny or it won't be

without contradicting ourselves. If we can consistently say this in a branching universe, the branching picture does not seem as antithetical to our ordinary ways of thinking about the future as Lewis and others have supposed.

It is sometimes objected that supervaluationism secures this nice result only by making “or” non-truth-functional, which is an intolerable distortion of its meaning. I think this charge can be resisted. Truth functionality should be defined in terms of truth at a point of evaluation, not truth at a context. To see this, it is only necessary to reflect that disjunction is just as truth-functional when it appears between two open formulas as it is when it appears between two sentences. But open formulas do not have truth values relative to contexts, for there is no such thing, in general, as “the assignment of the context” (Belnap *et al.*, 2001: 233–4). (That is why we *quantify* over assignments in our definition of sentence truth at a context.) So the *only* sense in which disjunction can be said to be truth-functional is this one:

- (15) A binary connective $\star$ is *truth-functional* iff for every point of evaluation π and all formulas Φ, Ψ , the truth value of $\ulcorner \Phi \star \Psi \urcorner$ at π is a function of the truth values of Φ and of Ψ at π .

The supervaluationist does not deny that disjunction is truth-functional in *this* sense. Perhaps there is *something* objectionable about allowing the disjunction of two sentences to be true at a context even though neither disjunct is true

that could possibly be relevant to the truth of S at C . (And please don't say that there are still facts to be learned—facts about which branch will be actualized. See above on the moving car picture.) Another option (that of Belnap and Green (1994) and Belnap *et al.* (2001)) is to decline to assign truth values to sentences at contexts at all. Whether this option is viable depends on whether the theoretical work done by the notion of sentence truth at a context can be done in other ways. For a brief discussion, see MacFarlane (2003: 331).

at that context, but it is not that doing so requires taking disjunction to be non-truth-functional. Indeed, the supervaluationist can use the very same recursive clause for disjunction as the classical logician:

- (16) $\lceil \Phi \vee \Psi \rceil$ is true at $\langle C, w, a \rangle$ iff Φ is true at $\langle C, w, a \rangle$ or Ψ is true at $\langle C, w, a \rangle$.

Philosophers have also objected that supervaluational semantics has odd *logical* upshots—for instance, that $\lceil \text{Sett}_{\text{now}} : S \rceil$ is a logical consequence of S , even though $\lceil S \supset \text{Sett}_{\text{now}} : S \rceil$ is not a logical truth. This is said to be a counterexample to the classical rule of Conditional Proof. For similar reasons, it is alleged, we must reject classical rules for Case Argument, Reductio, and Contraposition.⁸ As Williamson puts the point, “supervaluations invalidate our natural mode of deductive thinking” (Williamson, 1994: 152).

But this criticism, like the charge that supervaluationism makes disjunction non-truth-functional, seems to me misguided.⁹ Conditional Proof is a proof rule, not a semantic claim. Moreover, it is a peculiar kind of proof rule, since it makes reference to a *subproof*: it says that a conditional may be deduced from a subproof that begins with the antecedent and ends with the consequent. (Significantly, all of the rules deemed difficult for the supervaluationist are of this kind.) So the supervaluationist will only have trouble with Conditional Proof if she takes the move from S to $\lceil \text{Sett}_{\text{now}} : S \rceil$ to be legitimate in a subproof. But why must she do that? Whether a deductive system is sound and complete has to do with what can be *proved*, not what can be subproved. No general logical or methodological principle bars us from using a proof system that licenses certain moves only in the top level of a proof (that is, only on lines that do not depend on any hypotheses). Doing this would have its costs, of course—we could no longer simply lift proofs into subproof contexts, without checking to see if they use the special rules—but it is often the case that the addition of expressive power to a language forces us to complicate its proof theory. (Think of the extra complexities involved in moving from natural deductions in propositional logic to natural deductions in first-order logic. Indeed, even in standard first-order deduction systems, proofs cannot be indiscriminately lifted into subproofs, since flagging restrictions may be violated unless variables or individual constants are relettered.)

It may be objected that even if the supervaluationist can retain the Conditional Proof rule by restricting the inference rules that can be used in subproof contexts, such a restriction would be ad hoc and unmotivated. For surely

- (17) If B can be obtained from A by moves that are guaranteed to preserve truth, $\lceil A \supset B \rceil$ is guaranteed to be true.

⁸ See Machina (1976: 52–3) or Williamson (1994: 151–2), substituting “ $\text{Sett}_{\text{now}} :$ ” for “Definitely.”

⁹ For kindred reflections, see McGee and McLaughlin (1998: 224–5 and 2004: 132–6.)

What principled reason could there be, then, to disallow the use of a truth-preserving rule in a subproof for Conditional Proof?

But (17) is ambiguous. By “guaranteed to preserve truth,” one might mean either “guaranteed to preserve truth at a context of use” or “guaranteed to preserve truth at a point of evaluation.” So (17) might mean either

- (18) If B can be obtained from A by moves that preserve truth at every context (i.e., by logical consequences), $\lceil A \supset B \rceil$ must be true at every context (i.e., logically true).

or

- (19) If B can be obtained from A by moves that preserve truth at every point of evaluation (i.e., by logical implications), $\lceil A \supset B \rceil$ must be true at every point of evaluation (i.e., logically necessary).

The supervaluationist can perfectly well accept (19). Although the inference from S to $\lceil \text{Sett}_{\text{now}} : S \rceil$ preserves truth at every context, it does not preserve truth at every point of evaluation, and so is no counterexample to the plausible principle (17), on one natural reading.

I conclude, then, that the logical and semantic “anomalies” of supervaluationism are not so anomalous when one looks into them, and certainly no reason to reject the view. For that we must look elsewhere.

4.5 RETROSPECTIVE ASSESSMENTS

In FCRT, I rejected the supervaluational semantics on the grounds that it gave incorrect verdicts for *retrospective* assessments of claims about the future. Suppose that yesterday I said, “It will be sunny tomorrow,” and suppose that at the time of my utterance both sun and clouds were open possibilities for today. According to supervaluationism, then, my utterance was not true. By (11), the sentence I uttered was neither true nor false at the context in which I uttered it. But surely that is the wrong verdict. I said that it would be sunny today, and look—it is sunny! How could it be, then, that what I said was not true?

To see how strange the supervaluationist’s verdict is, suppose that the Director of the Bureau of Quantum Weather Prediction now offers me an irrefutable proof that, at the time of my utterance yesterday, it was still an open possibility that it would not be sunny today. Would such a proof compel me to withdraw my assertion? Hardly. If I had asserted that it was *settled* that it would be sunny today, I would have to stand corrected. But I did not assert that. I just said that it *would be* sunny—and it is. My prediction was true, as we can demonstrate simply by looking outside.

But suppose that the Director had visited me yesterday, just after I made my assertion, and confronted me with exactly the same facts. Wouldn’t I have had

to acknowledge that my claim was untrue? For it would have been arbitrary to understand what I said to concern any particular one of the many worlds I occupied. (It is useless, recall, to appeal to the “actual world” in this context.) By showing that some of those overlapping worlds contained a sunny tomorrow, while others did not, the Director would have shown that there was no objective basis for calling my utterance true rather than false. I would have had to agree with the supervaluationist’s verdict.

When we consider the claim just after it was made, then, the Director’s proof seems to show that it was not true. But when we consider it from a different vantage point—from the vantage point of today—the Director’s proof seems altogether irrelevant to its truth. Of course, *justification* is well known to exhibit this kind of perspectival variation. The very same considerations that count as solid grounds for belief in one context may be utterly insufficient in another, because new evidence has become available. But that’s not what we’re dealing with here. It’s not that my new evidence that it’s sunny today somehow undermines the Director’s proof that it was not settled at the time of my utterance that it would be sunny today. No, the Director’s proof still stands. I still accept its conclusion. It’s just that, from my current point of view, this proof isn’t at all relevant to the truth of my claim.

This, then, is the puzzle:

- present claims concerning the future can be shown to be untrue by a proof of present unsettledness, but
- past claims concerning the present cannot be shown to have been untrue by a proof of past unsettledness.

4.6 ASSESSMENT SENSITIVITY

I suggested in FCRT that this odd perspective relativity was both the source of the staying power of the problem of future contingents and the key to its solution. The source of its staying power, because if one focuses on present claims concerning the future, it seems impossible that they could have a determinate truth value, while if one focuses on past claims concerning the present, it seems impossible that they could fail to have one. The key to its solution, because one has only to find a way to give both perspectives their due. The way to do that, I argued, is to relativize the truth of sentences not just to a context of use but to a context of *assessment*. One can then say that my assertion was true as assessed from the context I occupy today, neither true nor false as assessed from the context I occupied yesterday, and false as assessed from the (quite rainy) context I might have occupied today if the winds had blown a different direction.

I should emphasize that the qualifiers “of assessment” and “of use” relate to the *use* to which a context is being put semantically. They designate roles, not

kinds of context. The very same context can be considered as context of use or as context of assessment. These roles relate to a situation in which we are assessing an utterance (actual or possible) for truth or falsity. The context of use is the context in which the utterance was made. The context of assessment is the context we occupy in assessing it.

From a technical point of view, the addition of contexts of assessment requires only minimal changes to our semantic theory. A point of evaluation can remain a $\langle \text{context of use, index, assignment} \rangle$ triple. Semantic clauses for expressions of the language, including the historical modalities, can also remain the same. The only real difference between the “relativist” semantics and the supervaluational semantics comes in the definition of truth for an occurrence of a sentence. We must now relativize sentence truth to *two* contexts (of use and assessment). Whereas on the supervaluational semantics, we look at all the worlds that overlap at the context of use, on the relativist semantics, we look only at those that *also* overlap at the context of assessment—provided that context can be reached from the context of use by going forward in time along one of the branches. (If it cannot be so reached, either because it is in the past of the context of use or because it is on another branch of the tree entirely, then we just look at the worlds overlapping at the context of use.)

In stating this definition formally, it will help to have some notation:

(20) Where C is a context, $W(C) =$ the set of worlds that overlap at C .

(21) Where C_1 and C_2 are contexts,

$$W(C_1|C_2) = \begin{cases} W(C_1) & \text{if } W(C_1) \cap W(C_2) = \emptyset \\ W(C_1) \cap W(C_2) & \text{otherwise} \end{cases}$$

We can then say that

(22) An occurrence of a sentence S at C_U is true as assessed from C_A iff for every world $w \in W(C_U|C_A)$, S is true at $\langle C_U, w, a \rangle$.¹⁰

This definition gives results that accord with the intuitions elicited above. Let us consider utterances of two sentences,

(23) Sunny(tomorrow)

It will be sunny tomorrow.

(24) $Sett_{\text{now}}$: Sunny(tomorrow)

It is now settled that it will be sunny tomorrow

both made by me, yesterday. Let C_1 be the context in which I uttered them and C_2 the context from which I am now assessing these utterances. Suppose that C_1

¹⁰ I will use “ S is true at (or relative to) context of use C_U and context of assessment C_A ” interchangeably with “An occurrence of S at C_U is true as assessed from C_A .”




Figure 4.1.

belongs to precisely four possible worlds, w_1 , w_2 , w_3 , and w_4 . Due to branching, C_2 belongs to just two of them, w_3 and w_4 . Suppose finally that in w_1 it is rainy today, while in w_2 , w_3 , and w_4 it is sunny today. (See Figure. 4.1)

We then get the following (ignoring the assignment parameter of points of evaluation, since we are dealing with sentences):

$$W(C_1|C_1) = W(C_1) = \{w_1, w_2, w_3, w_4\} \quad W(C_1|C_2) = W(C_2) = \{w_3, w_4\}$$

Truth value . . . at points of evaluation . . . at contexts (use/assessment)

S	$\langle C_1, w_1 \rangle$	$\langle C_1, w_2 \rangle$	$\langle C_1, w_3 \rangle$	$\langle C_1, w_4 \rangle$	C_1/C_1	C_1/C_2
(23)	False	True	True	True	Neither	True
(24)	False	False	False	False	False	False

Notice that (23) is *assessment-sensitive*—its truth value varies with the context of assessment—while (24) is not.¹¹

Finally, we need to define logical truth and consequence. (The definitions of logical necessity and implication can remain the same as before.) Instead of quantifying over just contexts of use, we quantify over both contexts of use and contexts of assessment:

- (25) A sentence S is a logical truth iff for every pair of contexts C_1 and C_2 , S is true relative to context of use C_1 and context of assessment C_2 .

¹¹ For this terminology, see MacFarlane (2005).

- (26) A sentence S is a logical consequence of a set Γ of sentences iff for every pair of contexts C_1 and C_2 , if every member of Γ is true relative to context of use C_1 and context of assessment C_2 , then S is true relative to context of use C_1 and context of assessment C_2 .

One consequence of this change is that $\ulcorner \text{Sett}_{\text{now}} : S \urcorner$ is no longer a logical consequence of S . We've just seen a counterexample: at context of use C_1 and context of assessment C_2 , (23) is true but (24) is false.

4.7 UTTERANCE TRUTH AND PROPOSITION TRUTH

The argument in FCRT for taking future contingents to be assessment-sensitive rested entirely on the retrospective assessment argument I presented in the last two sections. In the next three sections, I will explain how the supervaluationist might deal with retrospective assessments. I will then consider a different sort of argument for the assessment sensitivity of future contingents.

Consider two different ways of presenting the argument from retrospective assessments:

- (27) Yesterday I uttered the sentence "It will be sunny tomorrow."
 It is sunny today.
 So my utterance was true.
- (28) Yesterday I asserted that it would be sunny today.
 It is sunny today.
 So what I asserted was true.

The difference is that in (27), truth is predicated of an utterance of a sentence, while in (28), truth is predicated of "what I asserted"—a proposition.

Which of these arguments did I intend in FCRT? That is not at all clear. In the technical part of the paper, I talked of sentences and utterances, not propositions. Officially, then, the argument ought to have been something like (27). But this is not what one finds in the (less formal) part of the paper where the argument is first presented:

- (29) Jake asserted yesterday that there would be a sea battle today.
 There is a sea battle today.
 So Jake's assertion was true. (325)

Here Jake's speech act is described using indirect discourse, as in (28), by way of its propositional content rather than its sentential vehicle.¹²

¹² "Jake's assertion" in (29) denotes what Jake asserted, not Jake's act of asserting it. Although the word "assertion" can be used to refer either to an *act* of asserting or to the content of such an act, it is doubtful that we ever predicate truth of *acts* at all, even if they are speech acts (MacFarlane,

I think there is a reason I slipped into proposition talk in giving the retrospective assessment argument, despite my efforts to avoid it elsewhere. I was trying to elicit the intuition that the retrospective assessment of Jake's prediction as true was a natural one—something no ordinary person would reject. And in ordinary speech, truth and falsity are almost invariably predicated of *propositions*, as in the following:

- (30) What he said is false.
- (31) Nothing George asserted in his talk is true.
- (32) I know you believe he's dishonest, but that's false.
- (33) It's true that it has been a hot summer.
- (34) That was a false claim.¹³

Aside from a few, relatively isolated examples, like

- (35) A truer sentence was never spoken,

people do not apply the predicates "true" and "false" to sentences or utterances, except in areas of philosophical incursion. It surprised me a bit when I realized this, because it is very common in *philosophical* prose to predicate "true" and "false" of sentences and utterances. These uses, however, must be understood as technical.

As we have seen, supervaluationism gives the "wrong" retrospective assessments of truth for past utterances of future contingents. But if I am right that utterance truth is a technical notion that plays no important role in our ordinary thought and talk, then the supervaluationist can accept these consequences without being revisionist about our ordinary future-directed talk. What really matters is whether supervaluationism can vindicate our retrospective assessments of the truth of *propositions*.

4.8 PROPOSITIONS

To explore this question, we need to introduce ways of talking about propositional truth, both in our semantic metalanguage and in the object language.

Propositions are often said to be the "primary bearers of truth value." What this means is that anything else that is true or false is so in virtue of expressing a proposition that is true or false. It does *not* mean that propositions have their truth values absolutely or intrinsically. In standard frameworks, propositions are

2005: 322–3). Think about how natural it sounds to paraphrase "Jake's assertion was true" as "What Jake asserted was true," and how unnatural it would sound to paraphrase it as "What Jake did in asserting that was true."

¹³ Here "claim" denotes the content—what is claimed—not the act of claiming. See n. 12, above.

taken to have truth values only relative to “circumstances of evaluation”—here, possible worlds.¹⁴ In our semantic metalanguage, then, we will have a two-place propositional truth predicate: “proposition p is true at circumstance of evaluation (or world) w .”

However, this two-place truth predicate is not what we need to evaluate ordinary retrospective assessments like “What I said yesterday was true.” “True” occurs here in the object language, as a *monadic* predicate of propositions. To make progress, we need to understand *its* semantics, which are simple and straightforward:¹⁵

- (36) “True” applies to x at a point of evaluation $\langle C, w, a \rangle$ iff x is a proposition and x is true at w .

(To avoid confusion with the various metalinguistic truth predicates in play, I’ll capitalize the monadic object-language truth predicate.)

This definition has two consequences that cohere well with our ordinary use of the truth predicate. First, it implies that every instance of the following *disquotational schema* is true at every point of evaluation (and so both logically necessary and logically true):

- (37) $\forall x(x = \text{the proposition that } S) \supset (\text{True}(x) \equiv S)$.¹⁶

Second, there is no argument place for a time in “True” as defined by (36). Thus, even if there is branching, so that worlds are eliminated as live possibilities as time passes, it is never correct to say things like

- (38) What you said yesterday is True today, but it wasn’t True yesterday.¹⁷

(This is so not only on the supervaluationist’s semantics, but on the relativist’s, provided the same definition of “True” is used.) Although we do use both “is True” and “was True” as predicates of propositions, the tense does not appear to have any independent semantic significance: it is determined, rather, by the grammatical context. When a past event is being discussed, as in

- (39) What you said yesterday was True,

we tend to use the past tense with the truth predicate. When a present event is being discussed, as in

¹⁴ Generally, circumstances of evaluation will contain the same parameters as the indices of points of evaluation. Thus, for example, Kaplan (1989) takes them to be world/time pairs.

¹⁵ Of course, once we introduce “True” into the object language, something needs to be done to ward off paradox. But this issue is orthogonal to the one we’re concerned with here, so for simplicity I’ll ignore it.

¹⁶ It is here assumed that if x is the proposition expressed by S at C , x is true at w iff S is true at $\langle C, w, a \rangle$ (for any assignment a).

¹⁷ I won’t deny that we can find a use for such a sentence—but not when “what you said yesterday” denotes a proposition and literal truth is at issue.

(40) What you are saying is True,

we use the present tense.

Let's use (36) to see what the supervaluationist should say about retrospective assessments of propositional truth. Yesterday I said that it would be sunny today. Today I say

(41) What I said yesterday was True.

Have I spoken truly? On the supervaluational account, I have spoken truly iff (41) is true at every point $\langle C, w, a \rangle$ such that C is the current context and $w \in W(C)$. By (36), (41) is true at a point $\langle C, w, a \rangle$ iff the proposition denoted at $\langle C, w, a \rangle$ by "what I said yesterday" is true at w . But for all assignments a and all worlds $w \in W(C)$, the denotation of "what I said yesterday" at $\langle C, w, a \rangle$ is the proposition that it would be sunny today. So I have spoken truly iff the proposition that it would be sunny today is true at all the worlds that overlap at the present context. But this proposition *is* true at all such worlds (just look outside). So according to the supervaluational semantics, I have spoken truly in retrospectively judging yesterday's claim to have been true.

Here is another way to see the point. (41) is true at a point of evaluation $\langle C, w, a \rangle$ where $w \in W(C)$ iff

(42) The proposition that it would be sunny today was True

is true at that point. Since (in this framework) propositions are not tensed and propositional truth is not time-relative, (42) is true at a point iff

(43) The proposition that it is sunny today is True

is true at that point. And given the disquotational property, (43) is true at a point iff

(44) It is sunny today

is true at that point. So the supervaluationist—or anyone else who employs a truth predicate defined by (36)—must assign the same truth value to (41) at C as she does to (44). Since the supervaluationist takes (44) to be true at C , she must say the same about the retrospective assessment (41).

Thus, when the retrospective assessments are thought of as judgements about the truth of a proposition—as I have argued they should be—they do not pose a problem for the supervaluational semantics.

4.9 DETERMINATE TRUTH

What should the supervaluationist say about *contemporary* assessments of contingent claims about the future? For example, what should I have said *yesterday*

about the truth of my prediction? One might expect the supervaluationist to say that yesterday, after the Director's visit, I could truly have uttered

(45) What I just said (when I said that it would be sunny tomorrow) is not True.

But the supervaluationist *can't* say this. Given the disquotational properties of "True," (45) is true at the context in question if and only if

(46) It will not be sunny tomorrow

is true at that context. Since on the supervaluationist account one cannot truly utter (46) or its negation when it is still unsettled whether it will be sunny the next day, it follows that one cannot truly utter (45) or its negation, either. Thus the semantic fact recorded in the metalanguage by the observation that neither (46) nor its negation is true at such a context is *ineffable* from the "internal" point of view. To express it, one must deploy the semanticist's technical notions of utterance truth or sentence truth relative to a context.

For those who do not think that a proof of unsettledness should compel withdrawal of an assertion about the future, this result might actually be welcome. From their "internal" points of view on the branching tree of histories, speakers will not be able to acknowledge proofs of unsettledness as grounds for asserting that what they said was not true—since these proofs would then also be grounds for asserting the negations of what they said.

But for those supervaluationists who *do* think that a proof of unsettledness should compel withdrawal of an assertion about the future, there is an easy solution. We can introduce a "determinate truth" predicate:

(47) "DetTrue" applies to x at a point of evaluation $\langle C, w, a \rangle$ iff x is a proposition and x is true at every world $w' \in W(C)$.

Using this predicate, our speakers can correctly characterize propositions whose truth is still unsettled as "not Determinately True." Whether they take a proof of unsettledness to compel withdrawal of an assertion about the future will then depend on whether they think retraction is required by a proof that the assertion is not *Determinately* True.

Like "True" and "False", "DetTrue" and "DetFalse" are not time-indexed. So it won't ever be correct to say that a proposition "is now Determinately True, but wasn't Determinately True yesterday." On the other hand, it may be that a proposition that can now be correctly called "Determinately True" could yesterday have been correctly called "not Determinately True." "DetTrue," unlike "True," is use-sensitive.¹⁸

It now appears that, contrary to what was alleged in FCRT, the supervaluationist *can* account for the asymmetry between contemporary and retrospective assessments of contingent claims about the future. She can acknowledge that I

¹⁸ "True," by contrast, is assessment-sensitive (for the terminology, see MacFarlane (2005)).

can now truly assert “What I said was true,” even though I couldn’t truly assert this yesterday. And she can acknowledge that I can now truly assert “What I said was determinately true,” even though yesterday I could have truly asserted “What I just said is not determinately true.” The only thing she can’t do is acknowledge today that my *utterance* yesterday was true. But the considerations raised in Section 7 suggest that this need not worry us much. So the argument from retrospective assessments is not sufficient to show that we need to take the novel step of relativizing sentence truth to contexts of assessment.

Does that mean that there is no need to appeal to contexts of assessment in semantics for branching time? No. As I will now argue, the supervaluationist does not have resources for a proper treatment of “actually.” It is this that requires us to relativize truth to contexts of assessments.

4.10 ACTUALITY

In standard (non-branching) frameworks, the actuality operator works as follows:

- (48) $\ulcorner \textit{Actually} : \Phi \urcorner$ is true at $\langle C, w, a \rangle$ iff Φ is true at $\langle C, w_C, a \rangle$, where w_C is the world of the context C . (See Lewis (1970b), Kaplan (1989: 545).)

No matter how deeply embedded we are, no matter how far the world of evaluation has been shifted, the actuality operator returns it to the world of the context of use. Of course, this only works if there *is* a unique world of the context of use—as there is not when worlds can overlap and branch.

(48) respects a plausible constraint on actuality operators, which I call *Initial Redundancy*:

- (49) An operator $\star$ is *initial-redundant* just in case for all sentences S , $\ulcorner \star S \urcorner$ is true at exactly the same contexts of use (and assessment) as S (equivalently: each is a logical consequence of the other).

If “*Actually* :” were not initial-redundant, it might sometimes happen that you could truly utter a sentence S , but not $\ulcorner \textit{Actually} : S \urcorner$ (or perhaps vice versa). But that does not seem to be possible. When you can truly say, “It will be sunny tomorrow,” you can truly say, “It will actually be sunny tomorrow,” and when you can truly say, “It will actually be sunny tomorrow,” you can truly say, “It will be sunny tomorrow.” This is not because “actually” has no effect on truth conditions, but because of a delicate relation between the semantics for “actually” and the definition of sentence truth at a context. The effect of adding an actuality operator to the *front* of a sentence is to shift the world of evaluation to the world of C . This has an effect on the sentence’s truth-at-points profile, but not on its truth-at-contexts profile, because (in standard, non-branching frameworks) a sentence is true at a context C just in case it is true at the point $\langle C, w_C \rangle$, where

w_C = the world of C . (Since we're dealing with sentences, I omit the assignment parameter here and in the following paragraph.)

In a branching framework, we no longer have a unique “world of the context of use,” so we need to generalize (48).¹⁹ Here we are tightly constrained by Initial Redundancy, which requires that for all sentences S and contexts C ,

(50) S is true at C iff $\ulcorner \textit{Actually} : S \urcorner$ is true at C .

To respect this constraint, the supervaluationist, who takes S to be true at C just in case it is true at $\langle C, w' \rangle$ for all $w' \in W(C)$, must say that

(51) $\ulcorner \textit{Actually} : \Phi \urcorner$ is true at $\langle C, w \rangle$ iff Φ is true at $\langle C, w' \rangle$ for all $w' \in W(C)$.²⁰

On the other hand, the relativist, who takes S to be true at context of use C_U and context of assessment C_A just in case it is true at $\langle C_U, w' \rangle$ for all $w' \in W(C_U|C_A)$, must define the actuality operator as follows:

(52) $\ulcorner \textit{Actually} : \Phi \urcorner$ is true at $\langle C_U, C_A, w \rangle$ iff Φ is true at $\langle C_U, C_A, w' \rangle$ for all $w' \in W(C_U|C_A)$.²¹

So the relativist's different definition of sentence truth at a context (or contexts) requires a substantially different definition of truth at a point of evaluation for the actuality operator.²²

To see why this matters, let us reflect on the upshots of the previous three sections. From an “external” point of view—the point of view of the semanticist assigning truth values to sentences relative to contexts of use and contexts of assessment—there is a clear difference between the supervaluational semantics of (11) and the relativist semantics of (22). But from an “internal” point of view—the perspective of speakers within the tree of worlds, making claims and

¹⁹ The first attempt to do this of which I am aware is Belnap *et al.* (2001: 246–7). Belnap defines two actuality operators. The first, “*Actually*₁,” is essentially (51), which is initial-redundant in supervaluational semantics but not in relativist semantics. The second, “*Actually*₂,” which is initial-redundant for the relativist as well, will be discussed briefly below (n. 20). Belnap is neither a supervaluationist nor a relativist; he declines to give truth values to sentences at contexts at all.

²⁰ Initial Redundancy would also be secured by the clause: “*Actually* : Φ ” is true at $\langle C, w \rangle$ iff Φ is true at $\langle C, w \rangle$. But this makes the actuality operator redundant in all contexts, not just initially. It may be that there is a use of “actually” in English that behaves this way—Lewis's (1970b) “non-shifty” use—but we're after an operator that makes a difference in embedded contexts. Belnap's “*Actually*₂” operator is a kind of hybrid between (51) and the redundant operator: it behaves like one or the other, depending on whether it is evaluated on a history that contains the moment of the context of use.

²¹ Note that the addition of an actuality operator has forced us to bring the context of assessment into our points of evaluation, where it was not needed before.

²² The introduction of an actuality operator brings with it “anomalies” similar to those discussed at the end of section 4: for example, “*Actually* : S ” is a logical consequence of S , though “ $S \supset \textit{Actually} : S$ ” is not a logical truth. But are these anomalies really unwelcome? After all, we designed our actuality operator precisely to ensure that “*Actually* : S ” was a logical consequence of S and vice versa (Initial Redundancy). If we want our language to contain an actuality operator, then, we are going to have to get over our qualms about these logical “anomalies.” Fortunately, these qualms are misplaced (see the end of Section 4).

commenting on the truth of other claims—this difference is effectively invisible. Provided that “actually” is not in the language, the supervaluationist and the relativist can use exactly the same recursive definition of truth at a point of evaluation (even for the predicate “True” itself). The differences between their positions arise only when we need to define truth at a context of use (and context of assessment). But these are rarified metatheoretic notions: the folk assess claims for truth and falsity from *within* the language, by saying things like “That’s false” or “What she said was true.” If we stick resolutely to this internal point of view, the supervaluational theory looks indistinguishable from the relativist theory.

In FCRT, I sought to connect the theoretical notion of truth relative to a context of use and context of assessment to the participants’ linguistic practice by giving it a role in a normative account of assertion. To assert a sentence, I said, is *inter alia* to commit oneself to responding to appropriate challenges, by providing grounds for the truth of the sentence relative to the original context of utterance and the context of assessment one occupies at the time of the challenge. This account gives “internal” significance to the difference between supervaluationism and relativism; if it is correct, then the two views would have different implications about precisely what one is committing oneself to in asserting that it will be sunny tomorrow. But is it correct? One might argue that if there are norms for assertion, they should be formulable from the participants’ internal perspective, using a monadic propositional truth predicate, not the semanticist’s doubly relativized truth predicate. For example: withdraw an assertion if what you asserted is shown to have been Untrue (or, alternatively, Determinately Untrue). And, as we have seen, these norms won’t distinguish between supervaluational and relativist semantics.

By adding an actuality operator to the language, however, we have dragged the “external” differences between the two views down to the “internal” level. In order to secure Initial Redundancy, the supervaluationist and the relativist must give different definitions of truth at a point of evaluation for the actuality operator. These differences will manifest themselves in different verdicts about the correctness of judgements of the form, “What she said was True.”

Take, for example, my claim yesterday that it would be sunny today—henceforth, my “first claim.” As we have seen, the supervaluationist can agree with the relativist that it is correct, today, to say “What I said yesterday was True,” even though it wouldn’t have been correct, yesterday, to say “What I just said is True.” But suppose that yesterday I also said

(53) *Actually* : Sunny(tomorrow)

It will actually be sunny tomorrow.

Call this my “second claim.” According to the supervaluationist, it should be correct for me to say (now) that my first claim was True and my second claim False. On the relativist’s account, on the other hand, I can correctly say that both

claims were True. For on the relativist's account, "actually" tells us to look at the worlds that overlap at the context of *assessment*, not the context of use, and at all such worlds, it is sunny today. The relativist's verdict here more closely matches ordinary usage.

Suppose that I also made a third claim yesterday:

- (54) *Actually* : Sunny(tomorrow) $\vee$ *Actually* : Cloudy(tomorrow).

Either it will actually be sunny tomorrow or it will actually be cloudy tomorrow.

The supervaluationist and the relativist agree that *yesterday* I could not have correctly described this claim as "True." However, they differ about the *retrospective* assessment. According to the relativist, it is correct today to say that this third claim was "True." The supervaluationist must deny this. Here again, I think the relativist's view accords better with common sense.

Finally, consider an assertion, made yesterday, of

- (55) $Poss_{today} : \exists x \exists y (x \neq y \wedge \text{Weather}(\text{tomorrow}) = x \wedge \text{Actually} : \text{Weather}(\text{tomorrow}) = y).$

Today it is still possible that the weather tomorrow will be different than it actually will be.

On the supervaluational semantics, the embedded clause

- (56) *Actually* : Weather(tomorrow) = y

will be false at every point of evaluation that includes the context of use (assuming that the weather is, at that point, still unsettled). So the whole sentence will be false at every point of evaluation that includes that context. Thus the proposition expressed by the sentence at that context can never be correctly said to be "True"—not even later, from the vantage point of a sunny tomorrow. On the relativist semantics, by contrast, an occurrence of (55) yesterday is true as assessed retrospectively from a sunny today, and from that vantage point speakers of the language can correctly characterize the proposition it expressed as "True." Here again, the relativist's predictions accord with common sense, while the supervaluationist's do not.

It is often supposed that a resolute commitment to branching worlds would require us to give up talk of actuality. That is the dilemma Lewis offers us. He picks common sense; others pick branching (say, on the grounds that it provides the best interpretation of current physical theory). The supervaluational semantics for branching does saddle us with Lewis's dilemma, but it is not the only option. As I have shown, it is possible to give a semantics for "actually" that makes good sense of our thought and talk about the future and past, even if we live in a "garden of forking paths." But to do so, we need to relativize truth to contexts of assessment.

References

- Belnap, N. and M. Green (1994). 'Indeterminism and the Thin Red Line'. *Philosophical Perspectives* 8: 365–88.
- *et al.* (2001). *Facing the Future: Agents and Choices in Our Indeterministic World*. Oxford: Oxford University Press.
- Borges, J. L. (1964). *Labyrinths: Selected Stories and Other Writings*. New York: New Directions.
- Kaplan, D. (1989). 'Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals'. In J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan*, 481–566. Oxford: Oxford University Press.
- King, J. C. (2003). 'Tense, Modality, and Semantic Values'. *Philosophical Perspectives* 17: 195–245.
- Lewis, D. (1970a). 'Anselm and Actuality'. *Noûs* 4: 175–88.
- (1970b). 'General Semantics'. *Synthese* 22: 18–67. Rep. in Lewis 1983.
- (1980). 'Index, Context, and Content'. In S. Kanger and S. Öhman (eds.), *Philosophy and Grammar*. Dordrecht: Reidel.
- (1986). *On the Plurality of Worlds*. Oxford: Blackwell.
- MacFarlane, J. (2003). 'Future Contingents and Relative Truth'. *Philosophical Quarterly* 53: 321–36.
- (2005). 'Making Sense of Relative Truth'. *Proceedings of the Aristotelian Society* 105: 321–39.
- Machina, K. F. (1976). 'Truth, Belief and Vagueness'. *Journal of Philosophical Logic* 5: 47–78.
- McGee, V. and B. McLaughlin (1998). 'Review of Timothy Williamson, *Vagueness*'. *Linguistics and Philosophy* 21: 221–35.
- (2004). 'Logical Commitment and Semantic Indeterminacy: A Reply to Williamson'. *Linguistics and Philosophy* 27: 221–35.
- Thomason, R. H. (1970). 'Indeterminist Time and Truth-Value Gaps'. *Theoria* 36: 264–81.
- Williamson, T. (1994). *Vagueness*. London and New York: Routledge.

Margins for Error in Context

Denis Bonnay and Paul Egré

According to the epistemic theory of vagueness defended in particular by Sorensen (2001) and Williamson (1994: 237), vagueness is due to our limited powers of discrimination: looking at a particular shade of red fabric, I may not be able to recognize that it is red, as a result of the specific granularity of my perceptual apparatus, which for instance makes the shade look somewhere between orange and red to me. Conversely, whenever I am confident that a particular shade of color is red, then this means that a slight variation in color should leave intact the fact that the shade is red. In Williamson's account of vagueness, this idea is expressed in terms of what Williamson calls *margin for error principles* for knowledge: whenever my knowledge is inexact in the sense of being approximative, it requires a sufficient margin for error in order to hold. More abstractly, the margin for error principle says that in order for me to know that some property P holds of an object d , then a slight modification of some relevant parameter in the object d should leave it in the extension of P . Expressed in terms of propositions and contexts, this means that in order to know that some proposition p holds in a context w , then p should still hold in a context w' that is only slightly different from w .

An important consequence of Williamson's margin for error semantics for inexact knowledge concerns the problem of iterations of knowledge. When applied to knowledge itself, the margin for error principle says that for me to know that I know p in a context w , I should know p in all contexts

We wish to thank the Editors, Max Kölbel and Manuel García-Carpintero, for their invitation and for encouraging us to investigate the contextualist dimension of our epistemic semantics. At the time we came up with the first version of Centered Semantics, our concern was primarily to have a modular semantics for knowledge and a plausible model of knowledge iterations; we became aware of the link with indexicality issues only secondarily. Special thanks are due to P. Schlenker for pointing out the link with Kaplanian themes, to J. van Benthem for directing our attention to Kamp's temporal logic, and to W. Rabinowicz for referring us to his work with K. Segerberg. We also thank M. Cozic, J. Dokic, and an anonymous referee for helpful comments, as well as audiences in Lisbon (ENFA3), Prague (Prague Colloquium on *Reasoning about Vagueness and Uncertainty*) and Amsterdam (PALMYR 4).

sufficiently similar to w . If knowledge is taken to be positively introspective, namely to be such that I know that I know p whenever I know p , then one can build a soritical argument to the effect that if I know p in a context w , then step by step, I should continue to know p even in contexts that are informationally very remote from w , including contexts in which p is false. For Williamson, the contradiction shows that knowledge, just like other mental states, is not luminous, namely that one can know p without knowing that one knows p .

In Dokic and Egré (2004), an axiomatic version of Williamson's argument was challenged, resting on the idea that higher-order knowledge and first-order knowledge need not obey the same margins of error. More recently, in Bonnay and Egré (2006), we showed how Williamson's margin for error semantics for knowledge can be modified in order to let the margin for error principle and the principle of positive introspection coexist. Despite their common inspiration, there remains a gap between the two approaches. On the one hand, as we shall see below, it can be shown that our semantic approach, like the one of Dokic and Egré, turns out to deny one of the premises of the syntactic version of Williamson's argument, namely the idea that the agent can have systematic knowledge of her margin of error. On the other hand, although both the syntactic approach followed by Dokic and Egré and the semantic approach put forward by Bonnay and Egré consist in adding a contextualist component to the evaluation of epistemic sentences, this common ingredient needs further articulation. In Bonnay and Egré (2006), in particular, we formulate a two-dimensional semantics for knowledge whose contextualist inspiration can be made more explicit if one introduces actuality operators. Moreover, the introduction of actuality operators will also allow us to handle the axiomatic version of Williamson's argument, while remaining closer to its initial premises. In what follows, our aim will therefore be to spell out the details of our contextualist approach to Williamson's paradox. The main point which we will develop concerns the idea that iterations of knowledge should remain anchored to the initial context of epistemic evaluation, in a way that blocks the soritical progression. In this respect, the paper should be seen as an effort to bring the two paradigms of epistemicism and contextualism about vagueness closer together.

The paper is structured as follows: in the first section we review the syntactic and semantic versions of Williamson's paradox; in Section 2 we present a solution to the semantic version in the framework of Centered Semantics, a two-dimensional semantics for epistemic logic. We discuss in the following section the connection between Centered Semantics and actuality operators, and use this connection for the analysis of the syntactic version of the paradox. We conclude with some comparisons in Section 4, in particular with Kamp's contextualist treatment of the sorites.

5.1 MARGINS AND ITERATIONS

Margin for error principles and knowledge iterations do not live well together, as they give rise to paradoxical conclusions. Williamson was the first to observe this phenomenon, and turned it into an argument against the view that knowledge obeys the (KK) principle of positive introspection, namely that whenever I know p , I am in a position to know that I know p . Others, however, in particular Gomez-Torrente (1997), and then Fara (2002), have expressed worries about this, by pointing out that for some propositions at least, the (KK) principle seems clearly to be valid, thereby putting into question the validity of margin for error principles. In this opening section, we want to review two versions of Williamson's paradox, first a semantic version, linked to Williamson's margin for error semantics for knowledge (Williamson 1994), and then a syntactic version, as exposed by Williamson (2000). The two versions will be needed in order to evaluate the limitations of our own semantic approach to the paradox.

Williamson's paradox can be presented in more than one guise. Williamson himself has given several versions of it (1992, 1994, 2000): in all cases, the scenario is that of an agent who is to give an estimation about some quantity, or to express a qualitative judgment depending on the variation of some parameter. For instance, it can be a situation in which I observe a group of people, and wonder whether it makes a crowd, depending on how many people I see (Williamson 1994); or it could be a situation in which I have to make an explicit estimation about the height of a tree that I observe at some distance (Williamson 2000); or it can be a situation in which I make judgments about whether I feel cold or not as the temperature varies (Williamson 2000). The examples can easily be multiplied. As yet another example in what follows, we shall consider a situation in which I observe a series of sticks at a certain distance, linearly ordered by size, and in which I am asked whether they fit in a certain box: in that case, I have to make a decision about the precise predicate "fits in the box". The predicate "fits in the box" is precise in the sense that whether an object fits in the box or not really depends on objective features of the world, and not on specific properties of the language. Indeed, by contrast to a predicate like "bald", "far", or even "crowd", the predicate does not seem semantically vague (namely such that it may fail to have a determinate extension).¹ We are more prone to admitting the existence of an objective cut-off point between the objects that fit in the box and those that do not, as assumed by the epistemic theory of vagueness more generally, than in the case of predicates like "bald"

¹ See Kölbel (forthcoming) for a detailed characterization of semantic vagueness.

or “far”, for which at least some ingredient of semantic vagueness seems to play an important role. For a scenario of this kind, the intuition is therefore clear that the predicate “fits in the box” has a determinate extension, and that our inability to say whether it applies to this or that object is imputable to the limitations of our discriminative capacities (as an effect of distance and other visual parameters).

5.1.1 The paradox, semantically

The easiest way to present Williamson’s paradox is to start with the semantic version. Intuitively, to say that knowledge obeys a margin for error principle means that in order to know that some proposition ϕ holds in a context w , ϕ must hold in all contexts that are sufficiently similar to w . For instance, consider a situation in which I reflect upon whether I feel cold over time (Williamson 2000). Suppose that in order for me to know that I feel cold at time t , not only must I feel cold at t , but I must feel cold already at time $t - 1$, and also at time $t + 1$. This gives us a margin for error principle expressing the idea that in order for me to become aware that I feel cold, my feeling of cold must last sufficiently long, or be sufficiently intense. Likewise, consider the situation in which I observe a series of distant sticks, such that all and only sticks of size less than 4 inches will fit in the box. Seeing the sticks at a distance, my judgments about whether they fit in the box need a sufficient margin for error in order to be reliable. By analogy to the temporal case, we can suppose that as a fact of my perception, in order for me to know that a stick of size n fits in the box, a stick of size $n + 1$ or $n - 1$ inches also has to fit in the box.

Both scenarios can be represented by means of the same linear Kripke model, a particular case of what Williamson 1994 defines as a *fixed margin model*, in which we let the worlds be named by the natural numbers: in the temporal case, the points represent successive moments of time; in the box scenario, they represent the different sticks ordered by sizes; more generally, the model can be used to represent the possible values of some scalar quantity.




Figure 5.1. A margin for error model

Two points i and j of the model are connected by an arrow if and only if they belong to the same interval of error, namely if $|i - j| \leq 1$. Letting p stand for “I feel cold”, then the model depicts a situation in which I feel cold from

time 0 up to time 3, and start not to feel cold from time 4 onward (or indeed somewhere between 3 and 4). If instead p denotes the property of fitting in the box, then the model represents the fact that all and only sticks of size less than or equal to 3 inches fit in the box. Like Williamson, we shall assume that the margin for error principle captures exactly the truth conditions of knowledge sentences. With respect to the model, this means that for every i :

(MS) $i \models K\phi$ if and only if for every j such that $|i - j| \leq 1, j \models \phi$

Suppose now that knowledge satisfies the (KK) principle of positive introspection. Then this means that for every i in the model:

(KK) If $i \models K\phi$ then $i \models KK\phi$.

It is now easy to see how a contradiction can be obtained from (KK) and (MS). Suppose that $n \models Kp$: for instance, at time n , I know that I am cold; or of the n -th stick in the series, I know that it fits in the box. By (KK), this entails that I know that I know it, namely: $n \models KKp$. From this and (MS), it follows in particular that $n + 1 \models Kp$. Thus, if I know that I am cold at time n , then I also know that I am cold at time $n + 1$. By induction, it means that if I know p at n , then I must know p at all subsequent points. Thus, supposing that I know I am cold at time 0, then I know that I am cold at all later times, a contradiction if the temperature can rise. Likewise, from my knowledge that a stick of size 1 fits in the box, it should follow that I know that sticks of any size however large fit in the box as well, which is a contradiction.

Several remarks can be made on the assumptions we used. Firstly, as Gomez-Torrente (1997) originally pointed out, the full generality of the (KK) principle is not needed to reach the soritical conclusion. Indeed, it suffices to assume that there is some i such that $i \models Kp$, and for which all iterations of knowledge hold. Propositions p satisfying this property are called *transparent* by Fara (2002). Likewise, the assumption that (MS) captures *exactly* the truth-conditions of knowledge is not needed either to derive the paradox. The “if and only if” in (MS) can be weakened into an “only if”, and the K operator may be given separate truth conditions.² Also, it can be noted that we only made use of the “upward” version of the margin for error principle in order to get a contradiction. In the temporal case, in particular, it would have been enough to work with non-symmetric margins for error, that is to suppose that in order to know some proposition at time t , then p must hold at t and at the next time, but not necessarily at the previous time. Similarly, if we think of the sticks scenario, if I know that a stick of size n fits in the box, then presumably I should know of smaller sticks that they fit too, but not so easily for larger sticks. The reason

² In the temporal case, for example, the upward version of the margin for error principle is expressible as $K\phi \rightarrow X\phi$, where X is the “next” operator. See Egré 2006 for a discussion of bimodal formulations of margin principles in terms of neighborhood modalities.

to assume symmetric margins of error, however, is that the accessibility relation in the above margin model can be taken to represent the relation of perceptual indiscriminability, which is standardly assumed to be symmetric, the crucial feature being the fact that this relation is reflexive and non-transitive. Here we assume symmetric margins of error, as they also allow us to get “downward” versions of the epistemic sorites we presented, but the reader should bear in mind that nothing substantial rests on this assumption, and may just as well replace the model of Figure 1 with a reflexive, right-directed model (in which iRj if and only if $j = i$ or $j = i + 1$). Finally, it is easy to see the same argument we presented can be used to show that (MS) conflicts with the principle of negative introspection, which says that if I don’t know ϕ at t , then I know that I don’t know that ϕ at t . In the model of Figure 1, $3 \models \neg K\neg p$. Assuming negative introspection from this point onward, it is easy to see that $\neg K\neg p$ should hold at all subsequent points, thereby contradicting the fact that $5 \models K\neg p$.

5.1.2 The paradox, syntactically

Let us now turn to the syntactic version of the puzzle, as appears for instance in Williamson (1992: §1) and Williamson (2000, ch. 5): by syntactic, we have in mind the reconstruction of the paradox by means of explicit axioms and rules of inference. This approach contrasts with the semantic one insofar as we do not need to start from explicit truth-conditions for knowledge. We now consider a language consisting of atomic propositions p_i , with $i \in \mathbb{N}$, with an epistemic operator K and the usual rules of constructions for formulas. From an intuitive semantic point of view (which the reader should keep in mind), each proposition p_i should now be taken to express an indexical proposition, for instance: “the temperature is i degrees” (over some appropriate scale), or: “the stick is i inches long”. This makes another important difference with the approach of the previous section: instead of using the indices i in the metalanguage, we now introduce indices explicitly in the object language. Furthermore, where $i \models Kp$ meant “when the temperature is i degrees, I know that it is cold”, Kp_i now means: “I know that the temperature is i degrees”; likewise, where $i \models Kp$ meant “looking at a stick of size i , I know that it fits in the box”, Kp_i means: “the stick I am looking at is i inches long”. Consequently, the form of the puzzle we discuss now involves quantitative instead of qualitative propositions.

In the presentation of the syntactic version of the paradox, Williamson considers a scenario in which a subject reflects upon his perceptual limitations. In particular, the subject is now supposed to know that his knowledge involves a specific margin for error. In the case where the margin is of 1 unit, we get a reflective version of the margin principle (ME), which we call (KME) (for Knowledge of the Margin of Error); that is, we have for every $i \geq 0$ (letting $i - 1 = 0$ for $i = 0$):

(ME) $K\neg p_i \rightarrow (\neg p_{i-1} \wedge \neg p_i \wedge \neg p_{i+1})$

(KME) $K[K\neg p_i \rightarrow (\neg p_{i-1} \wedge \neg p_i \wedge \neg p_{i+1})]$

The principle (ME) says that in order for me to know that the stick I am looking at is not i inches long, then the stick can't be $i + 1$ or $i - 1$ inches either; its iterative version, (KME), states a constraint of rationality on the part of the agent. Not only is the agent's knowledge subject to a margin for error, but the subject is aware that his knowledge obeys a margin for error principle. As a second rationality constraint, implicit in the semantic version of the puzzle, the knowledge of the subject is supposed to be closed under logical consequence, as expressed by the following axiom (K) (for Kripke's axiom):

(K) $K(\phi \rightarrow \psi) \rightarrow (K\phi \rightarrow K\psi)$

We suppose moreover that the subject's knowledge is positively introspective:

(KK) $K\phi \rightarrow KK\phi$

With the rule of modus ponens and the rule of substitution, the three principles (KK), (KME) and (K) are sufficient to replicate the foregoing soritic argument. Indeed, suppose now that $K\neg p_n$, that is the subject knows that the stick he is looking at is *not* of size n . By (KK), he knows that he knows it, $KK\neg p_n$. From (KME) and (KK), it holds in particular that $K(K\neg p_n \rightarrow \neg p_{n+1})$. From (K) and the last two formulas, it follows that $K\neg p_{n+1}$. Thus, from my knowledge that the stick is not of size n , I know that it is not of size $n + 1$. By induction, this means that for every n , I know that the stick is not of that size, a plain contradiction if the stick is supposed to be of a determinate finite size.

The syntactic version of Williamson's puzzle is slightly more involved than the semantic one, which only involved the two principles (MS) and (KK). The two versions of the puzzles are easily related, however: first, it is easy to see that the (K) principle is true at every point of the model of Figure 1, as a consequence of the semantic clause expressed by (MS).³ Likewise, if one assigns to every atom p_i the value true at i , and the value false everywhere else, as expected from the intended meaning of p_i , one can check in the same way that (KME)—like (ME)—is also universally true in the model.⁴ By contrast, (KK) is not valid in the model, since for instance $2 \models Kp$, but $2 \not\models KKp$. Thus, the two principles (K) and (KME), which may seem to express *prima facie* stronger principles of rationality than (ME) and (KK) themselves, are already validated by Williamson's margin for error semantics. Intuitively, however, (KME) does

³ Clearly, if $i \models K(\phi \rightarrow \psi)$ and $i \models K\phi$, then this means that at every j such that $|i - j| \leq 1$, $j \models \phi \rightarrow \psi$, and likewise $j \models \phi$, so $j \models \psi$, and thus $i \models K\psi$.

⁴ Suppose that $j \models K\neg p_i$; by (MS), this means that for every n such that $|j - n| \leq 1$, $n \models \neg p_i$, that is $j \neq i + 1$, $j \neq i$ and $j \neq i - 1$, and so $j \models \neg p_i \wedge \neg p_i \wedge \neg p_{i+1}$. So every j satisfies $K\neg p_i \rightarrow (\neg p_{i-1} \wedge \neg p_i \wedge \neg p_{i+1})$: this holds relative to every point k such that $|k - j| \leq 1$, so $K[K\neg p_i \rightarrow (\neg p_{i-1} \wedge \neg p_i \wedge \neg p_{i+1})]$ also holds everywhere in the model.

correspond to a more demanding condition than its non-reflective version (ME). The semantics we present in the next section will make that intuition precise.

5.2 CENTERED SEMANTICS

In the previous section we presented two versions of Williamson's paradox: first a semantic version involving qualitative predicates, then an axiomatic version involving quantitative predicates. In each case we get a paradox in the sense that, starting from plausible premises and plausible rules of inference, we end up with an obviously false conclusion. For Williamson, the paradox suggests that one of the premises must be rejected, namely the (KK) principle. In this section we present a different solution: in our view, neither (KK) nor (MS) is to be rejected strictly speaking. Rather, we make a more holistic move and propose a modification of the underlying semantic framework: roughly, the idea is that epistemic sentences ought to be evaluated with respect to more than one index in order to avoid spurious dependencies and to get a more realistic model of the agent's epistemic situation. Before introducing our semantics, it will be useful to review the general features of the standard epistemic semantics.

5.2.1 Iterations of knowledge and context-shifting

According to the standard Kripke–Hintikka semantics for knowledge, I know p in a context w if and only if p holds in all circumstances that are compatible with the information available to me at w (including the actual world, since knowledge is true). Standardly, a situation of knowledge is represented by means of a Kripke model $M = \langle W, R, V \rangle$, where W is a set of epistemic states or contexts, R is the epistemic accessibility relation, and V is a function specifying which atomic sentences are true at which worlds. Williamson's margin for error semantics for knowledge, as expressed for instance by clause (MS), is only a particular case of the standard semantics, according to which, given such a model M , $M, i \models K\phi$ iff for every $j \in W$ such that iRj , $M, j \models \phi$.

An important consequence of the standard semantics concerns sentences with iterated modalities: to check whether a sentence of the form $KK\phi$ holds at i , one needs to check whether ϕ holds at worlds that are two steps away from i along R in the model. More generally, for a sentence with n nested modalities, one has to visit worlds that are n -steps away from i along the accessibility relation. In the model of Figure 1, for instance, $1 \models KKp$, but $1 \not\models KKKp$, since 4 is a world accessible in three steps in which p is not true. To put it otherwise, whether $1 \models KKKp$ depends on whether $3 \models Kp$: however, it is arguable that whether $1 \models KKKp$ should depend only on whether $1 \models Kp$, and thus only on perceptual alternatives to context 1, not on perceptual alternatives to context 3.

With regard to the standard semantics, we may therefore say that iterations of knowledge shift the context, a phenomenon that is crucial to the generation of the soritical conclusion when positive introspection is supposed to be universally true in the model.⁵

This feature of the standard semantics is precisely what we think should be challenged to handle Williamson's paradox. In a margin model like that of Figure 1, the relation of epistemic accessibility can be taken to represent the notion of perceptual indiscriminability. As such the relation represents the fact that I cannot perceptually discriminate between pairwise adjacent sticks, but that I can discriminate between pairwise non-adjacent sticks. Thus, in order to know that a given stick fits in the box, all sticks that are adjacent to it must fit too. The question then is to what extent higher-order levels of knowledge have to depend on perceptual alternatives that are more remote than those needed to secure knowledge at the first level. For Williamson, such long-distance dependencies give an accurate picture of the attainability of higher-order knowledge. Indeed, Williamson sees iterations of knowledge as a process of "gradual erosion", precisely because, according to him, iterations of knowledge "widen margins of error".

We disagree with this conception. The context-shifting phenomenon is related to what we consider to be three main problems for this view. The first is a supervenience problem for higher-order knowledge: Williamson considers that higher-order knowledge supervenes on information in the world that lies beyond the information on which first-order knowledge supervenes. In our view, however, higher-order knowledge depends only on whether first-order knowledge occurs, without requiring the subject to gather more information in the world. The second problem concerns the computation of higher-order knowledge: Williamson is committed to the view that, depending on the context, I may know that I know . . . that I know *p*, and still fail to know that I know that I know . . . that I know *p*. From a cognitive point of view, this does not seem very plausible. Indeed, we become very quickly unable to keep track of levels of knowledge and it is even hard to say what a sentence like "Peter knows that he knows that it's raining, but does not know that he knows that he knows" really

⁵ Note that our use of the notion of context-shifting is consistent with the Kaplanian distinction between circumstance-shifting operators and context-shifting operators (See Kaplan 1989, and Recanati 2000 on the distinction). Standardly, "know" is described as circumstance-shifting because the truth of "*X* knows that *p*" in a context *w* depends on the truth of *p* in all circumstances or epistemic alternatives compatible with *w* for *X*. Precisely, we agree that *K* is circumstance-shifting, but it should not be context-shifting (when iterated). Our claim is that the truth of iterated knowledge modalities in a context *w* should depend only on the circumstances that are relative to *w* not on circumstances that are relative to contexts other than *w*, as they seem to do over non-transitive models. In S5 models of knowledge in which the accessibility relation is an equivalence relation, the accessibility relation is transitive, and this phenomenon of context-shifting does not arise. More abstractly, therefore, what we want and achieve in the next section can be described as a way of validating positive introspection over non-transitive models (Bonnay and Egré 2006).

means. The third problem is a problem of modularity: Williamson's account presupposes that higher-order levels of knowledge are subject to the same margins of error that apply at the first-order level. This supposition does not speak for itself, however, especially if we take perceptual knowledge and the reflective knowledge on one's perceptual knowledge to be two different modes of knowledge. In the next section we propose an alternative picture of higher-order knowledge, in which higher-order knowledge and first-order knowledge supervene on the same informational facts, in such a way that if I know p , I thereby know that I know p , without this second-order knowledge introducing an additional margin of error. This picture of higher-order knowledge is essentially the Cartesian picture, in which if I know something in a context w , then my knowing that I know remains anchored to the same context. Building on the previous intuitions, our aim is to show that the Cartesian picture is at least coherent and compatible with the idea that first-order knowledge comes about with a margin of error.

Before examining the details, however, it is fair to say that Williamson's own criticism of the Cartesian view gains much plausibility when the operator under discussion is the operator "it is clear that", which is the operator Williamson originally considers in his account of vagueness (Williamson 1994: Appendix). A central aspect of the vagueness phenomenon, in Williamson's account, concerns higher-order vagueness. Higher-order vagueness corresponds to the idea that some property can be clearly exemplified, without it being clear that the property is clearly exemplified: for instance, some object may clearly be red, without it being clear that the object is clearly red. Likewise, an object may not clearly be red, without it being clear either that the object is not clearly red.⁶ In the picture of higher-order knowledge that we defend, by contrast, if I know that the pen fits in the box, then I know that I know that the pen fits in the box. Our account also validates that if I don't know whether the pen fits in the box, then I also know that I don't know. The problem then is to determine to what extent the operators "it is clear (to me) that" and "I know that" can be considered equivalent in a situation of inexact knowledge like the one we are discussing.

Our response to this problem is that the two operators are most likely substitutable at the first-level, but not at higher-order levels. Certainly, if it is not clear to me that the pen fits in the box, then I don't know that it fits. Arguably, in the opposite direction, if it is clear that the pen fits in the box, then I know that it fits in the box. Is it contradictory, however, to suppose that it is not quite clear that it is not clear whether the pen fits in the box, and yet to say that I know that it is not clear whether the pen fits in the box? We don't see a contradiction here: if it is not clear whether the pen fits, and still not quite clear that it is not clear, I

⁶ We actually perceive an asymmetry between the two judgments: "if it is clear that p , then it is clear that it is clear that p " (schema 4), and "if it is not clear that p , then it is clear that it is not clear that p " (schema 5). In our view, the failure of schema 5 is much more intuitive than that of schema 4, and gives a better illustration of higher-order vagueness.

can perfectly reflect on my situation and be aware that it is not clear whether the pen fits or not, and yet, precisely because it is not clear that it is not clear, think that probably little would be needed for me to know whether or not the pen fits. In this regard, we agree with Halpern (2004) that a distinction should be made between the operators “it is clear that” and “I know that”. More precisely, what we say here is that even if they are taken to be synonymous in front of atomic sentences, they should no longer be taken to be such when iterated.

Consequently, we think it makes sense to preserve introspection principles such as “If I know that p , I know that I know p ”, and even “If I don’t know that p , I know that I don’t know”, without thereby committing ourselves to the denial of higher-order vagueness. By saying this, we consider that second-order knowledge describes a form of categorical awareness of one’s first-order epistemic state, more than an iterated judgment of clarity on content of that state proper.

5.2.2 Fixing the context

To avoid the context-shifting phenomenon induced by iterations of knowledge, we present a modification of the standard Hintikka–Kripke semantics, which we call Centered Semantics (Bonnay and Egré 2006), in which epistemic sentences are evaluated with respect to two indices. The semantics, as we shall see below, bears some interesting connections to two-dimensional formalisms developed earlier for temporal and for epistemic logic (Kamp 1971, Rabinowicz and Segerberg 1994, Halpern 2004). As just explained, the original motivation for the present approach was primarily to capture a notion of modularity for knowledge, however, namely to allow first-order knowledge and higher-order knowledge to obey distinct constraints (Egré 2006), and moreover to have a plausible model of knowledge iterations in situations of bounded rationality (Egré and Bonnay 2006).⁷

At the level of indices, this idea of modularity is cashed out in the following way: given a couple (w, w') , the second index w' determines the atomic information available at w' ; the first index w , by contrast, determines what the subject knows, including about himself, when he is located at w . Epistemic sentences are evaluated with respect to couples (w, w') such that w' is a world perceptually accessible from w : in this way, the index w serves to fix the perspective of the subject. Borrowing from the terminology of Rabinowicz and Segerberg 1994,

⁷ CS is a particular case of a family of resource sensitive semantics, which we call Token Semantics, of which the standard Kripke semantics is also a particular case (Bonnay and Egré 2006). We only focus on CS in the present paper, first because CS is the first semantics in the hierarchy of Token Semantics (with Kripke semantics at the other extreme), but also because it is the most appropriate framework for the discussion of the issues of epistemic indexicality connected to Williamson’s paradox. One feature of TS that is relevant for the problem of introspection concerns the idea that trivialization of knowledge iterations need not occur at level 2, as happens with CS, but only at higher levels.

we may call the first index the “perspective-world” and the second index the “reference-world”. Typically, however, we are interested in what the subject knows at w , when the primary information he receives comes from w itself. Hence, we define the semantics in two stages: first relative to arbitrary couples (w, w') , then relative to single indices, corresponding to the diagonal case in which $w = w'$. Given a Kripke model $\mathcal{M} = \langle W, R, V \rangle$, the satisfaction is defined inductively as follows:

- (i) $\mathcal{M}, (w, w') \models_{\text{CS}} p$ iff $w' \in V(p)$.
 - (ii) $\mathcal{M}, (w, w') \models_{\text{CS}} \neg\phi$ iff $\mathcal{M}, (w, w') \not\models_{\text{CS}} \phi$.
 - (iii) $\mathcal{M}, (w, w') \models_{\text{CS}} (\phi \wedge \psi)$ iff $\mathcal{M}, (w, w') \models_{\text{CS}} \phi$ and $\mathcal{M}, (w, w') \models_{\text{CS}} \psi$.
 - (iv) $\mathcal{M}, (w, w') \models_{\text{CS}} K\phi$ iff for all w'' such that wRw'' , $\mathcal{M}, (w, w'') \models_{\text{CS}} \phi$.
- (*) $\mathcal{M}, w \models_{\text{CS}} \phi$ iff by definition $\mathcal{M}, (w, w) \models^{\text{CS}} \phi$

According to clauses (i)–(iii), atomic sentences and boolean compounds thereof depend on the second index only, namely the reference-world; by contrast, the evaluation of epistemic operators depends essentially on the first index or perspective-world, as expressed by clause (iv). This feature becomes crucial for the evaluation of sentences with iterated modalities: in that case, clause (iv) ensures that we keep track of the initial point of evaluation, seen as the “center” of the worlds that count as relevant alternatives in the evaluation of epistemic modalities.

If we consider the model of Figure 1, we can check indeed that the semantics is equivalent to the standard semantics for propositional sentences and for sentences without nested modalities. For instance, $2 \models_{\text{CS}} Kp$, just as in the standard case, since $(2, 1)$, $(2, 2)$, and $(2, 3) \models^{\text{CS}} p$, and 1, 2, 3 are the only worlds accessible from 2. The situation changes for the case of iterated modalities: with the standard semantics, $2 \not\models KKp$; however, it is easy to check that $2 \models_{\text{CS}} KKp$: by definition, $(2, 2) \models_{\text{CS}} KKp$ if and only if $(2, 1)$, $(2, 2)$ and $(2, 3) \models_{\text{CS}} Kp$, if and only if all of these in turn satisfy p , which we have seen to be the case.

More generally, relative to a world w , the semantics bounds the number of possible moves in the model to those worlds that are accessible from w in one step. In this manner, whether KKp holds or not at a world w depends only on whether Kp already holds *there*, and not on whether Kp holds at more remote indices. In centered semantics, iterations of knowledge no longer give rise to context-shifting, but remain anchored to the context relative to which knowledge of the first-order is determined. Thus, in the same way in which it articulates a notion of modularity, centered semantics also captures an element of indexicality of higher-order knowledge.

Before saying more about this notion, it is now easy to see how the semantics blocks Williamson’s paradox. First, the reader can check that the semantics validates principle (KK) of positive introspection, that is $w \models_{\text{CS}} K\phi$ implies $w \models_{\text{CS}} KK\phi$. In Bonnay and Egré (2006), it is shown that the semantics is

sound and complete for the normal logic K45, namely for the axiom system which validates, besides axiom K and (KK), the principle of negative introspection (namely $\neg K\phi \rightarrow K\neg K\phi$). If we graft the semantics directly onto Williamson's fixed margin semantics, which validates the axiom T of veridicity ($K\phi \rightarrow \phi$), the resulting semantics becomes sound and complete for the system S5 of fully introspective knowledge. Thus, if we consider (MS) above in section 1.1, its formulation in the framework of centered semantics becomes:

(MS') $i \models_{CS} K\phi$ iff for every j such that $|i - j| \leq 1$, $(i, j) \models_{CS} \phi$

namely: I know ϕ in context i if and only if ϕ holds at every world j that is perceptually indiscriminable from i . With this enrichment of the semantics, the principles of positive introspection and of margin for error are compatible again. In particular, from $n \models_{CS} Kp$, it follows that $n \models_{CS} KKp$, since positive introspection is CS-valid. By (MS'), this entails in particular that $(n, n + 1) \models_{CS} Kp$. But from that, it does not follow that $n + 1 \models_{CS} Kp$, since this would mean that $(n + 1, n + 1) \models_{CS} Kp$. Another way to see the point is by considering the above model: in the model, $Kp \rightarrow KKp$ is CS-true at every point of the model, likewise Kp holds at point 0, but now this is compatible with the fact that p holds only in a limited segment of the model. As a consequence, the move from the standard semantics to centered semantics blocks the soritic progression. Despite this, (MS') remains as close as possible to (MS): indeed, (MS') still means that in order to know some proposition in a given context, that proposition must be true in all contexts that are sufficiently similar to it; but this dependence to the initial context is now more tightly articulated.

5.3 KNOWLEDGE AND ACTUALITY

Centered Semantics makes margin for error models and introspection principles compatible, and blocks the path to Williamson's paradoxical conclusion. What about the syntactic version? Centered semantics, as we shall see, turns out to invalidate one of the premises of Williamson's argument, namely principle (KME), by which the agent knows her margin of error. The rejection of that premise, however, is arguably too strong: the aim of this section is to examine whether this is so, and to give a more precise discussion of the status of this principle. To do so, we present a reformulation of Centered Semantics in terms of actuality operators. The idea that perceptual possibilities are indexical—what I can or cannot perceive only depends on what is indistinguishable from the actual situation I am in—is indeed hardwired in our semantics. The introduction of actuality operators allows us to make explicit this implicit indexical element, and to articulate the semantic and syntactic versions of the paradox more precisely.

5.3.1 Knowing the margin

Williamson's syntactic paradox shows the inconsistency of a set of four principles: in order to escape the contradiction, and short of changing the logic, it is necessary to reject at least one of the principles used in the argument. Given the correspondence between the syntactic and the semantic versions of the paradox, we must expect Centered Semantics to invalidate at least one of those principles. This can't be (KK), since Centered Semantics was designed to make it valid. This should not be (ME), since the whole point of (ME) is to account for the inexactness of perceptual knowledge. Hence there are only two principles left for rejection, namely (K) and (KME).

Centered Semantics does validate (K), and we think this is as it should be. (K) expresses closure of knowledge under logical inferences: if I know that p implies q and if I also know that p , then I know that q . As a rationality requirement, (K) is an idealization. Real life agents are not logically omniscient, they do not consciously know all the consequences of what they consciously know, and in this respect the principle is obviously very strong. For our purposes, however, rejecting (K) would not count as a real solution to the syntactic paradox: for the paradox could still be derived for a logically omniscient agent with inexact knowledge (even brilliant mathematicians get cold sometimes).

So what about (ME) and (KME)? Let us state these principles in their syntactic form as in paragraph 1.2:

$$(ME) \quad K\neg p_i \rightarrow (\neg p_{i-1} \wedge \neg p_i \wedge \neg p_{i+1})$$

$$(KME) \quad K[K\neg p_i \rightarrow (\neg p_{i-1} \wedge \neg p_i \wedge \neg p_{i+1})]$$

Until now we have used a qualitative margin for error model, in which a given proposition like "I feel cold" or "the pen fits in the box" held up to some threshold. In order to match the syntactic formulation of (ME) and (KME), we can easily turn this model into a quantitative model, in which atomic propositions like p_3 : "the pen is 3 inches tall" express the underlying quantities:




Figure 5.2. Margin model with quantitative propositions

(ME) is valid on this model. For example, $4 \models_{CS} K\neg p_2$. For (ME) to hold, it must hold that $4 \models_{CS} \neg p_1 \wedge \neg p_2 \wedge \neg p_3$, which is the case. By forcing here that $4 \models_{CS} \neg p_3$, (ME) implements my margin for error: if I was looking at

a pen of size 3, I would not be able to know that it is not of size 2, because my eyes cannot distinguish two pens whose sizes differ by no more than one inch.

However, surprisingly enough, the validity of (ME) on a model does not imply the validity of (KME) on that same model. To see this, look at what happens at point 4 in the model of Figure 2. For $K(K\neg p_2 \rightarrow \neg p_3)$ to hold at 4, we must have $(4,3) \models_{CS} K\neg p_2 \rightarrow \neg p_3$. Clearly, $\neg p_3$ is false at $(4,3)$, because $(4,3)$ is the perceptual alternative to 4 in which the pen is of size 3. But $K\neg p_2$ is true at $(4,3)$, just because $\neg p_2$ is true at all perceptual alternatives relevant to the situation 4 in which I am in, namely $(4,3)$, $(4,4)$ and $(4,5)$.

To make things clear, it will be useful to define the following notions in full generality. We refer to Kripke semantics as KS in what follows.

Definition 1. Let $\mathcal{M} = \langle W, R, V \rangle$ be a Kripke model, S a semantics, a formula ϕ is $\mathcal{M}$, S -valid iff for all $w \in W, w \models_S \phi$.

Definition 2. Let $\mathcal{F}$ be a class of frames, S a semantics, a formula ϕ is $\mathcal{F}$, S -valid iff for all models $\mathcal{M}$ based on a frame $F \in \mathcal{F}, \phi$ is $\mathcal{M}$, S -valid.

The following facts are well-known:

Fact 1 (Model necessitation for KS). Let $\mathcal{M} = \langle W, R, V \rangle$ be a Kripke model, if ϕ is $\mathcal{M}$, KS -valid, then $K\phi$ is $\mathcal{M}$, KS -valid too.

Fact 2 (Frame necessitation for KS). Let $\mathcal{F}$ be a class of frames, if ϕ is $\mathcal{F}$, KS -valid, then $K\phi$ is $\mathcal{F}$, KS -valid too.

Of course, these facts hold as special cases for S5-models and S5-frames. However, there is major difference here between KS and CS. Frame necessitation holds for both KS and CS:

Fact 3 (Frame necessitation for CS). Let $\mathcal{F}$ be a class of frames, if ϕ is $\mathcal{F}$, CS -valid, then $K\phi$ is $\mathcal{F}$, CS -valid too.

By contrast, as we just saw, although (ME) is valid over the model of Figure 2, the formula (KME) is not. This suffices to show that model necessitation fails for CS, namely:

Fact 4 (Failure of model necessitation for CS). There exists a Kripke model $\mathcal{M}$ and a formula ϕ such that ϕ is $\mathcal{M}$, CS -valid and $K\phi$ is not $\mathcal{M}$, CS -valid.

It would be possible to recover model necessitation using a stronger notion of validity: following Rabinowicz and Segerberg (1994), a formula may be called *strongly valid* on a model according to CS if and only if it is true at all pairs (w, w') in the model. But note that this notion of validity is more like a model-theoretic artefact. CS is about what is true at one world. Pairs of worlds are just a notational

device meant to implement the indexical character of knowledge. In this respect, the natural notion of validity on a model is just truth according to CS at all worlds (that is truth at all pairs (w, w)).

To summarize, margin for error models do validate (ME)—again, this is as it should be—but they do not validate (KME). This is because model necessitation, a familiar feature of standard Kripke semantics, fails for our centered semantics. The failure of the semantic paradox in CS is thus reflected at the syntactic level by the failure of one of the principles needed to derive the paradox. Is this a good thing? Do we have good reasons to reject (KME) while accepting (ME)? Or to put the question in the reverse direction: should the acceptance of a principle like (ME) commit us to the acceptance of a principle like (KME)?

We don't think it should be so. To use a distinction common in cognitive science, while a principle like (ME) expresses a constraint on the acquisition of knowledge of a subject whose discriminative capacities are limited, this constraint may very well be operative at the subpersonal rather than at the personal level. By contrast, a principle like (KME) requires that the subject be aware not only that her knowledge is constrained by a margin of error, but even be aware of what the value of the margin is. There is a distinction to make, in this regard, between a principle like (KK) and a principle like (KME): although both principles are iterative principles, the latter requires more than the former, since it requires that the agent be aware of the structural constraints that impinge on her knowledge. A principle like (KK), when considered in a Cartesian perspective, can be taken to reflect a notion of inner sense: whenever the subject is in a state of knowledge, then she can *thereby* know that she is in a state of knowledge. But there is no reason why a rational subject of this kind, when her knowledge is systematically constrained by a margin of error, should *thereby* know the value of the margin, even if we assume her to know that her knowledge is imprecise. To suppose otherwise would be to suppose a capacity of self-knowledge even more powerful than Cartesian luminosity (as expressed by KK).

To be sure, Williamson himself never presents a principle like (KME) as a principle that ought to hold systematically. In his statement of the syntactic version of the paradox, Williamson simply considers that the subject “can reflect on the limitations of his eyesight and ability to judge heights”. But as emphasized in Dokic and Egré (2004), the source of this knowledge is not discussed by Williamson. We certainly agree with Williamson that the subject can reflect on his eyesight and ability to judge heights, and more generally that the subject can come to know her margin of error, but then such a knowledge is most likely acquired empirically, or simply revealed to the subject, rather than acquired a priori by some kind of inner sense. By denying the validity of (KME) over the model of Figure 2, we therefore deny that the subject can have immediate knowledge of her margin of error when (ME) holds systematically. More precisely, we saw that for every world in the model of Figure 2, there is

an instance of (KME) that the subject fails to know: consequently, the subject does not know her margin of error (even though she may know particular instances of (KME): for instance, at 5, she knows $K(K \neg 10 \rightarrow \neg 11)$, but that is far from sufficient, and then she in fact knows more, namely $KK \neg 10$, and $K \neg 11$).

That being said, the question remains open how one should represent the knowledge of a subject who knows her margin of error systematically. To clarify this issue, we shall first present an equivalent version of (ME), which is such that it is known everywhere in the model. As we will see, however, this modified version of (ME), and the modified version of (KME) that it induces, are too weak to enable the subject to reason on her own limitations.

5.3.2 Actuality operators

A remarkable feature of the knowledge operator K in the framework of CS is its kinship with actuality operators. Actuality operators are operators like “now” or “actually”, which allow us to preserve the reference to the time of utterance within the scope of context-shifting operators such as future or past tense operators.⁸ From a model-theoretic point of view, actuality operators can be seen as “resetting” the context of evaluation to the context of utterance. One way to implement this idea is precisely within a two-dimensional framework (Kamp 1971, Segerberg 1973). In Kamp’s temporal semantics for “now”, sentences are evaluated relative to ordered pairs (t, t') , where t represents the utterance time, and t' the event time. An atom p is true relative to (t, t') if p is true at the event time t' . Likewise, $(t, t') \models G\phi$ if and only if ϕ is true at all times later than the event time t' . By contrast, $(t, t') \models N\phi$ if and only if $(t, t) \models \phi$, namely precisely when ϕ is true at the time of utterance. In other words, “now” resets the event time to the utterance time.

Looking at the behavior of K in CS, we can see that it commands to reset the context of evaluation to the initial context in an analogous way: to evaluate a formula of the form KKp at world i , for instance, one first moves to worlds accessible from i in accordance with the standard semantics, to check whether Kp holds there; at that point, however, instead of moving to more remote worlds to check whether they satisfy p , one backtracks to i to ensure that Kp already holds there. In CS, the K operator thus behaves as a combination of a standard knowledge operator with an actuality operator: “I know ϕ ” is true if and only if ϕ is true at all *actual* epistemic possibilities. A closer comparison with Kamp’s semantics is appropriate to emphasize this point: in CS, an atomic sentence p is true at an ordered pair (w, w') if p is true relative to the second index, just as in Kamp’s logic; likewise, the boolean connectives are evaluated

⁸ See Edgington (1985: 561), from whom we borrow this formulation. We are indebted to J. van Benthem for the suggestive “resetting” image which follows.

in exactly the same way; by contrast, $(w, w') \models_{\text{CS}} K\phi$ if and only if ϕ is true at all epistemic alternatives to w , and not to w' . Strictly speaking therefore, K behaves neither exactly like N , nor exactly like G , but as a combination of the two: like G and unlike N , K is a universal modality quantifying over a certain set of accessible worlds; unlike G but like N , K refers these worlds back to the first index.

Let us call $\mathcal{L}(K)$ the basic modal language with K as the only modal operator, and $\mathcal{L}(K, A)$ the same language augmented with a second modality A (for “actually”). Interpret $\mathcal{L}(K)$ according to the rules of CS; conversely, interpret $\mathcal{L}(K, A)$ according to the rules of the standard two-dimensional semantics, namely Kamp’s semantics in which, given a model $M = \langle W, R, V \rangle$: $(w, w') \models A\phi$ iff $(w, w) \models \phi$, and $(w, w') \models K\phi$ iff for every w'' such that $w'Rw''$: $(w, w'') \models \phi$. Finally, consider the function $*$ from sentences of $\mathcal{L}(K)$ to sentences of $\mathcal{L}(K, A)$ defined by: $p^* = p$, $(\neg\phi)^* = \neg\phi^*$, $(\phi \wedge \psi)^* = (\phi^* \wedge \psi^*)$, and $(K\phi)^* = AK\phi^*$. The function is a satisfaction-preserving map, as the following proposition shows:

Proposition 1. $M, (w, w') \models_{\text{CS}} \phi$ iff $M, (w, w') \models \phi^*$

The proof is by induction over the complexity of formulas. The atomic and boolean cases are immediate. For $\phi = K\psi$:

$$\begin{aligned}
 & M, (w, w') \models_{\text{CS}} K\psi \\
 & \Leftrightarrow \text{for every } w'' \text{ such that } wRw'': M, (w, w'') \models_{\text{CS}} \psi \\
 & \Leftrightarrow \text{by induction hypothesis, for every } w'' \text{ such that } wRw'': M, (w, w'') \models \psi^* \\
 & \Leftrightarrow M, (w, w) \models K\psi^* \\
 & \Leftrightarrow M, (w, w') \models AK\psi^* \\
 & \Leftrightarrow M, (w, w') \models (K\psi)^*.
 \end{aligned}$$

As an immediate corollary, we obtain that $M, w \models_{\text{CS}} \phi$ if and only if $M, (w, w) \models \phi^*$. We thus obtain a translation from $\mathcal{L}(K)$ into $\mathcal{L}(K, A)$ which shows that the non-standard knowledge operator, as interpreted by CS, really is the composition of a standard knowledge operator with an actuality operator. This translation casts a new light on the results about CS mentioned in the previous section. The reader can check directly, for instance, that the complex operator AK satisfies both positive and negative introspection over pairs (w, w) with respect to the standard two-dimensional semantics, namely $AK\phi \rightarrow AKAK\phi$, and $\neg AK\phi \rightarrow AK\neg AK\phi$. Moreover, it satisfies Kripke’s axiom and the rule of necessitation over arbitrary couples: if ϕ is strongly model-valid (true relative to all pairs in a model), then so is $AK\phi$. Relative to diagonal pairs (w, w) , this is no longer the case: to repeat the point, every pair (w, w) of the model of Figure 2 satisfies the translation $(\text{ME})^*$ of (ME), namely

$AK\neg(p_i \rightarrow \neg p_{i-1} \wedge \neg p_i \wedge \neg p_{i+1})$, but not the translation $(KME)^*$ of (KME), that is $AK[AK\neg p_i \rightarrow (\neg p_{i-1} \wedge \neg p_i \wedge \neg p_{i+1})]$.

5.3.3 Margins actualized

By showing the correspondence between K relative to CS and AK relative to the standard two-dimensional semantics, it may seem we did not gain much. We did gain an important insight, however, since the introduction of an actuality operator shows that the problematic principles can be given a more articulated formulation without giving rise to paradox. In this section, we present another formulation of the margin principle using the actuality operator, which we call (MEA): like $(ME)^*$, (MEA) is logically equivalent to (ME) with respect to the standard semantics and to CS; when embedded under the operator K , however, (MEA) yields a modified version of (KME) that comes out valid in the model of Figure 2.⁹ For ease of exposition, we only consider the upward version of the margin for error principle in order to formulate (MEA):

$$(MEA) \quad K\neg p_i \rightarrow A\neg p_{i+1}$$

(MEA) says that if I know that the pen is not i inches tall, then, it is actually the case that the pen is not $i + 1$ inches tall. Since the actuality operator precedes an atomic sentence, (MEA) has the same content as upward (ME). When embedded under K , however, the principle is likely to express something different from the upward version of (KME), namely:

$$(KMEA) \quad K(K\neg p_i \rightarrow A\neg p_{i+1})$$

(KMEA) says that in all my epistemic alternatives, if I know *there* that p_i is false, then p_{i+1} is *actually* false. Note that (KMEA) is much weaker than (KME)—even though, as we just said, (MEA) and (ME) equivalent. The reason is that, with (KME), the information $\neg p_{i+1}$ holds at all our epistemic alternatives in which the consequent $K\neg p_i$ is satisfied. By contrast, in (KMEA) the information $\neg p_{i+1}$ is no longer referred to the epistemic alternatives, since the actuality operator semantically extracts it from the scope of the operator K . The phenomenon is exactly analogous to the fact that, by contrast to an anaphoric pronoun like “he”, an indexical pronoun like “I” cannot be bound by the attitude-holder when embedded under an attitude operator: while “he” in “John thinks he is right”

⁹ To complete the picture, one should also consider the formulation of upward (ME) in which both antecedent and consequent are prefixed with an actuality operator, namely: $AK\neg p_i \rightarrow A\neg p_{i+1}$, which we may call (AMEA). We only focus on (MEA), however, because (MEA) and (AMEA), and their respective K -prefixed versions, are equivalent in the framework of CS.

can refer back to “John”, “I” in “John thinks I am right” can only refer to the speaker and not to John.

As can be expected, using (KMEA), it is impossible to fall into the kind of soritic induction obtained by Williamson when (KME) is used instead. To be sure, let us consider what we get from $K\neg p_i$:

$K\neg p_i$, hypothesis
 $KK\neg p_i$, from (KK)
 $K(K\neg p_i \rightarrow A\neg p_{i+1})$, by (KMEA)
 $KK\neg p_i \rightarrow KA\neg p_{i+1}$, from (K)
 $KA\neg p_{i+1}$, by (MP)

We derive $KA\neg p_{i+1}$ instead of $K\neg p_{i+1}$. But $KA\neg p_{i+1}$ does not express any new real knowledge, for it says that each of my epistemic alternatives is such that at the actual world p_{i+1} is false, a case of vacuous quantification. $KA\neg p_{i+1}$ does *not* say that $\neg p_{i+1}$ is true in all my epistemic alternatives, hence, it does not say that I know that $\neg p_{i+1}$.¹⁰ Formally, one cannot conclude $K\neg p_{i+1}$ from $KA\neg p_{i+1}$ (our model in Figure 2 would be a case in point).

If we now assume CS as background semantics for the richer language $\mathcal{L}(A, K)$ —in order to validate (KK)—with the additional rule that $M, (w, w') \models_{CS} A\phi$ if and only if $M, (w, w) \models_{CS} \phi$, we can see that like (KK) and (K), (MEA) and (KMEA) are both CS-valid over the model of Figure 2. The reader should compare the situation with the counterexample given earlier: we do have now that $4 \models_{CS} K(K\neg p_2 \rightarrow A\neg p_3)$, because in particular $(4, 3) \models_{CS} A\neg p_3$ (where previously we had: $(4, 3) \not\models_{CS} \neg p_3$).

With (KMEA), we are thus able to regain the idea that a rational subject can know that her knowledge is subject to a margin of error. This kind of self-knowledge does not enable the subject to increase her perceptual knowledge, however. In CS, (KMEA) and (MEA) turn out to be equivalent: both the interpretation of K in the antecedent of (MEA) and the interpretation of A in the consequent make it so that, when knowledge is iterated, the evaluation remains anchored to the actual situation. Likewise, if we consider the conclusion of the above derivation, $KA\neg p_{i+1}$ is not CS-equivalent to $K\neg p_{i+1}$, but is in fact CS-equivalent to $A\neg p_{i+1}$: within the scope of K , the actuality operator blocks the shift to further alternatives. The subject's knowledge of her margin of error no longer gives rise to paradox. While the introduction of actuality operators is fairly drastic, the resulting set of axioms is nevertheless interesting on one aspect: arguably, it gives a model in which the knowledge of one's margin of error is automatically satisfied, without

¹⁰ As pointed out by Rabinowicz and Segerberg (1994), $\phi \rightarrow KA\phi$ is weakly valid and $A\phi \rightarrow KA\phi$ strongly valid more generally. For that reason, $KA\phi$ should not be read as “I know that ϕ holds actually” (from which, in natural language, we would normally infer: “I know that ϕ ”). Rather, when true, $KA\phi$ means that the epistemic alternatives of the agent are alternatives to a situation in which ϕ holds.

allowing the subject to gain knowledge simply by reflecting on his perceptual capacities.

5.4 COMPARISONS

In order to close this paper, some brief comparisons will be in order, both to evaluate the plausibility and the limits of our solution to Williamson's paradox, and also to connect the present treatment with other contextualist approaches to the phenomenon of vagueness more generally.

5.4.1 Varieties of knowledge

By reformulating (ME) and (KME) into (MEA) and (KMEA), we found a way to connect one plausible reformulation of Williamson's syntactic principles with one plausible background semantics for knowledge, namely Centered Semantics. The good thing about this approach is that the subject can be assumed to know even the exact size of her margin of error; the limitation is that our solution may seem to restrict too much what the subject can know on the basis of this second-order knowledge. Indeed, from her knowledge that the stick is not of size 0, and the fact that she knows her margin for error, the subject cannot infer that the stick is of size 1. The kind of higher-order knowledge afforded by a principle like (KMEA) does not allow the subject to learn more than she did from her first-order knowledge.

In Dokic and Egré (2004), the principle (KME) is also pointed to as the suspect principle at the origin of the syntactic version of Williamson's paradox. However, instead of using actuality operators, it is argued that Williamson failed to distinguish between kinds or methods of knowledge in his assumptions. In Dokic and Egré's approach, Williamson's principles are reformulated in terms of two knowledge operators, one for perceptual knowledge K_π , and the other for reflective knowledge K , with some bridge axioms to connect the two. From the assumption that $K_\pi \neg p_i$, it is possible to infer $K K_\pi \neg p_i$, and to reach the conclusion $K \neg p_{i+1}$. From that, however, it does not follow that $K_\pi \neg p_{i+1}$, as the account predicts that something can be known reflectively, without being known perceptually. Likewise, on the account presented in the previous section, from the assumption $K \neg p_i$, we reach the conclusion $KA \neg p_{i+1}$, without inferring $K \neg p_{i+1}$.

Technically, the two substitutes to (KME) bear some similarity, since both of them bring several operators into play (two knowledge operators in the one case, a single knowledge operator versus the combination of that operator with an actuality operator in the other), and both of them replace the problematic inference from $K \neg p_i$ to $K \neg p_{i+1}$ by an analogous rule involving distinct operators. Conceptually, there is a significant difference, however, since in the framework of

Dokic and Egré the transition from $K_{\pi}\neg p_i$ to $K\neg p_{i+1}$ means that, even though it is not known perceptually that the stick is not of size $i + 1$, this can be known at least reflectively. Thus, by reflecting upon her perceptual limitations, the subject can at least make an inference about the size of the stick. In the present approach, any increase in perceptual knowledge is blocked in the same way, but arguably $KA\neg p_{i+1}$ denies the subject any kind of inferential knowledge about the stick, since in CS it is equivalent to $A\neg p_{i+1}$. To avoid this trivialization, one would need to modify the interaction of the operators A and K , in such a way that $KA\neg p_{i+1}$ can still express knowledge, without thereby collapsing into $K\neg p_{i+1}$.

The semantics proposed by Rabinowicz and Segerberg (1994), in which the accessibility relation for K is defined directly over pairs of worlds, addresses precisely this problem and it would be interesting to see how to adapt it to the case of CS. If a modified account were successfully carried out, however, the outcome would still be that K and KA do not express exactly the same kind of knowledge. The approach of Dokic and Egré and the present one are perfectly consistent with each other, consequently: indeed, what we initially put into question in Section 3.1 is the idea that a subject whose knowledge is constrained by a margin of error principle can thereby know her margin of error. More specifically, by reformulating (ME) and (KME) into (MEA) and (KMEA) in Section 3.3, we gave a model in which the reflective knowledge of the margin is automatically satisfied, but without allowing the subject to extend her knowledge. This model is consistent with the view that, without an additional empirical input, reflection upon one's perceptual capacities should not allow one to extend one's knowledge. But if, finally, the subject is to *learn*, from experience or from some oracle, that her knowledge obeys a specific margin of error, then we agree that the subject might be able to use this additional knowledge and extend her initial perceptual knowledge.

In that case too, however, we remain in agreement with the conclusions of Dokic and Egré: a model in which the agent learns her margin should be a model in which the sources and varieties of knowledge are explicitly distinguished. Suppose the subject were to learn from some infallible oracle that her initial knowledge is subject to a margin of error, revealing to her the size of the margin: then obviously this second-order knowledge does not stand on the same footing as her initial knowledge, since it is acquired from an external source and not subject to any error by definition. But even in the case in which the subject learns from experience and careful empirical reflection upon her knowledge that her knowledge is subject to a margin of error, there are good reasons to doubt that this second-order knowledge is of the same kind as her initial knowledge. This new kind of knowledge may also be subject to error and approximations, but the method by which this knowledge is acquired is most likely not the same as that which underlies her first-order knowledge.

Our account may therefore be summarized by the following list of options: if the subject's knowledge obeys (ME), (KK) and (K), then the subject may simply fail to know what her margin is, as is reflected when CS is adopted as background semantics; but if the subject is to know her margin a priori, as happens when (ME) and (KME) are rephrased as (MEA) and (KMEA), then the subject's knowledge of her margin should not give her more information than is given to her by her first-order knowledge; finally, if a subject whose knowledge is described by (ME), (KK) and (K) is to learn the information that her knowledge obeys (ME), then we agree that this information may take her initial knowledge further, but our claim is that this knowledge is not homogeneous with the initial knowledge to which (ME) applies.

5.4.2 Contextualism and the sorites

The second question we want to address concerns semantic relativity, and the link of the present account to contextualist treatments of vagueness. By semantic relativity, we refer to the fact that the validity or invalidity of principles like (KK) and (KME) depends on a choice between distinct semantic characterizations of knowledge. Williamson's paradox shows the incompatibility of the introspection principle (KK) when margin for error principles are assumed to hold unrestrictedly for knowledge. In particular, we saw that Williamson's margin for error semantics, based on the standard semantics for knowledge, invalidates the axiom (KK). By contrast, our Centered Semantics makes (KK) valid, but we saw that it does not validate all margin for error principles: thus, the model of Figure 2 validates (ME), but invalidates (KME). Which semantics is the right semantics? The first thing to say is that both semantic characterizations of knowledge may equally fail to be right. Indeed, both treatments of knowledge rest on strong idealizations and consequently both are likely to yield an overly simplified description of the situations of inexact knowledge that we discussed.¹¹ But the question then should be: which of the two accounts of knowledge seems more likely to be right? To our minds, Centered Semantics gives a richer model of the context-dependence of margin for error principles for the kind of scenarios we discussed here. To a large extent, the soritical progression induced by the iteration of *K* modalities is an artefact of the standard Kripke–Hintikka semantics. To repeat the point of the previous section: knowing that ϕ , in Centered Semantics, means that ϕ holds at all *actual* epistemic alternatives. This anchoring to the initial context, we believe, gives a more realistic account of the situated character of knowledge and belief.

By way of conclusion, and to give another perspective on the contextualist epistemic semantics proposed here, we may point out an analogy between our treatment of Williamson's paradox and Kamp's influential treatment of the sorites

¹¹ One such idealization concerns the use of fixed margin of errors, irrespective of the quality or intensity of the stimulus in our model.

in Kamp (1981).¹² In this paper, Kamp presents two thought experiments, in which a subject faces a screen divided in small squares, together displaying a gradual transition from green to yellow. In the first thought experiment, the screen is completely visible and the subject is asked of each successive square whether it is green. In the second thought experiment, the screen is covered, and only pairwise adjacent squares are uncovered successively. Kamp supposes that the judgments will vary from one experiment to the other. In the first experiment, in particular, Kamp notes that the extreme squares can serve as “anchor points” to establish a comparison with the focal square, while only the adjacent square can do so in the second experiment. He is led to think for that reason that the judgments will shift more quickly from “Green” to “Non-Green” or “Uncertain” in the first experiment than in the second experiment. Kamp’s semantics is thus an attempt to make each instance of the tolerance principle $\forall n(P(n) \rightarrow P(n+1))$ true, while making the universal claim itself false. His idea is that each instance holds locally, as it were, but not globally.

Interestingly, we get the same kind of effect in CS when we compare (ME) to (KME): in the model of Figure 2, we saw that the conditional $K\neg p_i \rightarrow \neg p_{i+1}$ is universally true; however, its universal generalization $K(K\neg p_i \rightarrow \neg p_{i+1})$, which in combination to (K) and (KK) plays exactly the role of an induction principle, is not valid. If we ponder on the reasons of this failure, we can observe that they depend in an analogous way on a shift in “anchor point” (the reference point in the terminology we used here). Thus, $K\neg p_2 \rightarrow \neg p_3$ holds at 4 because it holds at (4, 4). By contrast, $K(K\neg p_2 \rightarrow \neg p_3)$ does not hold at 4, because its truth now depends on the pair (4, 3), which makes $K\neg p_2 \rightarrow \neg p_3$ false. We thus get the same kind of contextual effect which Kamp imagines in his thought experiments, albeit on a smaller scale: when the perspective point and the reference point coincide, (ME) is valid; when the reference point deviates sufficiently from the perspective point, (ME) can be invalidated.

This analogy, of course, is not meant to undermine the profound conceptual differences between Kamp’s view of the sorites (which preserves a local form of the tolerance principle) and the view of the epistemicists (who simply deny that all instances of the tolerance principle are true). However, our aim here was not to present an overall solution to sorites paradoxes in general. Our aim was rather to untie the knot of one specific sorites, namely the epistemic sorites which results within the metatheory of the epistemicists themselves when margin for error principles—which are used to account for the origin of vagueness—are assumed in combination with introspection principles. Despite this, we are inclined to think that the ingredients of our own contextualist solution to Williamson’s paradox enforce the general paradigm of contextualist approaches to vagueness. It would be interesting, in particular, to examine whether the

¹² Other contextualist approaches include in particular Soames (1999), Graff (2000) and Shapiro (2006).

kind of multi-dimensional semantics we used for an epistemic language can be transposed to a regular first-order language in order to handle the sorites paradox more generally, a problem we leave for further work.

References

- Bonnay, D. and Egré, P. (2006), 'A Non-Standard Semantics for Inexact Knowledge with Introspection', *Proceedings of the ESSLLI 2006 Workshop on Rationality and Knowledge*, S. Artemov and R. Parikh (eds).
- Dokic, J. and Egré, P. (2004), 'Margin for Error and the Transparency of Knowledge', forthcoming in *Synthese*.
- Edgington, D. (1985), 'The Paradox of Knowability', *Mind*, 94: 557–68.
- Egré, P. (forthcoming), 'Reliability, Margin for Error and Self-Knowledge', forthcoming in D. Pritchard and V. Hendricks (eds.), *New Waves in Epistemology*, Ashgate.
- Fara, D. Graff (2000), 'Shifting Sands: An Interest-Relative Theory of Vagueness', *Philosophical Topics*, 28(1): 45–81. Originally published under the name 'Delia Graff'.
- (2002), 'An Anti-Epistemicist Consequence of Margin for Error Semantics for Knowledge', *Philosophy and Phenomenological Research*, 64: 127–42. Originally published under the name 'Delia Graff'.
- Gomez-Torrente, M. (1997), 'Two Problems for an Epistemicist View of Vagueness', *Philosophical Issues*, 8: *Truth*, 237–45.
- Halpern, J. (2004), 'Intransitivity and Vagueness', *Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR 2004)*, 121–9.
- Kamp, H. (1971), 'Formal Properties of "Now"', *Theoria* 37: 227–73.
- (1981), 'The Paradox of the Heap', in U. Mönnich (ed.), *Aspects of Philosophical Logic*, 225–77.
- (1989), 'Demonstratives', in J. Almog *et al.* (eds), *Themes from Kaplan*, Oxford: Oxford University Press.
- Kölbel, M. (forthcoming), 'Vagueness as Semantic', in R. Dietz and S. Moruzzi (eds), *Cuts and Clouds: Essays in the Nature and Logic of Vagueness*, Oxford: Oxford University Press.
- Rabinowicz, W. and Segerberg, K. (1994), 'Actual Truth, Possible Knowledge', *Topoi* 13: 101–15.
- Recanati, F. (2000), *Oratio Recta, Oratio Obliqua*, Cambridge, Mass.: MIT Press.
- Segerberg, K. (1973), 'Two-Dimensional Modal Logic', *Journal of Philosophical Logic*, 2: 77–96.
- Shapiro, S. (2006), *Vagueness in Context*, Oxford and New York: Oxford University Press.
- Soames, S. (1999), *Understanding Truth*, Oxford and New York: Oxford University Press.
- Sorensen, R. (2001), *Vagueness and Contradiction*, Oxford: Clarendon Press.
- Williamson, T. (1992), 'Inexact Knowledge', *Mind*, 101: 217–42.
- (1994), *Vagueness*, London: Routledge.
- (2000), *Knowledge and its Limits*, Oxford: Oxford University Press.

This page intentionally left blank

6

Relativism, Vagueness and What is Said

Manuel García-Carpintero

In a series of recent papers, John MacFarlane (2003, 2005a, 2005b) has formulated a version of truth-relativism, and argued for its application in some cases—future contingents, knowledge attributions and epistemic modals among them. Mark Richard (2004) also defends a version of relativism, which he applies to vagueness-inducing features of the semantics of gradable adjectives; Richard thinks that it provides a model to properly articulate contextualist intuitions about knowledge attributions. On MacFarlane’s characterization, and on Richard’s as I will be interpreting it, truth-relativist claims posit a distinctive kind of context-dependence, to be distinguished from the more familiar one that has been closely studied since Kaplan’s (1989) work in the 1970s. Radical truth-relativists posit the dependence of the evaluation *of an assertion* as true or otherwise on aspects of the context of *the evaluation* itself—in contrast with the context of the assertion.

The main motivation for truth-relativism is to make sense of alleged cases of “faultless disagreement”. These are cases of “substantive disagreement” (Richard, *op. cit.*, 219) between two thinkers, “in the sense that one asserts something that is inconsistent with what the other asserts” (*ibid.*: 218), while, nonetheless, neither of them is at fault, with respect to the fundamental norms to evaluate acts of assertion such as those they are engaged in, and in some unexpected way: it cannot just be that they are talking about different subject-matters, without realizing it, so that in fact there is no real inconsistency; or that both assertions are equally well justified, given the evidence available to the asserters, although at

Research for this paper is supported by the European Science Foundation EUROCORES program “The Origin of Man, Language and Languages” by the Spanish Government’s grant MCYT BFF2002–10164 and research project HUM2004–05609–C02–01, by the DURSI, Generalitat de Catalunya, SGR01–0018, and by a *Distinció de Recerca de la Generalitat, Investigadors Reconeixuts* 2002–2008. I am very grateful to José Díez, Max Kölbel, Dan López de Sa, Teresa Marques and Genoveva Martí for their suggested revisions on previous versions that led to improvements. The paper was presented at the Logos Relativizing Utterance Truth Conference and the Logos Seminar. I thank the audiences there for helpful criticism. I am also indebted to Michael Maudsley for his grammatical revision.

most one is true.¹ Throughout the paper I will further characterize these alleged cases of substantive faultless disagreements, as understood by truth-relativists, and I will argue that they have not given us good reasons to think that there are any; we must do with this negative description for the moment.

My main goal in this paper is to argue against the truth-relativist proposals for gradable adjectives in Richard (2004). I intend my points to apply *mutatis mutandis* to knowledge attributions and epistemic modals, but I will not discuss them here, because they introduce complexities of their own which require independent treatment. An ancillary goal (which, notwithstanding its ancillary character, will take up a good portion of the paper) is to better understand what is at stake in these debates, a matter on which I do not find current literature sufficiently clear. In particular, following the lead of Evans's (1985) earlier critical examination of the issues, I will distinguish a *moderate* from the *radical* version of the view that I understand MacFarlane and Richard defend. Unlike the radical ones, moderate versions relativize contents in unusual ways, but they do not go on to relativize in addition the norms of speech acts with those contents, in particular the truth of assertions with those contents. As a result, I will argue, although they are semantically in well standing, and they are correct in some cases, including in particular the case of gradable adjectives that I will be discussing here, they cannot make real sense of faultless disagreements. Radical versions can, precisely because they do relativize not just contents but the norms of speech acts in addition; but it is unclear whether they are intelligible, precisely on account of that.

The paper is structured in two sections. In the first, I will provide an initial characterization of what is at stake in the truth-relativism debates, and I will also outline what I take to be the fundamental worries about radical truth-relativism. In the second I will present the evidence for a relativist treatment of gradable adjectives, examining the extent to which Richard's proposal appears to be relativist in the radical sense—while, for instance, the similarly sounding one by Egan, Hawthorne, and Weatherson, (2005) counts only as moderate. Then I will propose a more orthodox semantic way of dealing with the evidence for truth-relativism about gradable adjectives, a moderate relativist account.

6.1 VARIETIES OF TRUTH-RELATIVISM

For reasons I hope to make clear presently, in order to properly characterize what is at stake in these debates we should appeal to the vexed notions of *saying* and

¹ Cf. MacFarlane (2007) for his own attempt at stating the phenomenon; the paper also attests to the relevance that the phenomenon has for him, even if he had previously avoided the notion.

what is said; vexed because, as several writers have argued, our practices in making claims involving those notions, and the intuitions that go with them, evince a truly complex picture.² This has been already pointed out in a paper I will be drawing on later: “Unless we give it some special technical meaning, the locution ‘what is said’ is very far from univocal” (Lewis 1980: 41). However, against what some of those writers suggest, there are theoretically significant notions that can be reconstructed out of those practices and intuitions, traditionally understood to be fundamental in semantic theorizing.

There are, in fact, two different theoretically significant notions of *what is said*, related to what I will distinguish as *locutionary* and *illocutionary* sayings: the main interpretative claim I will be defending here is that, in critically examining the relativism debates, we must distinguish *truth* as applying to illocutionary sayings, from its use in applying to locutionary sayings. Moderate truth-relativism has to do with whether or not truth for what is said in the locutionary sense is more or less relativized in unexpected ways, while keeping truth for what is said in the illocutionary sense absolute; radical truth-relativism has to do with whether or not truth *as a feature of illocutions* should in addition be relativized. There is a general point here, that in philosophical discussions of the concept of truth we should keep in clear sight, namely, the distinction between truth as predicated of speech acts such as assertions and truth as predicated of contents. An intuitive consideration is that we do not evaluate as true or false questions, promises or orders, even though they might have the same content as an assertion, say the same in the locutionary sense. I will not discuss further the general issue. What matters to us is that it is both the truth of locutions and illocutions, not just the latter, which is at stake in the relativism debate—and hence that issues exactly paralleling those we will be confronting arise regarding, say, *compliance* for promises and *obedience* for orders.

Grice appealed to the notion of *what is said*, in his struggle to isolate (as he reports in his 1987 “Retrospective Epilogue”, Grice 1989: 359) a “region of signification which has special claims to centrality”. Retrospectively, he identified two features “which I shall call respectively ‘formality’ and ‘dictiveness’”, with seemingly equally strong claims to provide for us a rationally reconstructed interpretation of the initially hazy feature of “centrality”; the first is saliently present in cases “in which the items or situations signified are picked out as such by their falling under the conventional meaning of the signifying expression rather than by some more informal or indirect relationship to the signifying expression”; the second, in “those instances of signification in which what is signified either is, or forms part of, or is specially and appropriately connected with what the signifying expression (or its user) *says* as distinct from

² Notoriously, Cappelen and Lepore (1997) point out this fact. They go on to invoke it, in my view unconvincingly, to defend truth-conditional theories of meaning from standard (in my view, compelling) criticisms.

implies, suggests, hints, or in some other less than fully direct manner conveys” (*ibid.*: 359–60).

I agree with Kent Bach (1994) that in order to get clear about what is semantically central we should distinguish two importantly different notions of *saying* and *what is said*. According to Bach, they can be usefully explicated by appealing to “Austin’s distinction between locutionary and illocutionary acts. Austin, it may be recalled, defined the locutionary act . . . as using certain ‘vocables with a certain sense and reference’ [. . .] That sounds a lot like Grice’s notion of saying, except that for Grice saying something entails meaning it: the verb ‘say’, as Grice uses it, does not mark a level distinct from that marked by such illocutionary verbs as ‘state’ and ‘tell’, but rather functions as a generic illocutionary verb that describes any constative act whose content is made explicit” (*op. cit.*, 143).³ I offer below my own elaboration of the notion of locutionary saying, to which I will henceforth refer as *semantic content* or just *content*. But we must first concentrate on the other notion, illocutionary saying, which, like Bach, I take to be the one Grice intended, the one captured by the two features of *formality* and *dictiveness*.

What is said, in this Gricean illocutionary sense, is intuitively related to what a competent speaker could take to have been made explicit in a literal utterance of a declarative sentence. Unlike Bach, however, I do not think that this should be understood as generic speech act, of which assertion is a paradigmatic species but not the only one—swearing, conjecturing, predicting, reminding and supposing would be others. Instead, I think (but will not defend here) that it is methodologically a better option to think of what is said in the illocutionary sense as the specific act of assertion: the act made by default when literally uttering sentences in the declarative mood. For present purposes, I will also count the mental *acts* (the judgments that assertions express and the belief-states that they in their turn, I take it, are intended to fix or settle on) which overt assertions express as assertions made “in one’s heart”, as it were. It is distinctive of illocutionary acts that they fall under norms, in fact it is constitutive of them to fall under certain norms; for instance, I think that something close to Williamson’s knowledge-rule is the constitutive norm of assertion.⁴ However, *truth*, in the fundamental sense in which I will be using it here, is at least a derivative norm constitutive of assertion.

The notion of *illocutionary saying* is thus pragmatic; but it is related to semantics, in that, as we are characterizing them by relying in part on Grice’s *formality* criterion, assertions are made by default with expressions having certain conventional meanings, and they are determined by those conventional meanings—they are made with sentences of a natural language (English), with semantic features ascribed to them by the best semantic theory for it (ST_E), and

³ Salmon (1991) and Ziff (1972) have argued for similar views.

⁴ García-Carpintero (2004) argues for this.

they have the feature they do in virtue of them. This semantic theory is, I will be assuming following Lewis, “one suited to play a certain role in a systematic restatement of our common knowledge about language . . . the detailed and parochial part—the part that would be different . . . if we were Japanese” (Lewis 1980: 23–4). Now, I understand the job of semantics to be (in accordance with what I think Lewis says here) to give a systematic characterization of what is said in the *locutionary* sense, i.e., of *content*.

More specifically, ST_E will regiment sentences at a *logical form* level (LF), at which it will make explicit syntactic features that contribute to determine semantic features; at that level, sentences by means of which assertions are made will be analyzed into a *force-marker* and a *sentence-radical*. The same sentence-radical occurs with different force-makers (John will be home tomorrow; John, be home tomorrow), and, of course, the other way around. As the more detailed view about semantics outlined later will make clear, what ST_E ascribes to the sentence-radical by means of which an assertion is made is a feature of content, of what is said in the *locutionary* sense. I will refer to those features of sentence-radicals as *sr-contents*, purposefully avoiding the usual ‘proposition’, unless this is convenient to make a point in familiar terms, because, as I will explain, I want to leave it open, as something independent from our discussion, whether or not sr-contents are more relativized than propositions are understood to be.

As I said before, the interpretative claim I will be defending here is that we must distinguish *truth* as applying to *illocutionary* sayings, from its use in applying to *locutionary* sayings. Moderate truth-relativism has to do with whether or not truth for sr-contents is more or less relativized in unexpected ways, while keeping truth for assertions absolute; radical truth-relativism has to do with whether or not truth *as a feature of assertions* should in addition be relativized. The main reason for this is that the motivation presented for relativism by its proponents comes from alleged cases of “faultless disagreement”; but in order to evaluate these cases, we have to do it in the terms that I propose, or corresponding ones.⁵ Relativists grant their opponents that it is *prima facie* difficult to make sense of (not to say that there are) those cases, after all their theory is supposed to be novel and exciting; but this can only be so, I will argue, if the truth whose relativization is posited is the truth of assertions, not just that of the contents that ST_E ascribes to the sentence-radical by means of which they are made.⁶

⁵ This is my motivation for taking a stand on controversial issues, such as those regarding *what is said*, as I am doing in these paragraphs without, of course, being able to provide an acceptable justification for them. To critically examine the debates requires, I will argue, some distinction along the lines of the one I make concerning truth as a property of contents and truth as a norm of speech acts; thus, I might as well invoke the one I think is best.

⁶ The truth-relativism issue arises not just about speech acts, but also about the mental acts they give expression to. I am assuming (and this is why I intend what I want to say to apply also to judgments, understood as “utterances of sentences in one’s heart”) that a distinction corresponding

ST_E, we have said, ascribes to sentences by means of which assertions are made based on what our common knowledge of English tells us about them. According to such common knowledge many such assertions are at most contingently true. An assertion made with an eternal sentence such as 'fluorite structure is a lattice' is true, but it might have been false had the actual world where the assertion occurs been different, in ways that (let us grant, for the sake of the example) all our common knowledge of English tells us, it might have been.⁷ This, given the bond between meaning and modality appreciated after Carnap's (1956) elaboration of Tractarian ideas, leads ST_E to ascribe to sentences (to sentence-radicals, on account of the facts about compositionality presently to be mentioned) functions from possible worlds to truth-values.

Given my main interpretive claims, it will be useful to distinguish henceforth, when the distinction matters, 'i-truth' from 'l-truth'. On account of contingency, and given the facts about the metaphysics of basic contents just rehearsed, ST_E will ascribe to sentence-radicals sr-contents that determine (henceforth, for convenience's sake, which *are*) functions from possible worlds to l-truth-values. These contents, being generalizations concerning distributions of l-truth-values across possible worlds, are thus *possible-world neutral*. No relevant relativity ensues for truth as the fundamental constitutive property of assertions, because we assume that the truth of an assertion made by combining a force-marker with a sentence-radical results from applying the sr-content *to the world where the assertion occurs*. We are assuming the contingency of 'fluorite structure is a lattice'. Take a possible world in which this is false, and imagine that English is spoken there, and that an informed speaker A assertively utters 'fluorite structure is not a lattice', while his counterpart B in another world closer to actuality utters 'fluorite structure is a lattice'. Nobody would want to sustain the possibility of *substantive* "faultless disagreements" supporting controversial philosophical views on such a case, even though A and B are asserting contradictory contents (contents inconsistent with each other), and even though both A's and B's assertions are true.

What explains the nonchalant attitude vis-à-vis this form of truth-relativism that these facts manifest? I submit that the reason is that, in the relevant sense, the proper specifications of the core correctness conditions *of the speech acts* are not really contradictory. For *those* specifications to be really contradictory, in any sense potentially engendering unexpected cases of faultless disagreement, it is not enough that they concern contents that our semantics represents as complementary exhaustive functions from worlds to truth-values; it is in

to the one I will be making here, between truth for contents and truth as a norm of acts, also applies straightforwardly to mental acts. This, I take it, does not require any controversial view about the relative priority of thought and language.

⁷ In any case, this is what I will assume for the sake of our goals here. If one thinks that the structure of compounds such as fluorite is essential to them, one should consider instead some other example. Here we just need any atemporal, contingent truth.

addition required that they be predicated *of one and the same actual world*. I propose to put this for present purposes as follows: although both A and B assert contradictory l-truth conditions, and although both assertions are i-true, A and B are not contradicting themselves because they are not asserting the same i-truth-condition.⁸ An l-truth-condition is what our semantics ascribes to the sentence-radicals by means of which assertions are made; they are shared by assertions and orders, and they are what sentential operators like truth-functions operate on. On account of the facts about contingency, and the relation between meaning and modality, they involve relative l-truth-values: they are non-constant functions from worlds to l-truth-values. However, the i-truth-conditions of assertions are still absolute, because they consist of the application of these l-truth-conditions to a particular possible world where the illocutionary act occurs.⁹

MacFarlane notices this (2005a: 326); perhaps this was also the reason why his earlier (2003) publication on the future contingents issue studiously avoids talk of propositions, officially speaking instead of utterances. Now, however, he thinks that this was wrong, on account of facts about ordinary usage: “there is something a bit odd about calling utterances or assertions, in the ‘act’ sense, true or false at all. We characterize actions as correct or incorrect, but not as true or false” (2005: 322). But, as Austin (1950: 119)—who had as good an ear for common usage as anybody—pointed out, it is at the very least as much far away from common usage to predicate truth of propositions, in the philosophers’ sense.¹⁰ In any case, this is neither here nor there. If there is anything clear that

⁸ It should be clear that—contrary to what Stanley (2005: 132 n.) claims in discussing MacFarlane’s use of this very same point—appreciation of this point does not at all depend on accepting Lewis’s modal realism. Stalnaker’s views on the metaphysics of modality are as contrary to modal realism as it can go, but his conception of “propositional concepts”, crucial to his philosophy of language, depend on evaluating assertions (including those with identical propositional contents) made at different worlds; and this is all I need for the point I make. The point itself could be made modally, without referring to or quantifying over possible worlds: an assertion made with ‘fluorite structure is a lattice’ is true, but it might have been false had the actual world where the assertion occurs been different, in ways that (let us grant, for the sake of the example) for all we are required to know to fully understand the assertion, it might have been. Of course, contingency is very common, and the assertions we are mostly interested in are made in the (one and only) actual world, which explains why we do not need in general to pay attention to the fact that the actual world is strictly speaking part of the truth-conditions for assertions. But it is worth appreciating this, because it helps us to see that the further, not-so-vacuous relativizations of contents we are about to consider are as much independent from the radical relativist claims as this widespread one is.

⁹ The main points made here are in agreement with Austin’s (1950) views about truth; but here we only need to assume them for convenience, to make sense of the relativism debate.

¹⁰ MacFarlane is very clear about the interpretative claim I am making (in fact, reading his work and talking to him greatly helped me to appreciate it myself); now that he wants to present relativism as a view about propositions, he is forced into solecisms which sound to me at least as weird as the ones that lead him to them: he needs to appeal to a distinction between the context of use and the context of assessment *of propositions*, contexts being in his view concrete historical situations. If propositions are atemporal and amodal abstract entities—as in Schiffer’s (1996) characterization—this only makes sense derivatively: the context of a proposition is that of an act in

these debates show, it is that *philosophy is necessary*; quietism is the position or attitude that is more clearly wrong. Commonsense intuitions about these matters are confused; a philosophically well-developed conceptual battery is needed to straighten them out.¹¹ I will use *assertion*, now in the *object* sense (in contrast with the act sense) for i-truth-conditions, as an alternative to make sense of the relativism issue to philosophers' *propositions* and to *content*. In an utterance of an eternal contingent sentence like 'fluorite structure is a lattice', the content expressed has only relative (l-)truth-value. But the assertion itself has nonetheless an absolute (i-)truth-condition, which correlates well with the perception that the disagreement between the inhabitants of different worlds A and B we considered before is not in any way worth noticing: in the present sense, they are making different assertions.

The merits of my interpretative claim can be further appreciated if we introduce now into the picture some of the facts of *indexicality*. As examples such as (1) and (2) below show, ST_E must deal with the fact that our common knowledge of English tells us that which assertions are made with many English sentences (what is said with them, in the Gricean privileged illocutionary sense) depends on facts about the contexts in which instances of the expressions are used to make them; in my view, it depends on properties of the *tokens* themselves, because in many cases it is their properties as physical items (their locations in time and space, their causes and effects) that are relied upon. The dialogues that follow are common tests, which reveal the ordinary perception that the assertions at stake have different contents; I use '?' to mark the perception of infelicity, whatever its source.

(1) (A): I am hungry.

(B): (?) A said/thought that I am hungry.

(2) (A, at 3 a.m.): now it is 3 a.m., October 3, 2000.

(B, at 4 a.m.): now it is not 3 a.m., October 3, 2000.

(A, reacting to B's assertion): That does not contradict what I said // (?) Then I guess I was wrong after all before.

Indexicality and contingency are related; however, as Kaplan (1989) and Lewis (1980) have shown, ST_E must treat them differently: the way in which indexicality makes the truth of assertions dependent on facts about the context in which they take place differs from the way in which contingency makes it so dependent. ST_E will capture this by means of something like Kaplan's (1989)

which the proposition is expressed or appraised. But then it is clearer to discuss directly the features of the relevant acts, as I am suggesting, soberly remembering that, as Fred Dretske once put it, it is the philosopher's fate "to talk funny" at some point or other.

¹¹ Although, I will suggest, in many cases—and we will find instances here—different philosophical proposals, whose proponents maintain fiercely hot debates, are to all relevant purposes notational variants and could do the job equally well. Because of this, I do not claim anything more than instrumental virtues for my own suggestions.

two-dimensional distinction of character and content, dependence of content on *context* and of content's truth-value on *circumstance of evaluation*; i.e., by somehow or other ascribing to sentence radicals semantic values describable by means of the two-dimensional matrices familiar from Stalnaker (1978). These two-dimensional matrices provide in my view a good elaboration for present purposes of Bach's notion of *locutionary* saying. On this proposal, the semantic content of a sentence radical is its full matrix; this is what we should expect for ST_E systematically to deliver.¹² Notice for later use that, as Stalnaker insists, they provide alternative propositions for sr-radicals to convey; in addition to the "horizontal" propositions, there are now the "diagonal" propositions, which, at least in some cases, allow for a reconstruction of Kripke's distinction between metaphysical and epistemic modalities—as in relatively uncontroversial examples such as '(actually, fluorite structure is a lattice) iff fluorite structure is a lattice'.¹³

Now, in spite of the indexicality of 'I' and 'now' revealed by (1) and (2), David Lewis (1979) has notoriously proposed to deal with the phenomenon of what Perry (1979) calls the "essential indexical" by in effect ascribing to sr-contents a further relativity in the *circumstances of evaluation* than that required by contingency. On the present characterization, Lewis's proposal is to relativize the truth of the sr-content for 'I am hungry' not just to worlds, but to individuals and times too (i.e., their sr-contents are (two-dimensional) functions from centered worlds to truth-values).¹⁴ These contents, unlike the traditional ones so far considered, are not just generalizations specifying distributions of truth-values over possible-worlds, but agents and times in addition; they are thus not just possible-world neutral, but also "world-centre" neutral: the values of those parameters are not supposed to be essential or constitutive of them. But still, according to Lewis when one asserts them, what one asserts is the application of those functions to oneself at the time of the assertion (in the world of one's assertion). The truth of assertions remains absolute.

This is as it should be: as in the contingency case, no unexpected case of "faultless disagreement" is perceived if A asserts 'I am hungry' and B asserts 'I am not hungry'. If Lewis is right, B denies the very same sr-content that A asserts, while what A asserts can be true, and so can what B asserts; but there is no substantive faultless disagreement involved here. In the terms I am proposing, the relativization posited by Lewis is moderate, because, in the end, it

¹² This is not strictly speaking correct, but it will do for present purposes. ST_E should also provide some information about the conventional meanings of moods, as well as about conventional implicatures or conventionally signaled presuppositions, and this cannot be accurately represented merely by means of the two-dimensional matrices.

¹³ See García-Carpintero and Macià (2006).

¹⁴ Lewis only considers relativization to an individual, but this is because of his four-dimensionalism about material objects; on a three-dimensional view, his proposal would require the further relativization, which is independently needed if 'now' creates similar problems.

is only l-truth relativization; i-truth remains absolute. In a first-personal assertion such as one made with 'I am hungry', content according to Lewis's proposal only determines relative (l-)truth-value. But the assertion has nonetheless an absolute (i-)truth-condition, which correlates well with the perception that the disagreement between our illustrative inhabitants of different contexts A and B is not in any way of an unexpected sort, even when we focus on Lewisian *de se* contents. Note that I am not at all assuming that Lewis is right about these contents. What I am saying is that whether or not he is right must be decided on the basis of considerations alien to the radical truth-relativism debate.¹⁵ The arguments that can and should be provided against radical truth-relativism do not touch Lewis's proposal, only a moderate truth-relativist proposal.

The final example I would like to consider to hammer home the interpretative proposal I am making comes from debates about temporalism. Let me go back for a moment to the issue of what semantics is about, to mention an aspect not emphasized so far. ST_E must account for the productivity and systematicity that our common knowledge about language manifests, by positing compositional structures. Sentence-radicals by means of which assertions are made are in addition, without any force-marker, immediate constituents of more complex expressions, and must therefore contribute to the semantic properties of the complex expressions of which they are immediate constituents with their semantic values. Relevant here are modal adverbs ('necessarily, fluorite structure is a lattice'), quantifiers ('some professor is such that he loves Mary') and temporal expressions ('Tom Cruise has been to Spain').

On the standard treatment, modals are *operators*; they take a sr-content as so far envisaged, and deliver another. On the standard treatment, quantifiers are not operators; syntactically what they combine with are not at LF sentence-radicals, even if they superficially seem so ('he loves Mary'), but *open* sentence-radicals, sentence-radicals with missing parts in them signaled by variables; they take as values functions from assignments to those variables to sr-contents, yielding sr-contents. There are non-standard views of modals as quantifiers, on which what seem sentence-radicals ('fluorite structure is a lattice') are in fact *open* sentence-radicals that have at LF a world-variable. Regarding temporal expressions, there is no standard view in the literature. Following Prior and others in the tradition of tense logic, Kaplan (1989) takes the operator-view, ascribing to sentence-radicals ('Tom Cruise is in Spain') complex contents akin to those envisaged

¹⁵ Stalnaker (1981: 146–7) offers such considerations, on behalf of an alternative account of the "essential indexical" in terms of diagonal propositions. I am sympathetic to the account, and will be offering a structurally symmetrical one for some cases involving gradable adjectives later. But I doubt that, given the evidence that philosophers can marshal (such as the one Stalnaker appeals to), proposals such as Lewis's, Stalnaker's (and the alternative Reichenbachian reconstruction that Perry offers in the "Retrospective Epilogue" to the 1993 reprint of Perry (1980), which is in fact the one I prefer) are anything more than notational variants.

by Lewis for *de se* attitudes, functions from pairs of worlds and times to truth-values; King (2003), following Higginbotham (1995) and others, argues for the quantifier view.

In defense of his view about temporal expressions, King (2003: 196) appeals in part to considerations akin to those from Evans I invoke later against radical truth-relativism:

[T]hough it seems correct to hold that the things I believe, doubt, etc. can change truth value across worlds (i.e. some of the things I believe are true though they would have been false had the world been different), it is hard to make sense of the idea that the things I believe may change truth value across time and location. What would it be e.g. to believe that the sun is shining, where what I believe is something that varies in truth-value across times and locations in the actual world? It seems clear that when I believe that the sun is shining, I believe something about a particular time and location, so that what I believe precisely does not vary in truth value over times and locations.

King argues here for an asymmetry between the idea that the things we believe change value across worlds, and the idea that they change truth-value across time: the former is correct, the latter makes little sense, he suggests. In support of this asymmetry, he asks a rhetorical question and gives a small argument. I understand the rhetorical question as alluding to what still is in my mind the main reason against truth-relativism on which I will elaborate a little later; i.e., that it does not make sense to relativize *i*-truth, the *i*-truth-conditions of beliefs, because of the very same reasons that it does not make sense to relativize assertions.¹⁶ But the small argument is a non-sequitur. It is true that “when I believe that the sun is shining, I believe something about a particular time and location”, but this is compatible with temporalism, and, in fact, it is exactly what temporalists like Kaplan contend. In the terminology I am using here, this is because temporalism is a view about the relativization of *sr*-contents, i.e., about the relativity to times (in addition to worlds) of *l*-truth-values. It is compatible with this to go on and give perfectly absolute *i*-truth-conditions for the acts and states with those contents, by having them include not just the *sr*-content, but also its application to the time at which the particular assertion occurs.

And it is only natural to do so, because we would not feel any faultless disagreement if A asserts at t_1 ‘it is raining’ and B asserts at t ‘it is not raining’, independently of whether or not there is, as the temporalist claims, a common content (a temporally neutral one) affirmed in the first utterance and denied in the other. The situation is exactly parallel with respect to the relativization of contents to possible worlds. It seems equally clear that, when we believe that snow is white, we believe something about a particular possible world, even if this hardly needs mentioning; for it is conceivable to have the very same belief in

¹⁶ Remember that I am assuming that judgments are conscious episodic acts, prototypically effected by “saying in one’s heart” or inwardly accepting an aural mental image of a sentence, and that beliefs are the kind of state prototypically fixed by those acts.

a green-snow world, and there it would be false. The actual world then goes into what determines the truth-value of the belief, together with the sr-content, in an exactly parallel way to that in which, according to the standard temporalist, the time of the belief, together with its temporalist sr-content, goes into it. Hence, it is a mistake to make this kind of appeal to, in effect, the absoluteness of the truth of beliefs against proposals to, in effect, relativize their contents; for these proposals are perfectly compatible with an absolutist account of truth for beliefs (and assertions expressing them), and, in fact, that is the account that their proponents provide.¹⁷

I would like to insist that I am not claiming at all that proposals to relativize l-truth such as Lewis's about *de se* contents or Kaplan's for temporalism should be accepted without further ado. There are important syntactical and semantic issues, on which the question whether temporal expressions behave like standard quantifiers or rather like operators truly depends (which King (2003) also interestingly invokes in the bulk of his paper). Just after providing the characterization of the goal of a good grammar that I approvingly quoted at the outset, Lewis (1980: 24) goes on to disparage more ambitious goals, such as "that a good grammar should be suited to fit into a psycholinguistic theory that goes beyond our common knowledge and explains the inner mechanisms that make our practice possible". I strongly disagree with him about this, and, because of that, I would find highly relevant to deciding these issues, say, the kind of cross-linguistic information that contemporary linguists bring to bear on these issues.¹⁸ My claim here is merely that it is a confusion to think that they turn on what is here at stake, the philosophical standing of radical truth-relativism.

Let me sum up. There is no significant truth-relativity deriving from contingency. If the operator view of modals is correct, this is because the truth of an assertion of 'fluorite structure is a lattice' depends on its sr-content, and also on the actual world—which, with respect to *i-truth-conditions* plays the role of an "unarticulated constituent". If instead a quantifier view of modals were correct, then the same result ensues, in this case because at LF the sentence has a variable, which when it is asserted behaves like an indexical referring to the world in which the assertion occurs; it still goes into *i-truth-conditions*, but this time as an articulated constituent. There are linguistically important differences between

¹⁷ King (2003: 196) additionally supports his view in that "powerful arguments have been given against the view that the objects of belief are things that change truth-value over time", referring for such powerful arguments to Richard (1981). The argument there turns on alleged difficulties for temporalism which, as Richard himself admits in section 3 of his paper, vanish when a distinction equivalent to the one I am making between relativized sr-contents and unrelativized *i-truth-conditions* is applied to both beliefs and assertions. In the end, he only appears to be questioning temporalism on the basis of an alleged deficit of motivation. This is even clearer in his more recent reappraisal of the argument, Richard (2003: 38–42). I think that Prior's original motivation, close to Lewis's for *de se* and *de nunc* contents, is an adequate one.

¹⁸ Hence, I agree with (and have learned much from) King's criticism (*op. cit.*, 226–8) of Lewis's disparaging attitude towards what he (1980: 32–3, 39) calls the "schmentencine strategy".

the two views, but no debatable truth-relativism is at stake. The same applies, *mutatis mutandis*, to the temporal and the first-personal case.

So, what is the real issue? We have found so far no reason to relativize truth *as a property of assertions*, only considerations to have a more or less relativistic conception of sr-contents. The philosophically substantive issue is whether there is any reason for the radical thesis purporting to relativize truth *as a property of assertions*. This is, I take it, what MacFarlane (2003) argues for, on account of the possibility that the world might be genuinely ontologically indeterministic. If this is so, he contends, an evaluation of my assertion that tomorrow there will be a sea-battle contemporary with my assertion as untrue might be correct, while an evaluation tomorrow *of the very same assertion* as true might be equally correct. In order to make sense of this, he proposes relativizing the truth *of assertions* to “contexts of assessment”, in this case positions in time from which truth is ascribed *to one and the same assertion*. Kölbel (2003) appears to be arguing for a similar view, on account of cases of faultless disagreement, in which a subject makes an assertion, another makes a contradictory one thereby rejecting the first assertion, while both conflicting appraisals of the very same assertion are correct. If truth is, as I am assuming, a fundamental constitutive norm of assertions, this does not make sense unless, as Kölbel proposes, truth for assertion is relativized to, say, personal “perspectives”.

I believe these views are unacceptable, because they clash with intuitions about the nature of intentional acts that they give us no good grounds to abandon. My reasons here are, essentially, Evans’s (1985: 349–50): we are not properly told what we should do, if we are told that orders should be obeyed, or promises complied with, or assertions should be true “from a given perspective”, i.e., as Evans puts it, now correct, but later incorrect, “according to the state of the weather”. This is, I guess, the criticism that King’s rhetorical question in the quotation before hints at, “What would it be e.g. to believe that the sun is shining, where what I believe is something that varies in truth-value across times and locations in the actual world?”.

We take speech acts like assertions as central cases of intentional, rational action. Rational agents should be at least by default in a position to do whatever is necessary to perform them correctly; they at least should have some measure of control about that. But it is unclear how this could be the case, if i-truth were relativized in the way suggested by proponents of i-truth relativism. How can I rationally take responsibility for making correct assertions, if the correctness or otherwise of my assertions depends on parameters set at different contexts of evaluation about which I lack whatever information, in ways over which I have no control? Consider a related case; imagine that A makes a promise uttering ‘I myself will take B to the airport’, and let us assume that there are good reasons to take the relevant content of the uttered sentence (the one contributing to the fulfillment condition of the promise) in this case to be *de se*, a “property” in Lewis’s sense: a function from centered worlds to truth-values.

On a moderate relativist view about *de se* contents, such as Lewis's own proposal, the complete fulfillment condition of the promise consists in the application of this property to A. On an alternative radical relativist view, however, the commitment that A is taking here can only be fully specified at different contexts of assessment, which would determine different subjects to which the *de se* proposition would be applied. But, obviously, it would be irrational for A to commit himself to the fulfillment of the application of such proposition to most other subjects distinct from himself. We could have made the same point with respect to the relevant present time relative to which it is determined the later time at which A commits himself to fulfill the promise, now considering a temporalist view of the content of the uttered sentence.

MacFarlane has considered this type of objection, for the specific case of assertion, advancing a Brandom-inspired view of the norms or goals of assertion, to replace plain absolute truth. The norms to which an asserter subjects his act include things like vindicating its truth *at any relevant context of assessment* in response to legitimate challenges, or retracting it if shown untrue at such a context, or authorizing assertions at similar contexts. In my view, however, the point previously illustrated with the case of promises applies, *mutatis mutandis*, to the case of assertion. But I will not try to establish this here; it is enough for present purposes to keep clearly in sight the distinction between moderate and radical relativist views, because the fact that the latter are rather more doubtful than the former is enough to motivate adopting moderate views as opposed to radical ones, if both can equally account for whatever evidence is adduced.

Evans (1985) already made the distinction I have been insisting upon in his critical examination of the temporalist views by Prior and others. He distinguishes there several interpretations of the temporalist claims; on one of them (Evans's T_1), what temporalists propose to relativize is the truth of assertions. (Evans says 'correctness', instead of truth, I suppose in order to distinguish it from his privileged use of 'truth', which is for contents or Fregean thoughts.) According to them, on this interpretation, "a proper appreciation of the semantic functioning of tense requires us to abandon the idea that particular historical utterances of tensed sentences are assessable, once and for all, as correct or incorrect. Rather, we must acknowledge that the evaluation of particular utterances must change as the world changes" (*op. cit.*, 347). On another interpretation (Evans's T_3)¹⁹ truth-relativism merely "registers the commonplace idea that the evaluation of an utterance depends upon the time the *utterance* is made (and does not vary)" (*ibid.*, 348). This temporalist proposal corresponds to the one I have labelled before as moderate; it relativizes l-truth to times, but not i-truth. Accordingly, Evans's previously mentioned argument is addressed to T_1 ; he does not make any suggestion that such an argument in any way disposes also of T_3 . Against

¹⁹ We do not need to consider here Evans's T_2 .

such a semantic proposal, he instead goes on to provide the sort of considerations that, I suggested before, would be really relevant; i.e., semantic arguments such as whether or not temporal operators would be on the T_3 -temporalist proposal akin to Kaplan's "monsters", whether or not there is something wrong with that, and so on. We will now examine a particular case for relativism, keeping these facts in mind.

6.2 RICHARD'S TRUTH-RELATIVISM FOR GRADABLE ADJECTIVES

Examples such as (3) and (4) below suggest the indexicality of gradable adjectives (adjectives that admit the comparative and superlative, intensifiers like 'much' and 'very' and so on), as the application of the previous tests in (5) and (6) shows:

- (3) (A, assuming Bill is 1,96m tall, informally discussing the height of basketball players): Bill is short.
- (4) (B, assuming the same about Bill, discussing the height of common Spaniards): Bill is not short.
- (5) (B, in his context): (?) A said/thought that Bill is short.
- (6) (A, in B's context, reacting to B's assertion (4)): That does not contradict what I said/(?) Then I guess I was wrong before after all.

The information about differential standards of shortness which accounts for the intuition that different contents are being affirmed and denied in (3) and (4), provided by context in those examples, can in some other cases be explicitly articulated in the uttered sentence:

- (7) (A, as before): Bill is short for a basketball player.
- (8) (B, as before): Bill is not short for a common Spaniard.

The evidence so far can be handled by means of a *l*-truth-relativist proposal, taking into account suggestions about the semantics of gradable adjectives in the literature such as Kennedy (1999) and Kennedy and McNally (2005). On a version of this view, 'short' denotes a measure function—a function from objects to degrees on a scale (in this case, of height), which in its turn is a partial ordering of degrees. This allows a natural account of the truth-conditions of the typical phrases related to gradable adjectives, such as comparative claims like 'Chicago is larger than Rome'; 'more' (or the corresponding suffix) is interpreted so that the sentence is true just in case the degree that the interpretation of the adjective ascribes to the interpretation of 'Chicago' exceeds the degree that it ascribes to the interpretation of the phrase headed by 'than'. Similarly natural truth-conditions are given for sentences such as 'John is 2m tall' and 'Chicago is very large'. On this proposal, it is (paradoxically) a bit more complicated to treat

the positive form. On one way of dealing with this, in the syntax of a sentence such as (3) there is an absolute morpheme that combines with the measure function denoted by 'short' to yield a function from individuals to truth-values. The function takes an individual x to the true just in case the degree of height of that individual is at least as great as the average degree for the reference class, contextually given or, as in the examples (7)–(8), made explicit by the 'for a SN' PP.

This can be represented in the framework described before, by assuming that ST_E will in effect add a further parameter to the circumstances of evaluation for sentences like 'Bill is short' — a "standard"-parameter, corresponding to a class or property, giving specific ways of being short. Relative to the standard appropriate for basketball players, 'short' will have a certain extension with respect to a world and time; relative to the standard appropriate for common Spaniards, another, and so on and so forth. In fact, Lewis (1980) suggests including such a standard in his "indexes" — corresponding to Kaplan's circumstances of evaluation. He calls it a "standard of precision", because he is thinking of more or less strict ways of being hexagonal, or empty; perhaps the phenomena are not strictly speaking the same, but in principle the case appears to be exactly analogous.²⁰ In fact, we could have presented the phenomenon by contrasting (3) with (9) instead; and then (10) appears to articulate the relevant standard:

- (9) (C, assuming Bill is 1,96m tall, discussing the height of basketball players with basketball coaches, after concluding from careful study that players shorter than 1,9587 are not ideal for the game): Bill is not short.
- (10) (C, as before): Bill is not short for a basketball player when applying such-and-such strict standards.

These examples do not thus appear to introduce any truth-relativism of the radical kind. We can treat them as just suggested, by further relativizing the l -truth-values of sr -contents to (more or less precise) standards of comparison, thus in effect counting expressions of the form *for a SN* as operators; when gradable adjectives occur without them, as in 'Bill is short', a standard is provided by the context of the assertion, so that the i -truth-conditions remain, for all these expressions call for, absolute.²¹ A similar treatment can be given to Travis's

²⁰ Kennedy and McNally (2005) argue, on the basis of interesting evidence, for a semantic difference between "relative" gradable adjectives such as 'tall' and 'rich' and "absolute" gradable adjectives such as 'empty' or 'impure'; and Kennedy (MS, sec. 3.2.1) argues correspondingly for a distinction between the vagueness of the former and the mere imprecision of the latter in some cases. I think that granting such a distinction would not affect any of the arguments in this paper.

²¹ Alternatively, we could count sentences like 'Bill is short' as in fact open sentence-radicals at LF; the standard provided by context would then be articulated as the referent of a hidden indexical. Jason Stanley has a well-known general argument to this effect, based on independent semantic considerations. I am skeptical, for reasons close to those in Hawthorne (2004: 98–100 n.). Nevertheless, as I argued in the previous section, nothing relevant for the present debate hinges on this.

(1997) well-known example: an utterance *u* of ‘these leaves are green’, with *w* the actual world where *u* is made, *w* being such that the leaves demonstrated at *u* are naturally red leaves that have been painted green. Intuitively, if *u* is made for the benefit of a photographer, the leaves would (to use Travis’s words) *count as* green, they would be in the extension of ‘is green’ w.r.t. *w*; however, if *u* is made for the benefit of a botanist seeking green leaves for a study of green-leaf chemistry, the leaves do not *count as* green. To deal with these cases, we can ascribe to the sentence a sr-content such that, when evaluated w.r.t. *w* and a “count as” parameter selecting things that are green from a photographer’s perspective, delivers truth; when evaluated w.r.t. *w* again, but now together with a “count as” parameter selecting things that are green from a botanist’s perspective, it delivers falsehood. The standards in circumstances of evaluation will determine what counts for a three-dimensional entity with many potentially layers and surfaces to be a certain color from a given perspective—in other words, which of the many potentially relevant parts count, as suggested by Szabó (2001).²² The context of the utterance will select one, when none is explicitly indicated in the sentence.

However, a closer look at cases such as (3) and (9) raises doubts; for, if we apply to them the tests for indexicality applied before, the intuitive results contrast sharply with those in (5)–(6):

- (11) (C, in his context): A said/thought that Bill is short.
- (12) (A, in C’s context, reacting to C’s assertion (11): (?) That does not contradict what I said//Then I guess I was wrong before after all.

To account for this, Richard (2004) distinguishes between two kinds of context-dependence for gradable adjectives:

‘tall’ is context sensitive insofar as it requires a reference class—something is tall for a tree/building/two year old; nothing is tall *simpliciter*. Other adjectives, while perhaps not requiring a reference class in order to get assigned an extension, clearly accept one. I may call a field flat and obviously be comparing it to the fields in Bucks County, or call a pond flat with the clear intent that we evaluate what I say relative to the class of naturally occurring large bodies of water. A quite different sort of context sensitivity results from the fact that (even after a reference class is determined) candidates for the extension of a term such as ‘flat’ or ‘rich’ are ordered by the degree of the relevant property (flatness, wealth), with the extension of the term being determined by “how one draws the line” within the ordering—such line drawing being something that can be done differently in different contexts (*op. cit.*: 237).

Other writers who argue in favor of truth-relativism,²³ have also contended that there are two forms of relativism in gradable adjectives, the relativism

²² MacFarlane suggests this account of the case in a footnote to his manuscript “Semantic Minimalism and Non-Indexical Contextualism”, commenting on Stefano Predelli’s (2005) proposals. I have greatly benefited from discussion with Predelli.

²³ I will later claim that they are in fact arguing for a form of moderate relativism; they explicitly commit themselves to the absoluteness (as far as their argument goes) of i-truth.

to a comparison standard, which lends itself to a straightforward contextualist treatment, and another one not so tractable. Thus, Egan, Hawthorne, and Weatherston (2005: 150) discuss this example:

- (13) (Ant Z): He's huge (said of 5f 3, 141lb NBA player Muggsy Bogues).
(Andy): Ant Z thinks that Muggsy's huge.

This is what they have to say about the example: "... the report is accurate, or at least extremely natural. And ... it would have been inappropriate for the reporter to continue 'and he's right'. But crucially, ... [it is not] *clear* that the original speaker made a mistake. ... From Ant Z's perspective, Muggsy Bogues is huge. We assume here, a little controversially, that there is *a* use of comparative adjectives that is not relativised to a comparison class, but rather to a perspective. Ant Z does not say that Muggsy is huge for a human, or for a NBA player, but just relative to him. And he's right. Even Muggsy is huge relative to an ant." The example that Richard is actually discussing is in fact quite similar: "Suppose, to take an example, that Mary wins a million dollar lottery. Didi is impressed, and remarks to a friend 'Mary's rich.' Naomi, for whom a million dollars is not really all that much, remarks in a conversation disjoint from Didi's, 'Mary is not rich at all'. Suppose the salient comparison class is the same in both cases. (Both are taking New Yorkers to be the relevant fields of comparison.) Suppose that there is no difference between the two conversations in the point of assessing people as rich or otherwise" (*op. cit.*: 218). Graff (2000: 55–6) has rightly insisted on the importance of such sorts of cases for understanding the phenomenon of vagueness; she also thinks that they involve a different kind of context-sensitivity than the relativity to a comparison class, to context-shifting significances of items relative to the interests of speakers.

It is to deal with the second type of context-sensitivity in gradable adjectives—the sensitivity to "how one draws the line", giving rise to the intuitions that (11)–(12) reveal—that Richard proposes a relativist theory. It is a radical relativist view by my lights, because it is not just a matter of appealing to the relativity to a "standard" parameter in sr-contents already countenanced before, while the truth-conditions for assertions is given not just by the sr-content, but by its application to a particular standard or class thereof given together with the assertion. Richard rejects this on the basis of the contrasting intuitions in (11)–(12). What he suggests is, instead, that the truth *of the relevant assertions* is relativized to "contexts of evaluation" providing different standards of precision. A's assertion (3) would be true from his own "precision-perspective", and false from C's:

What are we to say, then, about gradable adjectives such as 'rich'? Well, we *could* take appearances at face value. ... Didi and Naomi disagree. So there is something which Didi says and Naomi denies. Within the confines of each woman's conversation, each use of 'is rich' is correct. So Didi says something true when she utters 'Mary is rich', Naomi something true when she utters the sentence's denial. This is consistent with the

two disagreeing over the truth of a single claim, if the truth of the claim may be relative, so that it may be “true for Didi, but not for Naomi” (*op. cit.*: 225).

Although it is difficult to be totally sure, on account of the ambiguities on which I have been insisting in all those phrases—‘what Didi says’, ‘what Naomi denies’, ‘a single claim’—it appears that Richard is propounding here the radical variety of truth-relativism that Evans was objecting to with the argument mentioned at the end of the previous section, not the milder one for which he reserved his properly semantic critical considerations. The impression is reinforced when Richard gives us his views about *truth*, which he does in order to avoid a potentially threatening inconsistency. If *we* say that Didi spoke truly, given “the banality that whoever utters ‘Mary is rich’ says that Mary is rich” (*op. cit.*, 233), we appear to be committed to: it is true that Mary is rich, and then to: Mary is rich. But we also want to say that Naomi spoke truly, which, on the basis of a similar “banality”, together with one more application of the truth equivalence schema, ends up having us committed to: Mary is not rich. “To respond to this, the contextualist must, I think, clearheadedly embrace relativism. . . . Once the contextualist accepts the banality . . . he *must* use a relativized notion of truth to formulate contextualism” (*op. cit.*: 232–3).²⁴

Richard appears to assume the following modified disquotational principle (T):²⁵

(T) If an utterance (assertion) *u* says that *P*, then *u* is true iff *P*.

If we apply (T) to Didi’s utterance, from Didi’s perspective, it follows that ‘true’ applies to it; if we apply (T) to the very same utterance, from Naomi’s perspective, the result is that ‘true’ does not apply to it. In sum, as Richard suggests, ‘true’ has the same kind of relativity that ‘green’ has. Since he does not provide any other notion to evaluate assertions, it seems clear that what he is defending is genuine truth-relativism as it was characterized before.

Egan, Hawthorne, and Weatherson (2005) differ at this crucial juncture, separating moderate and radical forms of truth-relativism. They also find some use for a disquotational notion of truth satisfying (T) (*op. cit.*: 155), and they point out its relativistic character. However, they explicitly indicate about this relativist notion of truth that “we shouldn’t restate the norms of assertion in terms of it, because that will lead to the appropriateness of assertion being oddly relativised” (*op. cit.*: 160). And previously they have worried, “it is not obvious how to apply some of [the norms giving conditions for correct assertion] if utterance truth is contextually relative, because one of the norms is that one should say only what is true”. Then they go on to reject the suggestion

²⁴ Note that Richard uses ‘contextualism’ for ‘relativism’, given that he takes himself to be offering the best philosophical reconstruction of contextualism about knowledge-attributions.

²⁵ For reasons to prefer such a conditionalized version of the disquotational scheme to the more standard version, see Williamson (1994: 187–98).

that “utterance appropriateness is, like utterance truth, relative to a context of evaluation”, and contend instead that “the correct norm is that one should only say something that’s true when evaluated in the context you are in” (*op. cit.*: 153).

Hence by my lights their proposal belongs in the moderate camp. Truth for assertions (what they call ‘appropriateness’) is absolute, and depends on applying the relativist l-truth-conditions they ascribe to utterances to the context of the assertion. We should not take too seriously the rhetoric of “faultless disagreement” about cases such as (13) which I quoted after the example (in fact, they merely say that “[it is not] *clear* that the original speaker made a mistake”). For, put in my terms, what Ant Z i-says differs from what Andy would i-say, were he to add to his report “and he’s right”. It is just that Egan, Hawthorne, and Weatherson’s semantics allows them to recover a common l-relativist content, akin to the temporalist common content for utterances of ‘it is raining’ in different contexts. There is nothing wrong with this (and the proposal I will be making to deal with intuitions such as those in (11)–(12) is structurally similar). But this is no relativism of the offending sort, unlike what Richard and MacFarlane propound.

As I have been insisting, this is no mere terminological dispute. The issue is this: there is a conception of truth-relativism—the one to which Evans was objecting—which is *not* the one content-relativists like Lewis about *de se* attitudes, or Kaplan about tenses, are making; Evans’s argument does not touch this view (and he was not under the illusion that it does). Now, proposals such as Egan, Hawthorne, and Weatherson’s for ‘huge’ (and for epistemic modals) have exactly the same shape structurally. This is obscured by their talk of relativizing “utterance truth”, and by their claiming agreement with MacFarlane; in fact, as we have seen, they carefully distinguish correctness or appropriateness for assertions, and for that (exactly as the Kaplanian temporalist would do for “utterances”) they give absolute conditions. Hence, they are not advancing a relativism of the kind that people find problematic for Evans’s sort of reason.

Coherently with the interpretation of his view that I am assuming, Richard, like MacFarlane, does feel the threat of Evans’s argument, and the need to say something in response. As I said before, MacFarlane (2003: 334–5; 2005: 334–8) resorts to a more convoluted conception of the norms of assertion—based on the one put forth by Brandom—to confront the difficulty. I have already expressed doubts about the adequacy of this move; but, as I said, for present purposes it is enough to make out the extra burden on radical relativism. I will take up Richard’s response presently, but let me first outline the alternative moderate form of relativism that I would like to suggest to deal with the intuitions he relies on.

For easiness of presentation, let me refer to the intuitions of accurate cross-contextual report concerning indirect discourse reports such as (11), apparently

distinguishing gradable adjectives (with respect to the dimension that the writers we are discussing distinguish from relativity to a comparison class) from other indexicals, as *CCR*; and let me refer to the equally distinguishing intuitions about expressions of disagreement and retractions such as (12), as *D&R*. These intuitions conflict with the purely contextualist view I would like to defend, on which there is no relevant semantic difference between the relativity of the content-contribution of gradable adjectives to ordinary comparison standards, and their relativity to “perspectives” (Egan, Hawthorne, and Weatherson), to “how one draws the line” (Richard) or “interests” (Graff). There are two strategies that contextualists could pursue: (a) reject the very existence of intuitions *CCR* and *D&R*, and (b) reject their authority, ascribing “semantic blindness” to speakers, accounting for it on the basis of specific properties of these cases. My proposal is to combine the two strategies.

As for the first strategy: I think that the intuitions in (12) (and, *mutatis mutandis*, those in (11)) should waver if A is allowed to explain himself better, especially if he is reflective enough and has come up with the proper tools:

- (14) (A, in C’s context, reacting to C’s assertion (11)): That does not contradict what I said; I was just saying that Bill is short for a basketball player *on rough estimates for the purposes of coffee talk*; we were not contemplating your levels of precision; thus I was not wrong.

Similar maneuvers are available for Richard’s ‘rich’ example, and for Egan, Hawthorne, and Weatherson’s ‘huge’ example.²⁶ I have not taken polls, but apparently my intuitions here at least agree with those of other philosophers, even those otherwise sympathetic to relativism.²⁷

Still, I think that this is only half of the story; for the other half, we need to appeal to the second strategy. I agree that the content-relativity to comparison *kinds*, such as that of humans or that of basketball players, is much more salient to ordinary intuitions than that to “perspectives”, “interests” or “ways of drawing the line”. I also agree with Graff and Richard that the latter is more closely related to the phenomenon of vagueness than the former.

On the classical assumptions we have been taking for granted so far, an utterance of ‘A is rich (for an N)’ has a content such that the predicate divides the domain (in any possible world) into two mutually incompatible and jointly exhaustive classes. On the previously outlined semantics based on Kennedy’s

²⁶ The same applies to Graff’s (2000: 55–6) ‘bald’ and ‘blue book’ examples, which she offers to justify the context-sensitivity of gradable adjectives to something other than a comparison kind. I agree with Stanley (2005: 59, n. 8) that the examples do not ultimately contradict that, strictly speaking, there is just a form of context-sensitivity, to features or properties establishing comparison standards more or less difficult to articulate explicitly.

²⁷ Stanley (2005: 55–6) provides a similar reply. Hawthorne (2004: 104–7), to whom Richard expresses thanks for helping him appreciating “the importance of thinking about this sort of example” (*op. cit.*: 247 n. 1), appears also to concur that the intuitions *CCR* and *D&R* can be resisted in this way. MacFarlane (2005b: 214 n.) also agrees.

work, how this divide is effected depends on how the line in the scale of degrees of height is drawn, i.e., on what counts as greater than the average for the *N* in the context. Graff (2000) provides a contextualist-cum-epistemicist account of vagueness; on her view, what matters for the interpretation of the absolute morpheme (on the proposal sketched before) is not the average degree of the property (shortness, or shortness-for-basketball players), but what counts as significant for the conversationists' purposes in the context. Thus, with respect to a context *c*, 'short' would apply to *a* just in case the degree that the interpretation of the adjective ascribes to *a* exceeds the degree significant in *c*.²⁸

Now, I am prepared to concede (if this is a concession) to epistemicists like Graff that this is an accurate representation of what goes on in some very unusual contexts; imagine a use of 'tall for a flatlander' with respect to a cartoon-like situation in which most of the flatlanders in the domain are bidimensional beings without any height at all, such that having any height counts as satisfying the predicate.²⁹ However, the condition is not met in most ordinary contexts. We simply do not care to take decisions that are required, for the predicate to properly make a semantic contribution. Fortunately, the relativist proposals we have been considering before provide alternative contents available for assertion and other speech acts. Instead of assuming the context to fix just one of the standards (the "significant" one), we may allow it (modulo higher-order vagueness, which I put aside here) to leave it open which one among those in a more or less extended class counts as significant in the context. Then we would in effect have resort to a supervaluationist semantics to give us the truth-conditions for uses of the positive form in context.³⁰ An object *a* is in the extension of 'short', or 'short for a basketball player', just in case the degree that the interpretation of the adjectival phrase ascribes to *a* exceeds all the standards significant in the context. It is to contents of this kind, I submit, that the CCR and D&R intuitions respond. On the outlined proposal, the content that is being ascribed to 'Bill is short' in (11) is, essentially, that Bill falls under the extension of 'short' in any admissible precisification—admissible, that is, given the contextual information about comparison standards; and this is the claim that A retracts, in view of C's assertion (9).

On the alternative I suggest, thus, the intuitions in (11)–(12) are excusably confused. The content most usually i-said in, say, (3), should not be retracted just on the basis of accepting (9); and I think we lose the initial tendency to retract ourselves once allowed sufficient reflection. On the other hand, our semantic

²⁸ See Kennedy (MS) for an interesting and detailed elaboration of Graff's view.

²⁹ I owe these examples to Mario Gómez Torrente.

³⁰ Fine (1975) offers the classical supervaluationist account. Alternative accounts of vagueness might well find alternative ways of accounting for the anti-contextualist intuitions I am prepared to grant, but I need to offer a sufficiently specific one and this is the one I take to be correct. See García-Carpintero (2007) for elaboration and Kölbel (forthcoming) for a similar moderate relativist account of vagueness.

system provides an alternative content to be ascribed to (3), for (9) to contradict; but this is not the ordinary content that C appears to contradict.

This proposal is very close to one that Richard considers, and rejects:

Vague expressions may be made precise, for conversational purposes, so long as those in the conversation can agree on how to do so. If conversants can't agree, the expressions retain their vagueness. . . neither Didi nor Naomi's use of 'rich' is prohibited by the meaning of 'rich'. That is: in some contexts, 'rich' has an extension in accord with Didi's usage; in others, it has one in accord with Naomi's. . . The substantive dispute in [a situation in which the two turn from their respective conversations to an argument with each other over Mary's status] is something like this: whether for Didi and Naomi's purposes, the best way of sharpening the predicate 'rich' involves putting Mary in its extension (*op. cit.*: 221).

In summary, on the present proposal the intuitions in (11) and (12) do not support the existence of genuine cases of faultless disagreement, and therefore do not call for any truth-relativization of the problematic kind. On the first strategy, the speakers are not contradicting each other; the content one of them denies is not the same content the other asserts. On the second strategy, the denied content is the same one affirmed, but at least one of the speakers is at fault, in the assumptions they made about the standards prevailing in the relevant context. The latter strategy involves a form of relativization, but it is one of the moderate kind, unquestioned by Evans's argument against relativizing truth as a norm of assertions.

Richard rejects proposals such as this on the basis of two considerations that I do not find convincing (*op. cit.*: 222). The first is that this does not allow us to characterize what is going on in cases, such as the original Didi–Naomi example, in which the two speakers are engaged in independent conversations, and we are supposed to uniformly report the contents that one asserts and the other deny while making both assertions correct. The anti-relativist accepts this, for he insists (on the basis of Evans's argument) that there is no reason to accept this as a constraint on acceptable accounts of the cases. The second is an appeal to phenomenology; in cases like those discussed, "speakers do not feel that they are sharpening a vague concept in applying it; indeed, they would challenge the claim that they are sharpening indeterminate usage". But, if this is really a feeling that speakers have, I think the anti-relativist is within his rights not to grant it any authority (on the same grounds as before).

I said before that Richard does consider Evans's argument. But what he says in response does not help, as far as I can see:

The picture of assertion behind this objection is incomplete. Suppose that I assertively utter 'Mary is rich', when it is not antecedently settled for conversational purposes whether Mary is in the term's extension. My statement, that Mary is rich, is as much an invitation to look at things in a certain way, as it is a representation of how things are. In saying that Mary is rich, I am inviting you to think of being rich in such a way that Mary counts as rich. If you accept my invitation—that is, if you don't demur, and carry on the

conversation—that sets the standards for wealth, for the purposes of the conversation, so as to make what I say true. It is this idea—that an assertion can be as much an invitation to conceptualize things in a certain way, as a representation of how things are—that is missing from the picture of assertion on which the objection rests.

I take it that an “invitation” is a speech act akin to a stipulation, or a proposal for an agreement. If one combines the view about the vagueness in gradable adjectives suggested before with Stalnaker’s view of assertion, which I find congenial, assertions such as one of ‘Mary is rich’ do involve such invitations. On Stalnaker’s dynamic picture, an assertion, if unchallenged, changes the “context set” or “scoreboard”, by being added to it. If the assertion involves the application of a vague predicate to a borderline case of application (by which here I mean a case such that the conventional meaning of the predicate does not fix whether or not the predicate applies to it), and the Stalnakerian nature of assertion is at least tacitly known by the speakers, the assertion is at the same time a proposal or “invitation” to conceive of the standards of comparison significant in the context as all of them allowing for Mary to count as rich.³¹

However, this by itself does not help to provide an acceptable response to Evans’s concern. The question is, how does the possibility that the “invitation” is declined affect the conditions for the evaluation of the assertion as true or otherwise? If, as we have seen that Richard appears to maintain, this still requires a relativization of the problematic form, so that the assertion cannot be evaluated as true or otherwise absolutely, Evans’s concern remain unaffected. As I put it before, how could it be rational for me to assert ‘Mary is rich’, if whether or not the act I am carrying out is correct can only be decided relative to what is the case at “contexts of evaluation” (presumably, relative to whether or not the invitation implicit in my assertion is accepted in them), over which I lack information and control?

On the proposal I have made, on the other hand, the fact that my invitation is rejected can be taken in two different ways. I can still maintain that I have made an assertion with a content involving my own take on the relevant standards of comparison, and thus I have made a correct assertion; or it rather shows that I have made an assertion with a content concerning standards that I had not contemplated, and then I have made an incorrect assertion.³² That the latter option is available might be objectionable on the basis of the purely semantic kind of considerations that Evans contemplates concerning his T_3 reconstruction of temporalism, but not on the basis of the arguments he provides against his

³¹ Barker (2002) discusses assertions that can only be sensibly taken as such “invitations”, and provides a dynamic model for them which, although it uses a semantics for gradable adjectives slightly different from the one presupposed here, can be easily adapted to our suggestions.

³² An additional worry that I have with Richard’s suggestion is that I cannot really see how he can think that his proposal survives the “phenomenological” argument he provides against the one I take to be essentially the same as the I-truth-relativistic one I have made. Seeing a potential invitation in the assertions we are considering appears to require sensitivity to the fact that the relevant concepts allow for different sharpenings.

T₁ interpretation, which, applied to the present case, appears to be Richard's. This is independent of the issue of whether or not the assertion involves also the invitation, and therefore the appeal to that, by itself, does not help.

References

- Austin, John (1950): "Truth", *Proceedings of the Aristotelian Society*, supp. vol. xxiv; also in *Philosophical Papers*, J. O. Urmson and G. J. Warnock (eds.), Oxford: Oxford University Press, 117–33, to which I refer.
- Bach, Kent (1994): "Conversational Implicatures", *Mind and Language* 9, 124–62.
- (2001): "You Don't Say?", *Synthese* 128, 15–44.
- Barker, Chris (2002): "The Dynamic of Vagueness", *Linguistics and Philosophy* 25, 1–36.
- Cappelen, H. and Lepore, E. (1997): "On an Alleged Connection between Indirect Speech and the Theory of Meaning", *Mind and Language* 12, 278–96.
- Carnap, Rudolf (1956): *Meaning and Necessity*. Chicago: University of Chicago Press.
- Egan, A., Hawthorne, J., and Weatherson, B. (2005): "Epistemic Modals in Context", in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*, Oxford: Oxford University Press, 131–68.
- Evans, Gareth (1985): "Does Tense Logic Rest upon a Mistake?" In his *Collected Papers*, Oxford: Clarendon Press, 343–63.
- Fine, Kit (1975): "Vagueness", *Synthese* 30, 265–300.
- García-Carpintero, Manuel (2004): "Assertion and the Semantics of Force-Markers", in C. Bianchi (ed.), *The Semantics/Pragmatics Distinction*, CSLI Lecture Notes, Chicago: University of Chicago Press, 133–66.
- (2007): "Bivalence and What Is Said", *Dialectica* 61, 167–90.
- and Macià, Josep (2006): "Introduction", in M. García-Carpintero and J. Macià (eds.), *Two-Dimensional Semantics*, Oxford: Oxford University Press.
- Graff, Delia (2000): "Shifting Sands: An Interest-Relative Theory of Vagueness," *Philosophical Topics* 28, 45–81.
- Grice, H. P. (1989): *Studies in The Ways of Words*, Cambridge, Mass.: Harvard University Press.
- Hawthorne, John (2004): *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Higginbotham, James (1995): "Tensed Thoughts", *Mind and Language* 10, 226–49.
- Kaplan, David (1989): "Demonstratives", in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan*, Oxford: Oxford University Press, 481–563.
- Kennedy, Christopher (1999): *Projecting The Adjective: The Syntax and Semantics of Gradability and Comparison*, New York: Garland.
- (MS, 25/6/05): "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Predicates", at <http://home.uchicago.edu/~ck0/prose.html>.
- and McNally, L. (2005): "Scale Structure and the Semantic Typology of Gradable Predicates", *Language* 81(2), 1–37.
- King, Jeffrey (2003): "Tense, Modality, and Semantic Values", *Philosophical Perspectives*, vol. 17, *Language and Philosophical Linguistics*, 195–245.

- Kölbel, Max (forthcoming): "Vagueness as Semantic", in Richard Dietz and Sebastiano Moruzzi (eds.), *Cuts and Clouds: Essays in the Nature and Logic of Vagueness*, Oxford: Oxford University Press.
- Lewis, David (1979): "Attitudes *De Dicto* and *De Se*", *Philosophical Review* 88, 513–43. Also in D. Lewis, *Philosophical Papers*, vol. 1, Oxford: Oxford University Press.
- (1980): "Index, Context, and Content", in S. Kanger and S. Öhman (eds.), *Philosophy and Grammar*, Dordrecht: Reidel. Also in David Lewis, *Papers in Philosophical Logic*, Cambridge: Cambridge University Press, 21–44, from which I quote.
- MacFarlane, John (2003): "Future Contingents and Relative Truth", *Philosophical Quarterly* 53, 321–36.
- (2005a): "Making Sense of Relative Truth", *Proceedings of the Aristotelian Society* 105, 321–39.
- (2005b): "The Assessment Sensitivity of Knowledge Attributions", in T. S. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology* 1, Oxford: Oxford University Press, 197–233.
- (2007): "Relativism and Disagreement", *Philosophical Studies* 132, 17–31.
- Perry, John (1979): "The Problem of the Essential Indexical", *Noûs* 13, 3–21. Reprinted with a "Retrospective Epilogue" in his *The Problem of the Essential Indexical and Other Essays*, Oxford: Oxford University Press, 1993, 33–52.
- Predelli, Stefano (2005): *Contexts: Meaning, Truth, and the Use of Language*, Oxford: Oxford University Press.
- Richard, Mark (1981): "Temporalism and Eternalism", *Philosophical Studies* 39, 1–13.
- (2003): "Introduction to Part I", in A. Jokic and Q. Smith (eds.), *Time, Tense, and Reference*, Cambridge, Mass.: MIT Press, 25–45.
- (2004): "Contextualism and Relativism", *Philosophical Studies* 119, 215–42.
- Salmon, Nathan (1991): "The Pragmatic Fallacy", *Philosophical Studies* 63, 83–91.
- Schiffer, Stephen (1996): "Language-Created Language-Independent Entities", *Philosophical Topics* 24, 149–67.
- Sorensen, Roy (2000): "Direct Reference and Vague Identity", *Philosophical Topics* 28, 177–94.
- Stalnaker, Robert (1978): "Assertion", in P. Cole (ed.), *Syntax and Semantics* 9, New York: Academic Press, 315–32. Reprinted in R. Stalnaker, *Context and Content*, Oxford: Oxford University Press, 1999, 78–95.
- (1981): "Indexical Belief", *Synthese* 49, 129–51. Reprinted in R. Stalnaker, *Context and Content*, Oxford: Oxford University Press, 1999, 130–49, to which I refer.
- Stanley, Jason (2005): *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Szabó, Zoltán G. (2001): "Adjectives in Context", in I. Kenesei and R. M. Harnish (eds.), *Perspectives on Semantics, Pragmatics, and Discourse*, Amsterdam: John Benjamins, 119–46.
- Travis, Charles (1997): "Pragmatics", in C. Wright and B. Hale (eds.), *A Companion to the Philosophy of Language*, Oxford: Blackwell, 87–107.
- Williamson, Timothy (1994): *Vagueness*, London: Routledge.
- (2000): *Knowledge and Its Limits*, New York: Oxford University Press.
- Ziff, Paul (1972): "What Is Said", in D. Davidson and G. Harman (eds.), *Semantics of Natural Language*, Dordrecht: D. Reidel Publishing Company, 709–21.

PART II

THE METAPHYSICAL SIGNIFICANCE OF RELATIVISM

This page intentionally left blank

Relativism about Truth Itself: Haphazard Thoughts about the Very Idea

Crispin Wright

πάντων χρημάτων μέτρον ἐστὶν ἄνθρωπος,
τῶν μὲν ὄντων ὥς ἔστιν, τῶν δὲ οὐκ ὄντων
ὥς οὐκ ἔστιν¹

(Protagoras DK80b1)

The setting of relativistic ideas about truth in the general style of semantic-theoretic apparatus pioneered by Lewis, Kaplan and others has persuaded many that they should at least be taken seriously as competition in the space of explanatory linguistic theory, a type of view which properly formulated, may offer an at least coherent—and indeed, in the view of some, a superior—account of certain salient linguistic data manifest in, for example, discourse about epistemic modals, about knowledge and about matters of taste and value, and may also offer the prospect of a coherent regimentation of the Aristotelian “Open Future” (along with, perhaps, the Dummettian ‘anti-real’ past.) My main purpose here to enter a reminder of certain underlying philosophical issues about relativism—about its metaphysical coherence, its metasemantic obligations, and the apparent limitations of the kind of local linguistic evidence which contemporary proponents have adduced in its favour—of which there is a risk that its apparent rehabilitation in rigorous semantic dress may encourage neglect.

7.1 RELATIVISM INTUITIVELY UNDERSTOOD

Relativism, as a view about some discourse or subject matter, is one of the oldest of philosophical stances. Many, not merely in the academy but more widely,

¹ “Of all things the measure is man: on the one hand of things that are, that they are; on the other of things that are not, that they are not.”

would be willing to describe themselves as “relativists” about certain selected subject matters: about matters of taste or, more generally, about the substance of what I have elsewhere called “disputes of inclination”.² Post-Einsteinian physics has made a received view of relativism about motion and simultaneity. Is there a common idea between, say, ethical relativism and motion relativism? A widely acknowledged manifestation of an intuitive relativistic tendency is a willingness to take seriously the idea of faultless disagreement—disagreement about a shared content where not merely need there be no flaw of evidence or procedure on the part of the disputants but where, as it is supposed, *nobody need be wrong*.³ I do believe that this is an ineliminable aspect of anything plausibly purporting to be relativism in the Protagorean tradition. However it is, I think, more illuminating to place initial emphasis on a different and, I suggest, more basic indicator of ‘folk-relativism’, and then work upwards, so to speak, to try to make space for faultless disagreement.

The more basic indicator of folk-relativism is its characteristic expression in ordinary discourse in cognates of the idiom: “There is no such thing as simply being Φ ”. After Einstein, there is, for example, no such thing as simply *being in motion*: things move, or not, in relation to each other (or, more sophisticated, in relation to a frame of reference). More intuitively, there is no such thing as an object’s simply *looking red*: something looks red, or not, in particular viewing circumstances. Indeed there is presumably no such thing as simply *being useless*: something is useless or not in relation to a particular purpose. More controversially, there is no such thing as simply *being beautiful*: things are beautiful or not in (relation to) ‘the eye of a beholder’. There is no such thing as its simply *being four o’clock*—it is four o’clock or not in relation to a time zone (that is, a set of places). Note that it doesn’t go with the “no such thing as simply being Φ ” rubric that in all cases there is no content to the idea of something’s being *absolutely* Φ —say absolutely useless. But the sense of “absolutely” involved will involve quantification through the latter place of the relation (useless for *every* purpose), not a genuine drop in degree (useless *simpliciter*.)

I suggest we think of the “no such thing as simply being Φ ” idiom as gesturing at a kind of proto- or ground-level relativism. The underlying idea is that the circumstances that confer truth or falsity on a predication of Φ tacitly involve a *further parameter* of some kind—a parameter that goes unreflected in the surface syntactic structure of the relevant predication. In such cases, what makes a predication of Φ on x correct or not is actually x ’s standing in a certain relation to certain relevant items that is more complex than the surface expression reveals. The ground-level relativistic idea is that the satisfaction-conditions of a

² See Wright (2001).

³ That is, where neither of two irreproachably generated but (apparently) mutually inconsistent opinions need involve error. As is familiar, this idea is not easy to stabilize. For a developed discussion of the issues, see my (2006).

certain property or family of properties, though superficially presenting as unary, are actually implicitly relational—or more generally, are of a *higher degree* of relationality than is apparent in the surface syntax of the predications and other characteristic forms of locution in the relevant discourse. Proto-relativism is thus a thesis about *tacit additional relationality*.

Such a thesis is often advanced in a critical spirit, or as a reminder: ordinary thinking is alleged to make a mistake about, or to be unaware of, or to be aware of but fail to pay due heed to, the real degree of a certain characteristic (or family of characteristics; say, the values). One issue for such a critic, naturally, is where exactly to locate the putative error or oversight in ordinary thought. Is it that the discourse in question calls, by virtue of its *literal sense*, for a characteristic of a certain degree when all that is available is something more complex—so that, strictly, there is sweeping untruth across the discourse? Or is it that, when speaking literally, we do succeed in referring to the more complex properties, but misunderstand their nature (and hence the sense of our discourse)? I'll return to this briefly below.

In sum: one root relativist idea as applied to a given region of discourse is that its distinctive vocabulary is somehow associated with characteristics of higher degree than manifest on the syntactic surface. But this is just the start. How do we move from here to something more recognizably Protagorean? There are no less than, arguably, five further moves to make.

As remarked, a tacit relationality thesis need not be to the effect that a certain apparently unary property is in fact binary. It may be to the effect that a certain apparently n -ary relational property is in fact $n + k$ -ary, $k > 0$. We need a distinction here. To fix ideas, let's consider whether the kind of *subject-sensitive* or *interest-relative invariantism* about knowledge proposed by writers such as Hawthorne and Stanley should be counted as a form of relativism, intuitively understood.⁴ Everyone accepts that knowledge consists in some kind of relation involving a thinker and a true proposition. What Hawthorne and Stanley are naturally viewed as proposing is that the relation in question is more complex than traditionally supposed: that it involves an additional kind of parameter—one of needs, purposes, interests, or saliences—variation in whose values between two same-world subjects can make the difference between one being knowledgeable that P while the other is not, even though their evidential achievements are otherwise exactly the same. It may seem not unnatural to propose a relativistic caption for this kind of view—that “knowledge is relative to need”, or to saliences, or whatever. Still, I expect that the theorists in question would scorn the description of their proposal as “relativist”, and for good reason. For while X 's knowing that P , on such a view, is indeed constituted in a complex relationship involving X and the proposition that P , in which X 's needs, or interests, or whatever, are an additional, traditionally unrecognized relatum, that point is not enough to motivate

⁴ John Hawthorne (2004), Jason Stanley (2005).

a relevant instance of the “no such thing as simply being Φ ” rubric. There is no sufficient motive yet for the claim that there is no such thing as (X’s) simply knowing that P. There is, in fact, no difficulty at all, for subject-sensitive invariantism, in the idea of X’s simply knowing that P. True: what *determines* whether or not X knows that P may involve additional relational complexity, at least in comparison with traditional accounts. But the upshot—prescinding from other possible sources of indeterminacy—is still that X does, or doesn’t, know that P, *tout court*.

So: what is it about the original examples that we gesture at by the “no such thing as simply being Φ ” locution, and which goes missing in knowledge as conceived by subject-sensitive invariantism? The difference is this. The truth-conditions of judgements of the form, X knows that P, proposed by Hawthorne and Stanley do indeed involve an additional parameter, unreflected in the surface syntax of the judgement. But settlement of the semantic values of certain of the (surface) syntactic constituents of such a judgement is *itself sufficient* to settle the value(s) of the relevant additional parameter(s). Merely settle, that is, who X is, and what point in his/her history is being spoken about, and you presumably have enough to determine a specific fact about what his/her relevant needs and purposes, etc., are. In the “no such thing as simply being Φ ” cases, by contrast, this—the settlement of the semantic values of other syntactic constituents of the judgement—is, precisely, not enough. In these cases, determination of the value of the extra parameter requires a contribution from factors lying outside anything settled by the semantics of the judgement concerned. There is no semantic constituent of “The Sun is moving” which denotes a frame of reference, nor are the semantic values of the constituents of that sentence, on a particular occasion of use, themselves collectively sufficient to determine a frame of reference (though other aspects of the context of use may be). Likewise, *mutatis mutandis*, for an utterance of “This screwdriver is useless”. There are no semantic constituents of that sentence such that, once their semantic values—for a particular context of utterance—are determined, that must suffice to determine the relevant value of a purposes-parameter, relative to which the instrument’s uselessness is claimed.⁵

The salient next question is: what *does* determine the value of such alleged, additional parameters in relativistic cases? Of course we are hampered by the lack of uncontroversial relativistic cases—that is, examples where the operation of

⁵ Here is an opportune place to begin to emphasize the contrast between relativity of truth-value of the kind encompassed by the proposals so far made and contextuality (including indexicality) of *content*. Contextuality of content has to do with the mechanisms whereby the semantic values of the constituents of a token sentence are determined. The semantic value of a context-sensitive expression is determined as a function of its semantic character and, precisely, a context (normally, the context of an embedding sentence’s use.) Relativity of truth-value, of the kind we are concerned with, also has to do with the presence of a parameter that needs to receive a specific value before any truth-evaluable claim can be entered by a use of the sentence. But the difference is that the determination of the value of this parameter is precisely independent of that of the semantic values of the constituents of the sentence, whether or not *they* are determined as a function of context (of use).

such a parameter, and the values it takes, can be uncontroversially agreed to be independent of the semantics of the form of sentence concerned. But there is no shortage of examples where such a view has been seriously proposed, or might at least be. And in many of these, the answer is, broadly: the intentions of the speaker. Suppose, for example, we are relativists about motion. Asked whether an overtaking car was travelling quickly, we may properly answer affirmatively even though its speed exceeded that of the car in which we are travelling by no more than five miles per hour—the ‘frame of reference’ is supplied by the conversational intention rather than the physical circumstances of the conversants. A little reflection, however, discloses that, at least in the kind of case most often connected with intuitively relativist ideas, it is not the situation of a speaker, qua speaker, that matters. What matters is the situation of a *judge*. Often speaker and judge are one. But when they are not, it is often in the situation, broadly construed, of the judge that the factors determining the additional parameter are to be found. Not always. If I report John as having asserted that it was raining, and you ask me if he was right, it would be expectable that your question is asking after the probity of John’s assertion relative to the place of *his asserting* it.⁶ But if I report him as having asserted that stewed tripe is delicious, and you ask me if he is right, it is natural to hear you as asking after a judgement of the matter by *my standards*; so here the determination of the relevant parameter goes with the judge, not the assertor. This is the point that underlies the organization of the modern debates around a conception of truth as relative to *a context of assessment*.⁷

To take stock. So far, we have diagnosed an intuitive relativism about the characteristic claims of a given discourse as involving three components: first, that the truth or falsity of such claims is constituted in states of affairs whose fully explicit characterization involves an additional degree of relationality as compared with what their surface syntactic structure would suggest. Second, the determination of the value of the extra parameters involved in this additional relationality is not a function of the determination of the semantic values of the syntactic constituents in (a tokening of) the claim. Third, what does determine their value belongs with the properties of an assessor of the claim. The result is that whether such a claim is true will depend not just on its content, and the state of the world in other respects, but on who judges it and *their* state at the point of judgement.

This is still not Protagoreanism, however. For all the conditions so far imposed, it can happen that, for claims of a particular kind, the factors determining

⁶ I am not suggesting that we are intuitively inclined to a kind of Protagoreanism about the weather. (In fact, we’ll look at the example in more detail later.) Remember that at this stage we are merely working with the “no such thing as its simply being Φ ” rubric. “It is snowing” and its kin at least comply with this rubric: there is indeed no such thing as it’s simply snowing—it depends on *what place* is considered.

⁷ My impression is that this turn is more or less wholly owing to the influence of John MacFarlane’s writings.

the values of the additional parameters are such as to settle the *same* values for all participants in any particular conversation (in however broad a sense of ‘conversation’). Should that be so, there will be no possibility of intra-conversational fault-free disagreement—or at least, none occasioned purely by the relevant form of relativity. In order to provide for faultless disagreement (if indeed it is possible to make coherent provision for it), whatever settles the values taken by the extra parameters must allow them to take on distinct values for distinct assessors of a single claim—thereby generating the possibility of disagreement occasioned just by differences in the values taken. And if, in addition, disagreement is to be possibly *faultless*, the distinct values assumed for different thinkers must likewise be determined in a manner that provides no grip for the idea of fault—that is, in a fashion that does not involve bringing the assessors into liability of reproach nor underwrite any disadvantageous comparisons between the values taken (as when, for example, although X and Y are both right by their respective standards, X’s standards are properly adjudged superior to Y’s).

We thus arrive at the following proposal about what an intuitive full-blown, Protagorean relativism about a property Φ —goodness, beauty, justice, obscenity—may, as it were incrementally, consist in. (We restrict attention, for simplicity’s sake, to the range of atomic predications schematized by ‘ Φa ’.)

- (i) The circumstances conferring truth or falsity on any single token of ‘ Φa ’ involve an additional parameter, V, unreflected in the overt syntactic structure of the sentence.
- (ii) The value taken by V is not settled by determining the semantic values (in the context of use) of the syntactic constituents of the sentence.

These first two clauses secure a basis for the “no such thing as simply being Φ ” idiom.

- (iii) The value taken by V is determined instead by characteristics of one who assesses the token of ‘ Φa ’, at the point of assessment.
- (iv) These characteristics may vary between distinct assessors, for whom accordingly, as a result of different values taken by V, differing assessments may be appropriate of the token in question.

These two clauses secure a minimal Φ -relativism, involving the possibility of correct but conflicting judgements about its application in different contexts of assessment. However they do not ensure the possibility of faultless disagreement, for which we require in addition:

- (v) The variation in values taken by V may afflict participants in a single conversation.
- (vi) There is, or need be, no sense in the idea of the superiority of the assessments constrained by one value of V over those of any other.

7.2 TRUTH-RELATIVISM

So much for relativism about Φ . Our topic, though, is relativism about *truth*. We can, of course, apply the successive steps in the foregoing more or less mechanically to the idea of truth itself. The resulting truth-relativism will hold, in the first instance, that ‘*is true*’ (i) ascribes to a statement *S* a property of tacit relationality, involving (ii) an additional parameter whose values are independent of—i.e. left unsettled by determining—the semantic values of the constituents of *S*. Since—apart from those (deflationists) who do not reckon it to be a property at all—anyone will view being (contingently) true as implicating at least a binary relationship of some kind—viz. being *made true* by relevant circumstances—something worth describing as truth-relativism ought to involve the view that truth is at least *ternary*, with the values of the additional parameter fixed (iii) by certain of the characteristics of an assessor of *S*’s truth. We can then escalate up through clauses (iv) to (vi), arriving eventually at the view that *S*’s truth-value may vary with certain properties of its assessors—properties that in turn may vary irreproachably among the protagonists in a single conversation.

That may seem more or less like the kind of thing one would want truth-relativism to be. However two issues require immediate note. First, there is the question, what exactly is truth being proposed to be ternary property *of*? What is the schematic range of ‘*S*’? The claim of tacit additional relationality is nothing terribly startling if restricted to truth-bearers whose *content* is variable (e.g. type-sentences). It is the merest platitude that truth is relative to (varies with) content. Truth-relativism becomes potentially interesting only when the truth-value bearers are conceived as *beliefs*, or as *thoughts* (one use of “proposition”, which I shall stick by here), or—as more commonly in the contemporary literature—as *utterances*, but where the historical context of the utterance (by whom, when and where it is made, and in what collateral circumstances), and the semantics of the language, are fixed; in short, when everything is fixed which is normally conceived as sufficient to fix an utterance’s content.

This reflection—that any interesting truth-relativism operates *after* content is fixed—signals one major strategic line of resistance to it: that of (what I propose we here and now decide standardly to call) *Contextualism*.⁸ Contextualism, where philosophically contentious, is the thesis of some (interesting and unobvious) form of relativity of *content*. Whenever there is a case for truth-relativism—whatever form such a case may assume—there has to be theoretical space in principle for a corresponding and opposed contextualism: a

⁸ The terminology in the literature is already getting horribly tangled, and matters have not been helped by the choice of “non-indexical contextualism” by John MacFarlane to denote a type of view that is actually a variety of truth-relativism. (I’ll say a little bit more about this later.)

thesis to the effect that the additional parameter to which the truth-relativist contends that truth is relative, is actually something variation in which leads to variation in *truth-conditional content*, with the resulting variable propositions then taking their truth-values in some unremarkable, non-relativistic manner. Thus it may be suggested, for example, that it is not that the truth of the proposition that stewed rhubarb is delicious varies as a function of the standards of taste of different linguistically competent English consumers, but rather that the *propositions* they respectively express by tokens of the sentence, “Stewed rhubarb is delicious”, may so vary, and that these propositions may then take different truth-values in a straightforward, non-relativistic manner. Strategically, it is expectable that what may look like evidence for relativity of truth-value may always be *explained away* by a content-relativity thesis.⁹

It might be thought that this might be challenged as follows. Have we not in effect excluded any such diagnosis by requiring that S’s apparent relativity in truth-value should manifest as a function of variation of features of assessors, or more generally of *assessment-context*? For by contrast the features that characteristically determine the content of an indexical, or otherwise context-sensitive judgement are standardly part of the *context of its making* (the context of its utterance, if it is made out loud, or of its inscription if made on paper), rather than a context of assessment. Surely, then, there will not in general be an option of explaining the kind of relativity in truth-value characterized by our clauses as turning on context-driven shifts in content.

Not so. The correct point is rather that the form of contextualism that always potentially stands opposed to truth-relativism is a *non-standard* contextualism. Any data that suggest that judgements of a certain kind do admit of variations in truth-value in the manner characterized by the clauses we listed will allow in principle of a contextualist, rather than truth-relativist explanation provided the contextualism in question is one that allows that the *content* of a single token statement need not be something unique but may vary as a function of the properties of an assessor—in other words, that what someone literally says is not determined just by standing semantic properties of the type-sentence and the circumstances of the saying but has its own relativity to parameters whose values may vary among assessment-contexts. Far from being unheard of, such a view belongs with a long ‘hermeneutic’ tradition in philosophical semantics.¹⁰ An important point for our present purposes is therefore that truth-relativism is heavily invested in a conception of content that requires this tradition to be misguided, for the two kinds of view promise to be co-predictive in crucial cases otherwise. I’ll say more about this below (Section 5).

⁹ This reinforces Williamson’s (2005) slogan that “Contextualism is relativism tamed” (though the point is not quite what Williamson had in mind).

¹⁰ Most recently represented by the “interpretationism” of Davidson.

The second salient matter is that if the general account proposed of relativism about an arbitrary property, Φ , is accepted, at least in outline, then it is, in a way, confused to seek an account of what specifically relativism about *truth* might consist in—since there is no specific such thing. For, providing that the so-called Equivalence Schema:

It is true that S if and only if S

is not called into question, the import of truth-relativism for a range of admissible substituends for 'S' is always going to reduce to relativism about their proper subject matter, whatever it is. In the presence of the Equivalence Schema, relativism about truth for ascriptions of beauty is just relativism about beauty. And in general, relativism about truth for ascriptions of the property Φ is just relativism about Φ .

Although this is apt to impress as a point of some significance, I doubt that it is so. Certainly it would be a mistake to conclude that there is nothing for relativism about truth to be, or anything of that sort. The right conclusion is only that relativism about truth is always relativism about something else. You can be a relativist about beauty and about no other subject matter. But you cannot be a relativist about truth and about no other subject matter. Still if you are a relativist about beauty, you are a relativist about the truth of ascriptions of beauty, and some of the implications of your view—for instance its implications for content—may be easier to bring into focus when it is taken in the latter guise.

A final point of emphasis. According to the preceding perspective, truth-relativism comes by incremental stages, the earlier among which need not involve that the postulated extra parameter(s) should implicate any kind of subjectivity or anthropocentricity. Of course many of the proposals that make up the 'usual suspects'—perspective, point of view, sensibility, standards of taste—explicitly do so. But that doesn't have to go with assessment-relativity as such: one possible counter-example (to be reviewed later) is provided by Aristotelianism about the future. Here the additional element of relativity is merely time: no subjectivity need be involved in the idea that a proposition depicting a sea-fight at t is true at t —if a sea-fight occurs—but is untrue at any time before its occurrence. The metaphysical view in question would equally apply to future contingents at a world altogether free of human agency or perspective. In general, the connections between truth-relativism and the 'pluralism' or 'perspectivalism' of the Protagorean tradition are not straightforward: the former is neither sufficient—nor, I would argue, necessary¹¹—for the latter.

¹¹ I have it in mind that the notion of a discourse's failing to exert Cognitive Command, in the sense of my (1992), is an adequate vehicle for a sober pluralism, and that its characterization requires no use of a relative truth-predicate. But I cannot elaborate this thought here. Cf. Wright (2006).

7.3 THE TRADITIONAL MISGIVING

There has always been a major question whether any interesting form of relativism about truth is so much as coherent. Everyone knows one classic kind of argument why not: if truth is relative, then the thesis that truth is relative is itself at best only relatively true, and that—for reasons which are seldom clearly stated—supposedly somehow makes it abrogate its own cogency, or dialectically undermine the position of one who advocates it. This is one form of (what we may call) the *traditional misgiving*.¹² Is it compelling?

Perhaps the most immediate observation is that the objection seems to presuppose that the relativity in relativism must be to something optional, so that the relativist may always be confounded by shifting the value of the parameter involved in such a way that the very relativistic thesis is itself incorrect. The idea seems to be that if the thesis of truth-relativism is itself only relatively true, then an opponent can always deflect the thesis by judicious parameter selection, as it were. We have already noted that this need not be so. The extra parameter(s) may be something whose values are settled independently of any contingencies of anthropology or personality. And even if not, the upshot need only be, at worst, that that truth of relativism would itself be somehow dependent on contingencies of human perspective. Such a result need no more “optionalize” the acceptability of the view than—supposing it is so—the relativity of colour to human sensibility makes it optional what to think about the colour of the sky.¹³

There are other pressures worth reviewing, distinct from the traditional misgiving, to think that global truth-relativism is incoherent. Two of them I’ll label, in a Quinean spirit, ‘from above’ and ‘from below’. Both can be elicited just by reflection upon the general form of the idea that truth is relative to a ‘context of assessment’—that it is, so to speak, ‘in the eyes of an assessor’—without

¹² Its locus classicus is of course Socrates’ dissection of Protagoras’ ideas in Plato’s *Theaetetus*. The modern discussion includes Jordan (1971), Meiland (1977) and (1979), Seigel (1986), and Swyer (1982).

¹³ There is a subtler response to this form of the traditional misgiving. Reflect that the assumption that relativism is committed to regarding its own proper statement as only relatively true is unfounded in any case unless it is taken that relativism is a global thesis, extending to *all* truth-evaluable contents. However the thought that if the thesis of relativism is merely relatively true—true relative to the particular value of some relevant parameter which the relativist instantiates—then it may be seen off by calibrating oneself, as it were, to instantiate a suitable different value, seems to presuppose that what is true once such a value is settled is itself an *absolute* matter; that is why the deflection requires a different selection. Whereas a global relativism must also embrace judgements of the form that the relativistic thesis is true relative to v_i , for variable values $v_1, \dots, v_n$, of the relevant parameter V . In that case, the contention that relativism is incorrect relative to some v_i will itself be merely relatively true. So when the anti-relativist triumphantly announces that he has so calibrated himself that, relative to his value of V , the thesis of relativism is incorrect, he is open to the rejoinder that such a claim is itself at most relatively correct and cannot therefore be presumed to hold in an arbitrary context of assessment—in particular, not that of his relativist opponent. I leave the further prosecution of this dialectic to the reader!

troubling to enquire how exactly the idea of a context of assessment is to be understood, what variables it encompasses.

The pressure from below emerges like this. There is among modern proponents no suggestion that truth depends *only* on the standards of an assessor, or other aspects of a 'context of assessment'. It will also depend on an *input*: the material to be assessed (the 'circumstances of evaluation', a *world*). This input must presumably allow of propositional articulation in turn, so the question arises how the truth of *these* propositions—those that characterize the worldly input—is to be conceived. If it is conceived non-relativistically, then we are not after all dealing with global relativism. But if it too is conceived relativistically, then a regress is launched: the truth (in a context of assessment) of each relativistic proposition depends on the truth of another proposition articulating the relevant circumstances of evaluation—and when relativism is global, this too is in each case a relativistic proposition. Note, however, that this regress ascends through an infinity of types of *logically independent* proposition; for nothing follows just from the description of the circumstances of evaluation—from the input propositions—at any particular stage about which among the relevant truth-relativised propositions are true; and conversely, or so one would suppose, nothing follows just from the true relativistic propositions at each stage about exactly which of the relevant input propositions are true. The result is an epistemological problem. Mastery of the relativistic truth-conditions of propositions at each level in the regress will presumably require grasp of the truth-conditions of propositions describing the relevant kind of circumstances of evaluation, or worldly input, for them; so the threat is that mastery of the former will involve an implicitly infinitary range of propositional understanding, not in the benign sense of something recursive with a finite base but in a sense that makes it unattainable.

The pressure from above is this. In general, that P is true relative to circumstances C and parameter V looks like an *absolute* claim, whereas global truth-relativism will have to construe the truth of such claims as relative in turn. So parameters will have to be identified to which the truth of these absolute-seeming claims will itself be relative. This is very implausible in most cases, and will be additionally so if the very same parameter is just repeated. Claims, for instance, about what is true relative to certain ethical standards, if they themselves have some implicit relativity, are at least, presumably, not relative to *ethical* standards. What coherent account can be given of such an ascending hierarchy of invoked standards, or other relativistic parameters? Does it stabilize in some one Master Parameter, to which all truths above a certain level are relative? Or is there a literally endless variety?

I don't know whether either of these misgivings, or others, can be developed to the point of conclusiveness. It would be of interest finally to resolve the question whether a global truth-relativism can be a coherent stance. But I do not think it is important to do so. Even if the global thesis is incoherent, the fact is that the

relativist views in contemporary debate are typically local—to epistemic modals, morals, knowledge ascriptions, or conditionals, for instance. As far as I am aware, each of the relativist positions in these local debates is quite comfortable alongside an absolutist view of the philosophical thesis that truth for the relevant locality is relative. (Indeed, as a limiting case, one might hold that truth in philosophy is absolute but relative everywhere else!)¹⁴

7.4 TRUTH-RELATIVISM AND REPRESENTATION

We are entertaining the idea that for certain kinds of proposition, *P*, truth is to be thought of as fixed not merely by a world—a set of circumstances of evaluation—but by a world-cum-context-of-assessment. Call this the *ternary model*. Thus it may happen, one would suppose, where standards, say, are the relevant variable within contexts of assessment, that *P* is true by *S*₁ in *w*_A, but false by *S*₂ in *w*_A. Considerations not far from the territory just reviewed now suggest that the ternary model is inconsistent with *P*'s possession of *representational* content. For what state of affairs might we think of *P* as representing? Nothing, presumably, that might be mentioned in a compendious description of *w*_A—since those matters are all, as far as they go, neutral on the question of the truth of *P*, whereas the obtaining of a state of affairs of a kind represented by *P* would have to suffice for its truth. But nor are we at liberty to think of *P* as representing some state of affairs of the form: *by standard so-and-so, such-and-such is the case*—for that would be to misrepresent its content. As noted, facts about standard relativity are not in general plausibly taken to be themselves standard-relative; and even when they are, they are not plausibly taken to be relative to the same standards. But what other candidate type of state of affairs is there for *P* to represent? It looks as though the ternary model excludes representationality.

This is hardly surprising if, as I think is so, representing is an essentially binary relation holding between a *representans* and the matter (if any) that it represents. In saying that, I am not overlooking that whether something successfully pictures, or encodes, or in some other way represents something else, may depend on conventions of representation, or a 'method of projection', variation in which may lead to the question of successful representation getting a varying answer. But of course that's an irrelevant consideration. Varying the 'method of projection' is to be compared with varying the *content*—and our interest all along has

¹⁴ An arresting argument against the coherence even of certain *local* forms of truth-relativism is offered by Paul Boghossian in his contribution to this volume. It applies to any form of truth-relativism where the claimed relativity is to standards and where standards are thought of—as seems most natural and intelligible—as consisting in a set of general propositions informing the assessment of particular cases. Considerations of space prevent me from including here the discussion of it that I offered in my presentation at the Barcelona workshop. I do so in Wright (2008a).

been in the nature of representation (and truth) *after* content has been fixed. That is why we are focused on truth for propositions. Representation, once the conventions of representation have been fixed, surely *is* essentially a binary relation. So having the kind of content fit to represent something is—in the case of propositions—having the kind of content that fits a bearer to stand in a binary *true-of* relation. Propositions for which the ternary truth model is appropriate cannot, on that account, have that kind of content.

I do not want to belabour the point, but maybe it is worth pausing over. Cannot representation, it might be objected, be *perspectival*? Indeed, it is suggestive that at least one writer in the vanguard of the recent surge of interest in relativism uses exactly the term, “perspective”, as a catch-all to denote the kind of parameters to which relative truth is supposedly relative *to*.¹⁵ Perhaps only perspective-free representation is essentially binary. But clearly the fidelity of, say, a visual representation—a diagram, or a photograph—to the shape of an object can depend upon settling a notional perspective. For example, these figures




Figure 7.1




Figure 7.2

are potentially both adequate representations of one and the same solid object once it is specified that the respectively implied angles of view are such that first the point A, and then the point C, are to be thought of as representing that part of the object that is closest to an observer; but not if it is given that the closest points are A and D. Successful representation can be and often is *perspectival*—that is, in the present kind of example, relative to an implied angle of view—in this kind of way. So, it may be thought, it does *not* go with the very idea of representation that it be a purely binary relation; and there is, correspondingly, no reason why a theorist who thinks of the truth of propositions in some region of discourse as involving relativity to an assessment-parameter should thereby foreclose on their representationality either.

The analogy has the additional feature, attractive for the relativist purpose, that there is a sense in which the above figures are *incompatible*: they cannot both be accepted as accurate representations of one and the same figure from any single perspective (implied angle of view). On the other hand, they can both be accurate from a suitable pair of perspectives, and a pair of thinkers, occupying those perspectives can, in respectively endorsing their accuracy, both be right.

¹⁵ Kölbel (2002).

Is it along the lines of this analogy that we should try to fashion our thinking about truth-relativism about propositions and the idea of faultless disagreement? Well, only if we can think of the ‘contents’ carried by such figures as comparable to propositions. We have an analogue of the sought-after idea—that of a species of proposition which is representational in content all right, but where the fact of actual representation turns on perspective, or point of view—only if we suppose that the content of the figures is invariant under changes in perspective or point of view. It is here, obviously, that the analogy breaks down. Of course *something* is invariant under changing perspective. The diagrams don’t change shape, and the lines continue to represent edges, and so on. But it is clear that what is unchanging is not representational content: i.e., how the object is being represented as shaped. Suppose you are told that Figure 1 *is* an accurate representation of an aspect of some solid object; what do you thereby learn about the object’s shape? The answer *varies with perspective*. If you are also told that the implied angle of view is such that point A on the diagram represents the nearest point of the object to the observer, you learn that three of its sides are configured in a way consistent with its being a cuboid. If you are told only that point A on the diagram does not represent the nearest point of the solid to the observer, you do not learn that and then there are any number of other possibilities. Perspective—the implied angle of view—in such cases is accordingly first a parameter in the determination of *representational content*, and only secondarily in the determination of whether the figure successfully represents. Whereas—to stress again—truth-relativism at the level of propositions sees truth-value as variable in tandem with variation in some relevant parameter *after* content has been fixed.

There is no doubt a lot more to be said, but—with a qualification to follow—I think we can reasonably take it that further reflection should be constrained by the following *Non-Representationality Conjecture*:

Truth-relativism, construed along the lines of the ternary model, excludes the representationality of the propositional contents concerned.

A number of questions are pertinent if the Conjecture is accepted. First, any good argument for a (local) truth-relativism at the level of propositions should implicate grounds for denying that the contents in question are representational. One crux for representational content for which I have argued elsewhere is the constraint of *Cognitive Command*: a constraint which a discourse satisfies just in case it’s a priori that differences of opinion within it—where they do not fall among certain admissible exceptions: difference of opinion due to vagueness, for instance, or idiosyncratically varying thresholds of evidence—will turn on something properly viewed as a *cognitive shortcoming*. If, as I have argued,¹⁶ Cognitive Command is a necessary condition for representationality, then one

¹⁶ Wright (1992), 88 and *passim*.

natural channel for argument for truth-relativism will be to make a case for a failure of Cognitive Command. It may not be the only possible channel. However, whatever argument is advanced if informed by the ternary model, must—in the presence of the Non-Representationality Conjecture—somehow carry the conclusion that the discourse in question is not representational.

There is also a good question about why, if the Non-Representationality Conjecture is correct, anyone should want to propose truth-relativism for any of the broadly anti-realist reasons that have been traditional.¹⁷ For if there is a solid case that a particular discourse is not representational in content, that *already* undermines the realism associated with absolutism. If ethical discourse, for instance, is not representational, we already have the conclusion that there are no absolute moral facts, since if there were such facts, they would have to allow of representation by moral thought. What need, then, to claim an additional relativity in the notion of moral truth? Some deflated or minimalist line of goods will serve all purposes worth serving.

The connections between relativism and representationality are complicated, however, by the realization that not everything that's been presented as an example of truth-relativism is happily thought of in terms of the ternary model. Consider, for example, the Aristotelian view of future contingents, recently argued by John MacFarlane to provide an example of a view which calls for the kind of truth-relativism he thinks may be coherent and stable.¹⁸ Suppose at t_0 that X asserts that there will be a sea fight at t_2 , = P. According to the Aristotelian, it will be correct for an assessor, Y, at t_1 , < t_2 , to assert that X's assertion that P was untrue. It will be correct because there is nothing at t_0 or at t_1 to render P true—the future is open. But at t_2 it will be appropriate for Y to assert that X's assertion that P was true if—as we may suppose—a sea fight is then in full swing. So the very same content—that of X's assertion at t_0 —is correctly assessed as untrue at t_1 but true at t_2 .

Now it is of course perfectly true that we *can* cast the example in the form of the ternary model, with *time of assessment* in the slot occupied by 'perspective', or 'standards', in the other kinds of case. We can, if we like, include time of assessment as one more variable under the umbrella of 'context of assessment' in the ternary model. But it seems to me that this is a bad idea if we want to avoid confusion. It's certainly a bad idea if we are persuaded of the correctness of the Non-Representationality Conjecture, since the Aristotelian does not at all mean to contest the representationality of ordinary tensed discourse about mundane events like sea fights. What he means to contest, rather, is the ordinary *static metaphysics* of the being of such events: his view is that the future unfolds not in the sense of *revealing* itself as we reach it in time, but in the sense of, literally, *becoming*.

¹⁷ Of course these contrast with the putative empirico-linguistic reasons prominent in MacFarlane's writings and those of other recent authors in the field.

¹⁸ MacFarlane (2003) and (this volume).

Aristotelianism is not at all a view about the *degree* of utterance/propositional truth—it is, for example, consistent with the most robustly correspondence conception of truth. On the contrary, it's a view about the metaphysical behaviour of the *truth-conferers*. We can, if we like, caption it by saying that it holds that truth is relative to time—but what it really holds is that the *population of facts that constitutes the actual world* changes with time. The ternary model, properly understood, fixes the worldly parameter and *then* allows truth to vary as a function of 'context of assessment'. By contrast, what's happening in the Aristotelian metaphysics of time is that the *world itself* is varying with time, and truth at a time is varying accordingly. There is no genuinely additional parameter of context of assessment. Aristotelian truth-relativism is a kind of Fool's truth-relativism: a thesis about the nature of the world misconceived as a thesis about the degree of the truth relation.

Once that distinction is in the open, it's salient that the Aristotelian model could also be extended to the standards-and-taste varieties of truth-relativism. What, in effect, the Aristotelian view provides for is a plurality not merely of possible but of *actual* worlds, as indexed by the class of propositions that are true at them. The world at t_1 and the world at t_2 are both actual—that is why P changes its truth status as the passage of time takes us from one to the other. But once we have the prototype of time supplying different locations for actual worlds, what is to stop us allowing other forms of locality as well? We might for example permit the actual world *at Williamson*—that is the actual world as reflected in Timothy Williamson's gustatory standards—to exist simultaneously with the actual world *at Wright*, that is, the actual world as reflected in Crispin Wright's gustatory standards. The proposition that stewed rhubarb is delicious can then be true at the one aesthetic location, so to say, and untrue at the other, just as the proposition about the sea fight can be true at t_2 but untrue at t_1 . No need, in this case, to resort to a context of assessment parameter in order to accommodate the truth-relativist impulse—truth can be old-fashioned truth-at-a-world, *simpliciter*. The relativism surfaces, rather, in the thought that there is no single actual world but a plurality of them.

No doubt these worlds may substantially overlap. The thought that truth-relativism can plausibly only be a local doctrine can be accommodated within this framework by insisting that some denizens of any actual world will be denizens of all. Much of each actual world, perhaps, is made for us, but—the picture will be—we, each of us, individually and collectively, by our sensibilities, tastes and perspectives, give rise to locally variable actualities, partially creating our own domains of truth-conferers.

Maybe the right conclusion is that we need to distinguish two utterly different forms of propositional truth-relativism. One offers a package of the ternary model of truth for selected regions of thought and discourse, together with non-representationality of content, an implicit distinction from non-relativistic areas of thought and talk, and a single actual world fit for representation. The other

says nothing unorthodox about truth or content; it is, for example, consistent with both correspondence and deflationist accounts as holding across the board. But it abandons the idea of a single, comprehensive Tractarian world—a totality of all facts.

No doubt this distinction complicates the question of motivation. It also complicates the question of obligation: a good argument for truth-relativism about some class of propositions must, it appears, either enforce the thought that they lack representational content—so some other conception of their content will be owing—or it must enforce a pluralistic conception of the actual. Again, we should keep these points in mind when reviewing the considerations marshalled by those who are sympathetic towards (local) truth-relativism.

7.5 WEATHER-REPORT RELATIVISM AND NON-REPRESENTATIONALITY

The point still stands, it seems, that a theorist who is drawn to relativism about some region of discourse but wants no traffic with a pluralistic conception of the actual, is committed to thinking of the contents with which the discourse in question characteristically trades as non-representational. Call such a content a NRP (non-representational proposition). It is clearly an obligation on the relativist to say more about this species of content. We already remarked on the relativist's obligation to make a case that these contents are genuinely shared across disputes in which the protagonists differ merely in the values taken by relevant parameters in their respective assessment contexts; equivalently, to argue against any form of opposed contextualist suggestion that it is *content* that varies across such disputes. An account is also owing of the aptitude of such contents for all the usual propositional-logical and attitudinal constructions: they must be negatable, conditionalizable, believable, supposable, doubtable, and so on, and a story must be told about the commitments of one who takes any of these attitudes to an NRP. It is my impression that much more needs to be said to address these issues. But I will not try to explore them further here.¹⁹

What I want to canvass in this section is one, in some ways disappointingly deflationary account of what NRPs could turn out to be. Consider again the example of reports of the weather of the type, "It is snowing", "It is windy" etc. For reasons to be outlined shortly, these do not seem promising as candidates for

¹⁹ Those who assert a relativistic proposition presumably commit themselves to its correctness in the assessment context that coincides with the context of its assertion. The crucial question is what commitments are undertaken with respect to other contexts of assessment. MacFarlane has proposed that such an assertor commits to the continuing assertibility of the proposition in question in *any* assessment context, no matter what the value of the relevant parameter in that context. There are theoretical pressures towards this proposal but it seems like an extravagant, even an irrational commitment to undertake. See MacFarlane (2005a), § 5.

truth-relativistic construal. But they certainly pass the most basic of our intuitive tests: there is indeed, plausibly, no such thing as its simply being the case that it is snowing or not—there is the possibility of a true or false verdict whether it is snowing only when the issue is understood to pertain to a particular location. So while no one has, to the best of my knowledge, seriously proposed a relativistic account of weather reports, it is very easy to envisage the beginnings of one. The relevant assessment-contextual parameter would simply be place: “It is snowing” as uttered by T would be correctly assessed by A as true just in case there is snow falling at the place and time of A’s assessment.

It will, of course, be immediately objected that this gets the truth-conditions of such utterances grossly wrong: that what, for competent judges, determines whether or not T speaks truly is not *their* meteorological predicament, at the time and place of assessment, but that of T at the time and place of the utterance. But that claim is certainly not exceptionless. Invited while in St. Andrews to express an opinion on the weather currently in New York, I may venture simply, “It is mild and raining”. So the immediate objection was only a first approximation. A better formulation would be that what determines the place with respect to which an utterance of “It is raining” is to be assessed is the intention of the speaker, not the context of an assessor.

Even this, however, is not the real objection to weather-report relativism. The point about speaker-intention can be conceded consistently with the broader spirit of relativism. That broader spirit requires only that “It is raining” and its family be viewed as expressing propositions whose truth-evaluability then depends upon settling a place. This place is left to be determined as *some* kind of function of context of assessment. But there is no particular reason why that function should require it to coincide with the place of assessment. It can be that the way in which context of assessment determines the relevant place is, sometimes or normally, via interpretation of the intentions of an utterer. This is still relativism: one and the same proposition is still conceived as expressed by utterances of “It is raining” in different locations and its truth-value is still conceived as relative to place. It is just that (the rules of assessment determine that) the relevant places are the places that the respective utterers intend.²⁰

Nevertheless, this whole direction stands opposed to what may seem like utter commonsense on the matter. According to commonsense, when you truly utter (A) “It is raining” in New York and at the same time I truly utter (B) “It is not raining” in St. Andrews, it is not that we disagree faultlessly—rather, we do (or need) not disagree *at all*. According to the extended relativist proposal, we respectively affirm and deny the same content—hence the disagreement—but that content requires assessment with respect to different places—hence faultlessness. And against this, it would be normal to deny the first: to report us instead as

²⁰ This is actually an example of the extended form of relativism rather unhappily termed “non-indexical contextualism” by John Macfarlane. See especially his (2005b).

having been *speaking of* different places, the lack of any overt reference to place in the uttered sentences notwithstanding. The relativist treatment mislocates the parameter of place when it leaves it to be settled by assessment-context, whatever the proposed mechanism of settlement. Rather, it belongs in the *content* of the propositions respectively expressed and is settled by that.

That's the respect in which weather-report relativism essentially offends commonsense intuition. Naturally, the commonsensical view engenders theoretical obligations. If it is right, then in uttering the weather-reports (A) and (B), we speak of different places using forms of words that explicitly contain no vocabulary of place. An account needs to be given of how this is achieved. Familiarly, the matter is contentious. One proposal is that a reference to place is an *unarticulated constituent* of the propositions expressed.²¹ Another is that the reference to place occurs at the level of the deep syntax of the uttered sentences—that there is, so to say, a hidden indexical expression for which context of utterance serves to supply a reference in the normal kind of way. There is much to say about the content of these contrasting proposals and their respective advantages and problems.²² But on both of them, the contents respectively affirmed by utterance (A) and negated by utterance (B), when identified with the propositions expressed, are not the same.

A serious relativistic proposal about weather descriptions will hold that the content affirmed by (A) and negated by (B) *is* the same, but that it is a content which is truth-evaluable only when relativized to a place. However, this proposal stands opposed to both those just adumbrated only when the mention of contents is construed as a reference to *propositions*; that is, to complete, truth-evaluable thoughts. The relativist proposal takes it that tokens of “It is raining” and its ilk express fully articulated, complete propositions, albeit propositions whose truth-values vary with a parameter which is unrepresented in their logical form. So there is an evident question: why assume that tokens of those sentences express any such *fully articulated* propositions? The alternative is to think of the assignment of a value to the place parameter as serving to *complete* the propositions, rather than set up the means to assess it—completing it, for example, by assigning a referent either to an unarticulated constituent or to a hidden indexical for place.

The exact details of the opposing view do not matter. The crucial point is that the relativist is making an assumption in taking his distinctive direction against such views. For the case of weather reports, it is the assumption that “It is raining” and its ilk provide a fully articulate and explicit expression of the logical form of the relevant, allegedly truth-relativistic propositions. In this case, it is, as noted, most people's inclination to regard the assumption as

²¹ See Perry (1998).

²² For further discussion, see Perry loc cit., Stanley (2000), Recanati (2002), and Cappelen and Lepore (forthcoming).

false. But the crucial point is that this general form of theoretical dispute must arise *in every instance*—every region of discourse—where some kind of case for assessment-relativity can be made. For wherever there is such a case, there will be the option of supposing either that the propositions expressed by the targeted sentences contain unarticulated constituents referring to values of the assessment-contextual parameter which features in the relativist account, or that they contain hidden indexicals with the same range of reference, whose values are settled by features of an assessment context. The latter would, admittedly, be an unusual kind of indexical. But perhaps it is to the existence of such unusual indexicals that linguistic data that might otherwise seem to support relativism would be best construed as pointing.

To fix ideas, suppose we are inclined to a relativistic account of taste-evaluations like “Stewed rhubarb is delicious”, according to which a single proposition thereby fully explicitly expressed is correctly evaluated as true by Tim Williamson’s gustatory standards but false by mine. According to the “hidden indexicals” proposal, there is no such single proposition so evaluated. Rather the proposition expressed is of the deep syntactic form: [stewed rhubarb is delicious by *x*], where ‘*x*’ marks a place for an indexical for standards, whose value is normally²³ fixed as the standards of one assessing the claim. So Tim and I evaluate different propositions, one making mention of his gustatory standards, and the other making mention of mine. And it is the same on the unarticulated propositional constituents proposal; the difference is merely that the references to standards respectively effected by the propositions concerned correspond to nothing in the syntactic structure, even the deep syntactic structure, of the sentence but are supplied by pragmatic factors which vary in the two assessment-contexts.

It is this latter train of thought that promises the “disappointingly deflationary” account of NRPs advertised at the start of this section. On this proposal, “Rhubarb is delicious” expresses *no* complete, truth-evaluable proposition. The content common to “Tim thinks that rhubarb is delicious” and “Crispin doubts that rhubarb is delicious” is indeed non-representational. But it is also non-truth-evaluable. What are truth evaluable are the propositions that uses of the above sentences pragmatically convey as the respective contents of Tim’s and Crispin’s attitudes. But these are not NRPs: they are regular relational propositions about what passes muster by standards they mention; they represent, and—we may presume—they take their truth-values absolutely. On this view, there are no NRPs. The forms of words that seem to express them do not, semantically, express propositions at all. Rather they are, so to speak, sawn-off expressions of proposition radicals—as it were, propositional *amputees*—whose completion is

²³ Evidence: suppose you have never tasted rhubarb and ask whether Tim spoke truly when he said that rhubarb is delicious. Your interest is unlikely to be in whether he spoke truly by *his* standards.

accomplished by the pragmatics of assessment, and whose completed products are regular propositions, true or false in regular ways.

7.6 CAN TRUTH-RELATIVISM BE MOTIVATED BY PURELY LINGUISTIC DATA?

Let us nevertheless persist with the idea that there are genuine NRPs—full-fledged propositions for which truth is relativistic in some parameter, V , of assessment-context. To each such proposition, corresponds a range of explicitly relational propositions (ERPs), each of the form: it is true that P relative to value v of V (“That rhubarb is delicious is true relative to Tim Williamson’s gustatory standards”, “That there could be snipers over the rise is true relative to the information possessed by the platoon commander”). ERPs, it is natural to think, will not in general themselves be candidates for relativistic truth, and indeed there are worries, some canvassed earlier, one to be outlined below,²⁴ about the coherence of views that allow otherwise. However they will all be associated, naturally, with possible assessment as true or false. So here is a question we can ask about a given species of (putative) NRPs and the corresponding range of ERPs: must any context in which it is possible competently to assess an ERP in the latter range itself supply a value, v , for each assessment parameter, V , relative to which the relevant NRPs take their truth-values? In other words, if you are in position competently to assess a claim of the form

P is true relative to value v of parameter V ,

must your situation itself incorporate a value for V , and so commit you to your own assessment of the truth-value of P ?

A quick review of the types of discourse for which modern truth-relativists have proposed the view shows a mixed picture. There is, for example, no particular reason why a subject with no gustatory standards or predilections—perhaps someone born with no taste-buds—should be hampered when it comes to assessing “Stewed rhubarb is delicious by Tim Williamson’s standards”; she can have exactly the same evidence—Tim’s sayings, actions and reactions—as the rest of us. Similarly, there seems no reason why an utterly amoral subject—one party to no moral standards whatever—should be hampered when it comes to the assessment of claims of the form, “Incest is wrong by Western standards”, or “Welfare-maximizing punishment of the innocent is permissible by strict utilitarian standards”. But suppose we take a relativistic view of, say, logic, holding that, notwithstanding a uniform understanding of its logical vocabulary, one and the same propositional schema, e.g. Peirce’s Law, is valid by classical standards but invalid by intuitionistic ones. Here the assessment of the relevant ERPs itself

²⁴ In the Postscript.

demands propositional inference, and hence implicates a set of rules of inference. So there has, presumably, to be an answer to the question whether Peirce's Law may be validated by just those rules. Whether that commits an assessor to a view about Peirce's Law depends, of course, on what the answer to that question is. If the answer is "no", the matter will turn on whether the assessor takes it that the rules he is using to assess the two ERPs are the strongest sound such rules. So logical relativism offers a kind of intermediate, *partially committed*, case, lying between *neutral* cases like, seemingly, ethics and aesthetics, where competent second-order assessment—assessment of an ERP—need involve no commitment to any particular first-order assessment—assessment of the associated NRP—and cases where there always has to be such a commitment. It is the latter that I am calling *fully committed* cases. Are there any examples of this kind?

Certainly, at least according to proposals made by contemporary relativists. Suppose, for example, that epistemic possibility claims, "It might be that P", take their truth-values relative to the state of information of an assessor.²⁵ The relevant ERPs then take the form, "It is true that it might be that P, relative to state of information I". But to assess any such claim, you have of course yourself to be in *some* kind of state of information; and that, whatever it is, will itself presumably mandate a verdict about the epistemic possibility of P. So if truth for epistemic possibility claims is assessment-relative in this way, then they provide an example of a fully committed case.

It will presumably be similar with the view that knowledge-attributions are assessment-relative²⁶—that whether X knows that P depends upon the standards of an assessor. Suppose you are in position to assess whether X knows that P relative to certain demanding standards for knowledge. Then your situation will presumably both involve awareness of what comprises X's relevant information and cognitive state and a set of standards—perhaps less demanding—of your own for the ascription of knowledge that P to X. So you will be committed to a view about whether X does know that P.

There is a third very salient example: precisely the (putatively) Aristotelian view that the truth-value of contingent propositions concerning singular events is relative to *time* of assessment. We are, ineluctably, stationed in time. So an Aristotelian thinker called upon to assess the truth-value of a proposition of the form, 'What Aristipides asserted at t_0 by saying, "There will be a sea-battle tomorrow", is (tenseless) true at t_1 ', will himself occupy a station in time relative

²⁵ They don't always—there are apparent cases of *deferential* uses, where an assessor reports, or colludes in a possibility for an audience or third party: a Red platoon commander may order his trainee soldiers to "Keep down: there could be snipers on the hill", even though he knows there are no Blue positions in the area. For sympathetic proposal of assessment-relativism for epistemic modals, see MacFarlane (2003a and 2006b), Egan, Hawthorne, and Weatherston (2004), Egan (2005) and Stephenson (2005). Critical discussions of epistemic possibility relativism include Dietz (this volume) and Wright (2007). We will focus on this example in a moment.

²⁶ See MacFarlane (2005a).

to which Aristipides' prediction will have or will not have (yet) come to fruition; and so will be committed to his own assessment of that prediction.

What is the interest of the notion of a fully committed case? That there is an interesting difficulty about characterizing what, in principle, would constitute good linguistic evidence for a truth-relativistic semantic treatment of any region of discourse where it would have the effect that the discourse in question was fully committed. Consider for example, the way that John MacFarlane and others have actually tried to motivate relativism about epistemic possibility claims. MacFarlane writes

Here's the puzzle. When we claim 'It might be the case that *p*' and later come to know that not-*p* (for example, by being told this by a trusted authority), we tend to *withdraw* our earlier assertion and regard what we said as false. That is, we tend to assess past assertions containing epistemic modals against our *current* epistemic situations, rather than the epistemic situations in which they were made. Consider the following dialogue:

Dialogue 1

Sally: Joe might be in Boston. (= It might be the case that Joe is in Boston.)

George: He can't be in Boston. (= It is not the case that it might be the case that Joe is in Boston.) I saw him in the hall five minutes ago.

Sally: Oh, then I guess I was wrong.

This is not what one would expect if the standard semantics for epistemic modals were correct. If the truth-conditional content of Sally's original claim is just that she does not (at the time) know that Joe is not in Boston, then she should not retract this claim upon learning that Joe is not in Boston. What she said is still true, even though she could not *now* truly say "Joe might be in Boston." So if the contextualist semantics is correct—the Simple version, at any rate—then we should expect the last line of the dialogue to be:

Sally: Oh, okay. So he can't be in Boston. Nonetheless, when I said "Joe might be in Boston," what I said was true, and I stand by that claim.

I hope you'll agree that it would be odd and unnatural for Sally to say this.²⁷

Let's grant—although I think the point is questionable—that it would be "odd and unnatural" for Sally so to speak. Then that is evidence that we do not understand Sally's original assertion, "Joe might be in Boston", as making a claim about the limitations of her then current state of information, and hence also, if the standard semantics implies otherwise, evidence *against* the standard semantics. But how exactly is this evidence *for* a relativistic semantic treatment? The pattern of the exchange is simply

A: P

B: [new info]

A: I was wrong.

²⁷ MacFarlane (2003a).

together with an impression that it would be “odd and unnatural” for A to replace his second remark with

A: OK, not-P. Nevertheless when I said ‘P’, what I thereby said was true and I stand by it.

If the latter is incongruous, that does indeed tell against imputing utterance context-sensitivity to the content of tokens of ‘P’. Nevertheless: how could data as austere as this motivate any form of relativism? What is schematized is merely something that fits the withdrawal of any defeasibly grounded, non-context-sensitive assertion.

If a propensity to exchanges of the kind MacFarlane illustrates is to provide any motive at all to start thinking of the truth of epistemic possibility claims as context-of-assessment relative, it has to be supplemented. And what it needs supplementation with, if relativism is to be the right way to go, is evidence of some kind that, her withdrawal of it notwithstanding, Sally’s initial claim is also regarded by competent speakers as having been in some way *perfectly proper*—and this not just in the sense of having been justified though incorrect when she made it. Correctness varies with context of assessment: that claim is the very heart and soul of truth-relativism. So that is what we need to see evidenced in linguistic practice if linguistic practice is to provide evidence for relativism. We have to identify some feature of what strikes us as acceptable linguistic practice which would somehow evince practitioners’ recognition of the propriety of Sally’s initial claim consistently with that of her subsequent retraction of it. Absent any such feature, there is, so far as I can see, absolutely no call for a truth-relativistic response to the (alleged) linguistic data adduced by MacFarlane. We should merely conclude that Sally made an epistemic possibility claim at t_1 that, at t_2 , she recognizes to have been incorrect. She was just wrong before.

The point is simple and general: impressive evidence for the assessment-relativity of truth in some region of discourse has to include some kind of evidence of co-variation of truth-value with the values of an assessment-contextual parameter. You can’t make a case based purely on patterns of *correction*. There has also to be some kind of evidence of *concession* to the propriety of claims made in other assessment contexts—a manifestation of relativistic *even-handedness*. The pattern of linguistic practice—again, I am not, for the sake of argument, challenging its reality—adduced by MacFarlane is merely one-sided, and accordingly unconvincing.²⁸

Very well. So now we need to ask what in principle could amount to such even-handedness suggestive evidence. What is needed is some characteristic means whereby speakers can express an appropriate kind of equality of respect

²⁸ Indeed, it is not merely one-sided. It is actually an *awkwardness* for the relativist that Sally corrects her former self—“Oh. Then I guess I was wrong”—and is not content merely to revise her verdict: “Oh. Then I guess that is wrong”. I discuss this further in Wright (2007).

for contrasting opinions which originate in differences in the values taken by relevant parameters of assessment. But the problem is that if the discourse in question is *fully committed*, any assessor of such an opinion is liable to be stuck with an opinion of his own on the matter in question. Sally has no option, after George's intervention, but to regard her initial claim, that Joe might be in Boston, as false. So the strongest thing she could say by way of respect for it is that she was right to endorse it in the informational state she occupied when she made it. Unfortunately, that kind of endorsement does nothing to distinguish it from any claim that is appropriately advanced on the basis of certain evidence and subsequently withdrawn on the basis of altered evidence. It is, to repeat, no evidence of truth-relativity. Moreover—the crucial point—because discourse of epistemic possibility is fully committed, no other assessor of her claim is going to be able to do any better. Anyone in a position to judge that she spoke correctly on the basis of the information originally available to her will be mandated in an opinion of their own about Joe's possible whereabouts at the time. If it conflicts with hers, they too will be restricted to the form of endorsement of her claim—viz. that she was right to say what she did in the informational state she occupied at the time—which she herself can offer after George's intervention. That's of no interest, as emphasized. But if it coincides with hers, it is of even less interest, from the point of view of one seeking evidence for relativism.

Things look somewhat better if the discourse is neutral. In a neutral discourse, it is still true that a thinker, T, for whom a relevant assessment parameter shifts to a different value will be in no position to endorse a now discarded proposition in any way that evinces the distinctive form of concession to its former relative correctness that we are looking for. But what may be possible is for a third party, U, whose own assessment context takes no value for the relevant parameter, correctly to judge both that each of T's verdicts is correct in the context of assessment in which it is made and—the crucial extra—that neither of T's verdicts is superior in any respect to the other; or superior, at least, in any respect assessable by U. It is the unavailability of the latter claim in fully committed discourses that sets up the obstacle to the relevant kind of even-handed evidence.

However we need a final pair of cautionary distinctions. First, notice that even in such neutral cases, it will take collateral philosophical work to build a case out of the accepted propriety of the relevant kind of even-handed third-party assessments that a relativistic conception of truth operates in the discourse in question. The evidence will be just that a competent neutral assessor, U, can encounter cases where he can recognize first, that each of two conflicting verdicts is sanctioned by differing values of the relevant parameter, V—different standards, let's suppose—operating in the respective contexts of their authors; and second, that neither verdict is based on any shortcoming that U can detect from her position of standard-neutrality. That is perfectly consistent with U's position of standard-neutrality being conceived by practitioners of the discourse as one of *cognitive limitation*; more abstractly, it may be that assessment contexts

which are neutral on the parameter V are, just for that reason, thought of as contexts from which the fact that makes one or other of a relevant pair of conflicting verdicts absolutely correct is *invisible*.

Second, it is in any case—surely—not for *linguistic* evidence to decide whether or not a relativistic view of, say, epistemic possibility claims is appropriate. The most that might be manifest in features of our linguistic practice involving such claims is that *we understand them* as trafficking in assessment-relative contents. We are some considerable way short, I have been suggesting, of a model of the kind of linguistic data that might distinctively manifest even that. But even if we had such a model, and found it exemplified in speakers' practice, that would indicate no more than that they *implicitly conceived* of truth in the region of discourse in question as assessment-relative. That conception could still be philosophically completely unmotivated, or utterly mistaken.²⁹

7.7 POSTSCRIPT

As emphasized, once any global Protagorean ambition is abandoned, and truth-relativism is viewed as an at most locally appropriate view, ERPs, it is natural to expect, will in general be regarded as absolute. But there is one striking exception. The metaphysics of the Open Future is *prima facie* in tension with the so-called Truth-Value Links—the principles which articulate certain necessary connections of truth-value between variously tensed utterances conceived as made at different times. Consider for instance

“There is a Cup Final today” is true at Day₂ if and only if “There will be a Cup Final tomorrow” was true at Day₁.

At first blush, it looks as though this must fail left-to-right on the Aristotelian view—for isn't the whole point that the left-hand side can be true today while the quoted sentence on the right was not true but *indeterminate* yesterday? Yet a willingness to swallow that upshot

... appears incompatible with acknowledging the existence of a systematic link between the truth values

—and hence the contents—

of differently tensed statements uttered at different times.³⁰

²⁹ Thanks to participants at the Barcelona workshop and to members of the Arché 2005–6 Pilot seminar on Contextualism and Relativism for comments and questions on earlier drafts. Special thanks to Paul Boghossian and Sven Rosenkranz for detailed critical discussion and to John MacFarlane for generous written comments.

³⁰ Michael Dummett, in “The Reality of the Past”, 366–7 of his *Truth and Other Enigmas*.

and hence puts in jeopardy the entire basis of our understanding of tense. So what Aristotelians tend to say is that properly to absorb the idea of the Open Future and of the truth-conferrers for future contingents as subject to becoming must involve subjecting truth-conferral to the very constraints expressed by the truth-value links. Thus if “There is a Cup Final today” is true today, then it is also true that “There will be a Cup Final tomorrow” was true yesterday—but the latter is *itself* something that has only *become* true today, with the actual occurrence of the game. What *was* the case is likewise subject to becoming as the future unfolds.

Set this proposal in the context of a truth-relativistic treatment of contingent propositions about singular events, where the relevant assessment-contextual parameter is just time of assessment. We already noted that, so conceived, the discourse in question is fully committed. Now in addition, we are introducing second-order relativity. The ERPs for this discourse take the form:

“P” is (tenselessly) true at t.

And what, it is being suggested, goes with saving the intuitive semantics of tense in conjunction with proper absorption of the Open Future is to hold that the truth values of these ERPs are *themselves* relative to time of assessment. So while we can truly say *now* that it is true that

(I) “There will be a Cup Final tomorrow” was true yesterday,

the proper assessment *yesterday* of the same content, viz. that then expressed by a token of

(II) “‘There will be a Cup Final tomorrow’ is true today’,

would have rated it false; for “There will be a Cup Final tomorrow” was then indeterminate.

Well, this is a simple contradiction. For if “There will be a Cup Final tomorrow” was true yesterday, so was “‘There will be a Cup Final tomorrow’ is true today”. So in endorsing (I), we commit ourselves to saying that (II) was true yesterday. So we cannot consistently go on to say that it wasn’t!

This is a reinforced version of the predicament of the epistemic possibility relativist latterly reviewed. There, the fully committed nature of the discourse meant that when a pair of thinkers in distinct assessment contexts correctly return conflicting verdicts on an epistemic possibility claim, there is no way that a third party can avoid disagreement with one of them, so no possibility of the kind of even-handed assessment of their disagreement that relativism ought intuitively to make possible. What the third party can at least do, in that case, is to endorse each of the corresponding ERPs. In the present case, however, even this is not possible. Whatever my station in time, D_1 , it will enjoin not only a view—indeterminate if $D_1 < D_2$; or true, let’s suppose, if $D_1 \geq D_2$ —of the truth-status enjoyed by “There will be a Cup Final tomorrow” at D_1 but also a view—likewise indeterminate so untrue if $D_1 < D_2$; or true if $D_1 \geq D_2$ —of

the truth-status of “‘There will be a Cup Final tomorrow’ is true at D_1 ”. Accordingly, the Aristotelian relativist—at least so long as he determines to preserve the letter of the truth-value links—not merely lacks the resources for an even-handed evaluation of conflicting but respectively context-relatively correct claims about contingent events. He cannot so much as give consistent expression to their being context-relatively correct.³¹

References

- Boghossian, Paul, (this volume): “What is Normative Relativism?”.
- Cappelen, H and Lepore, E (forthcoming): “The Myth of Unarticulated Constituents”, in M. O’Rourke and C. Washington (eds.), *Essays in Honor of John Perry* Cambridge, Mass.: MIT Press.
- Dietz, Richard (this volume): “Epistemic Modals and Correct Disagreement”.
- Dummett, Michael: “The Reality of the Past”, in his *Truth and Other Enigmas*, London: Duckworth (1978) at 358–74.
- Egan, Andy (2005): “Epistemic Modals, Relativism and Assertion”, in Gajewski *et al.* (2005), 35–61.
- , Hawthorne, John, and Weatherson, Brian (2004): “Epistemic Modals in Context”, in G. Preyer and P. Peters (eds.), *Contextualism in Philosophy*, Oxford: Oxford University Press, 131–69.
- Gajewski, Jon *et al.* (eds.) (2005): *New Work on Modality*, MIT Working Papers in Linguistics, v. 51.
- Hawthorne, John (2004): *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Jordan, J. (1971): “Protagoras and Relativism: Criticisms Bad and Good”, *Southwestern Journal of Philosophy* 2: 7–29.
- Kölbel, Max, (2002): *Truth Without Objectivity*, London: Routledge.
- MacFarlane, John (2003a): “Epistemic modalities and relative truth”, MS, University of Berkeley. URL: <<http://philosophy.berkeley.edu/macfarlane/work.html>>
- (2003b): “Future Contingents and Relative Truth”, *The Philosophical Quarterly* 53: 321–36.
- (2005a): “The Assessment Sensitivity of Knowledge Attributions”, *Oxford Studies in Epistemology* 1: 197–233.
- (2005b): “Semantic Minimalism and Non-Indexical Contextualism”, in G. Preyer and G. Peter (eds.), *Content and Context: Essays on Semantics and Pragmatics*, Oxford: Oxford University Press.
- (2006b): “Epistemic Modals are Assessment-Sensitive”, version of August 23, 2006; URL: <<http://sophos.berkeley.edu/macfarlane/epistmod.pdf>>, MS, forthcoming in B. Weatherson and A. Egan (eds.), *Epistemic Modality* (preliminary title), Oxford: Oxford University Press.
- (this volume): “Truth in the Garden of Forking Paths”.

³¹ This point is elaborated in Wright and Moruzzi (2008).

- Meiland, J.W. (1977): "Concepts of Relative Truth", *The Monist* 60: 568–82.
- (1979): "Is Protagorean Relativism Self-Refuting?" *Grazer Philosophische Studien* 9: 51–68.
- Perry, John (1998): "Indexicals, Contexts and Unarticulated Constituents", in *Proceedings of the 1995 CSLI-Amsterdam Logic, Language and Computation Conference*. Stanford: CSLI Publications.
- Preyer, G. and Peter, G. (eds.) (2004): *Contextualism in Philosophy*, Oxford: Oxford University Press.
- Recanati, F. (2002): "Unarticulated Constituents". *Linguistics and Philosophy* 25: 299–345.
- Seigel, H. (1986): "Relativism, Truth, and Incoherence", *Synthese* 68: 295–312.
- Stanley, Jason (2000): "Context and Logical Form". *Linguistics and Philosophy* 23: 391–434.
- (2005): *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Stephenson, Tamina (2005): "Assessor Sensitivity: Epistemic Modals and Predicates of Personal Taste", in Gajewski *et al.*, 179–206.
- Swoyer, C. (1982): "True For", in M. Krausz and J.W. Meiland (eds.), *Relativism: Cognitive and Moral*, Notre Dame: University of Notre Dame Press, 84–108.
- Williamson, Timothy (2005): "Knowledge, Context and the Agent's Point of View", in G. Preyer and G. Peter (eds.), *Content and Context: Essays on Semantics and Pragmatics*, Oxford: Oxford University Press.
- Wright, Crispin (1992): *Truth and Objectivity*, Cambridge, Mass.: Harvard University Press.
- (2001): "On Being in a Quandary: Relativism, Vagueness, Logical Revisionism", *Mind* 110: 45–98.
- (2006): "Intuitionism, Realism, Relativism and Rhubarb", in P. Greenough and M. Lynch (eds.), *Truth and Realism*, Oxford: Oxford University Press.
- (2007): "New Age Relativism and Epistemic Possibility", *Philosophical Issues* 17.
- and Moruzzi, Sebastiano (2008): "Trumping Assessments and the Aristotelian Future", in a special number of *Synthese* on Relativism, Berit Brogaard (ed.).
- (forthcoming): "Fear of Relativism?", contribution to a book symposium on Paul Boghossian's *Fear of Knowledge* in *Philosophical Studies*.

This page intentionally left blank

Three Forms of Truth Relativism

Iris Einheuser

The received view about propositions—what is said by indicative sentences when assertively uttered on an occasion of use—is that they have their truth-value absolutely: a proposition represents a way for part of the world to be. If the relevant part of the world is as the proposition represents it as being then the proposition is true, otherwise it is false. Several authors have recently challenged the received view, arguing that at least in some areas of discourse propositional truth ought to be *relativized* to contexts of assessment.¹ Amongst the areas of discourse which appear to require a relativization of propositional truth are discourse concerning subjective matters such as taste and humor as well as discourse involving epistemic modals.

In this paper, I explore an alternative account of the data appealed to by propositional relativists, an account which allows us to hold on to the standard Fregean view of propositions as absolute with respect to some of the controversial areas of discourse.² The account proposed, *factual relativism*, involves two key ingredients. First, an analysis of the relevant facts—facts about taste and epistemic facts—as partially “configured” through or “structured” by a perspective. Second, attention to the distinction between the propositional *content* of an utterance and the *judgment* expressed with the utterance. The apparatus employed by factual relativism allows us to model, in an intuitively appealing way, the contents of assertions about subjective matters of fact as well as the states of affairs asserted to obtain.

I will start by developing a factual relativist account of matters of taste and apply it to the data which might tempt one to relativize propositional truth. To show the potential of the approach, I sketch an analogous account of the facts

¹ See, for instance, MacFarlane [7] and [8], Kölbel [5] and Egan, Hawthorne and Weatherson [2]. Numbers in brackets refer to the References list at the end of chapter.

² I do not believe that every area of discourse for which a relativization of propositional truth has been urged can be analyzed as proposed here. In particular, future contingents (see MacFarlane [6]) do not appear to admit of such an analysis. Therefore, this paper presents at best a partial case against deviating from the Fregean orthodoxy.

reported by statements involving epistemic modals. I then compare propositional with factual relativism, examine the assumptions underlying both and argue that factual relativism is preferable as it preserves the close link between the propositional content of a statement and the fact represented by that content.

8.1 MODELING MATTERS OF TASTE

People differ in their assessment of what tastes good. Consider: Moira asserts

(1) Vegemite tastes great.

Clem retorts

(2) No it doesn't, Vegemite tastes awful.

Here, we take Moira and Clem to be disagreeing about whether Vegemite tastes great. Further, most of us are willing to grant that none of them has made a mistake. Their disagreement is, as is sometimes said, *blameless*.³

8.1.1 Three Strategies

Three types of strategy have been pursued to give a semantics for statements like (1) which do justice to our intuitions about exchanges like the above. (i) *Expressivism* has it that such statements do not express truth-evaluable contents at all. Rather, they merely express a positive or negative attitude to the taste of Vegemite. In what follows I will ignore this strategy. (ii) *Indexical relativism*, or *contextualism*, has it that (1) expresses different propositions on different occasions of use.⁴ One familiar way to implement this strategy is to take (1) to express roughly the same as 'Vegemite tastes great to me'. On this strategy, disagreement about matters of taste is merely apparent: When Moira assertively utters (1) and Clem assertively utters (2) they do not express contradictory contents. Finally, (iii) *propositional relativism*, has it that while (1) expresses the same proposition across different occasions of use, the proposition it expresses has different truth-values relative to different contexts of assessment. It is true when evaluated relative to Moira's context and false when evaluated relative to Clem's even though the world remains the same in all relevant respects. This represents a clear departure from the standard view according to which the truth of a proposition depends only on what the world is like.

I will develop a further strategy, *factual relativism*, which has it that the relevant *facts* vary across contexts of assessment, in a sense to be made precise.

³ See Kölbel [5]. I am not concerned here with arguing that the dialogue does in fact present a case of blameless disagreement. Rather, I will assume that it does and explore the theoretical options for accounting for it.

⁴ Cappelen and Lepore explore a variation of indexical relativism according to which (1) expresses different propositions on different occasions of *assessment*. See their paper in this volume for details.

8.1.2 Methodology

In our investigation, we need to keep two issues separate. On the one hand there is the *metaphysical* question of the nature of the facts in the area of interest, in our case the facts concerning taste. These are the things in the world that the truth of our statements is sensitive to. Second there is the *semantic* question of how the truth conditions of statements about the area of interest are to be given in terms of the facts of that area. *Facts* and *meanings* are theoretical concepts introduced to systematically account for our linguistic behavior and cognitive interaction with the world. As usual, there is a trade-off between complicating the metaphysics while simplifying the semantics on the one hand and complicating the semantics while simplifying the metaphysics on the other. I see the propositional relativist as complicating the semantics, making propositional truth depend on an additional parameter. Factual relativism, on the other hand, complicates the metaphysics and so can give a simpler account of the semantics of the statements in question.

There are, however, constraints on what trade-offs we should be willing to accept. A complication in the metaphysics should be justifiable independently, as should be a deviation in the semantics. In Section 3 I will argue that, unlike factual relativism, the modification proposed by the propositional relativist gives propositions a theoretical role without much intuitive anchoring.

8.1.3 The Meaning of Taste Statements and the Facts they Represent

Statements like (1) appear to report facts about what tastes good *simpliciter* (call these *general* taste facts) rather than facts about what tastes good to some particular group of tasters (call these *particular* taste facts). However, matters of taste are largely subjective. Whether a particular item tastes good depends not just on the physio-chemical features of that item, but also on physiological and psychological features of tasters.

Since matters of taste are subjective, one may hold that there are no general but only particular taste facts. Taken at face-value, all positive general taste statements would then be false, as the general facts they report do not obtain. In analogy with a similar thesis in ethics, this account would have to be accompanied with an error-theory: Speakers of English regularly and unknowingly utter falsehoods when they say of an item that it tastes good. Alternatively, general taste statements may be reinterpreted and their truth conditions accounted for in terms of the particular taste facts. Options for how this might be accomplished include the following:

- 'x tastes great' is true when uttered by *i* iff *x* tastes great to *i*
- 'x tastes great' is true iff $\forall t : x$ tastes great to *t*
- 'x tastes great' is true in context *c* iff $\forall t \in T : x$ tastes great to *t*, where *T* is a contextually salient community of tasters.

These all amount to a form of indexical relativism with the context of utterance determining the content expressed and thus the particular taste facts relevant for evaluating the statement. None of these is entirely satisfactory. On the first construal, Moira's and Clem's utterances do not express contradictory contents. They do not disagree. On the second construal, only very few, if any, positive general taste statements are ever true. Propositional relativists argue that the third construal clashes with the data concerning typical conversational exchanges about matters of taste.⁵

I propose a different stance. Matters of taste are subjective, but that does not force us to deny the existence of general taste facts: While there are no *objective* general taste facts, we may admit that there are *subjective* general taste facts. *These* are the facts represented by general taste statements. If we accept this option, our task consists in (i) giving an account of how to think about these facts, and (ii) giving a semantics for statements which represent such facts. The next two sections will tackle the first sub-task while the subsequent two address the second one.

8.1.4 The Structure of Taste Facts

Subjective general taste facts are determined by a combination of two factors: The world itself, independent of any particular tasters, determines the physio-chemical features of items on the bases of which tasters classify them as tasting good or bad. It also determines the physio-chemical features of individual *tasters* which are responsible for their classifying certain items as tasting good or bad. The world thus uniquely determines the particular taste facts yet not the general ones. Call the contribution made by the world a *substratum*. The second contributing factor is a collection of physiological and psychological features of *particular* tasters. I will call the contribution of particular tasters a (*taste-*)*perspective* and say that a perspective *induces* taste facts over a substratum. I don't want to commit myself here to a specific analysis of what the relevant substrata and perspectives are and how the two interact to induce taste facts. The details of that story depend on the physiology of taste and are beyond the scope of the present paper.

The taste facts jointly determined by the substratum and a perspective are *subjective* or *perspectival*, depending not just on the world but also on the perspective that induces them. Over the same substratum, one perspective induces the (perspectival) fact that Vegemite tastes great, while another perspective (mine for instance) induces the (perspectival) fact that Vegemite tastes awful. In effect, the analysis given reconceptualizes *objective particular* taste facts as *subjective general* taste facts.

⁵ See Kölbel [5].

8.1.5 Representing Possible Worlds

I now introduce the representational apparatus to model possible worlds in which some facts are induced by a perspective. Let S be the collection of possible substrata, that is the possible contributions the world itself can make to the taste facts. A substratum, we may assume, determines all the physical facts of a world, including the distribution of physio-chemical properties over potential candidates for taste-ascriptions. We may further assume that it fixes all the *particular* taste facts such as the fact that Vegemite tastes great to Moira. Let P be the collection of possible perspectives—collections of the features of a particular taster (or homogeneous group of tasters) which determine how things taste to that taster.

A possible world with determinate general taste facts can be represented by a pair $\langle s, p \rangle$, where s is a substratum and p is a perspective. The actual substratum from my perspective, represented by $\langle s@, p_I \rangle$ —with $s@$ being the actual substratum and p_I being my perspective—is a world in which the general taste facts are determined by my sensibilities. It differs from, say, the actual substratum from Moira's perspective, $\langle s@, p_M \rangle$: In $\langle s@, p_I \rangle$ Vegemite tastes awful while in $\langle s@, p_M \rangle$ it tastes great. Different possible worlds with the same substratum coincide in their objective facts, but differ in their subjective (general) taste facts, as these are induced by different perspectives. For ease of expression, I will use '*the world*' to refer to the actual substratum, so that '*the world from Moira's perspective*' refers to $\langle s@, p_M \rangle$.

So far I have talked about the nature of the relevant facts. They are understood as subjective and modeled by combining a subjective taste-perspective with an objective substratum. What is the semantic pay-off of this construal? Is it worth it, from a semantic point of view, to buy into the metaphysics of subjective facts?

8.1.6 Propositions and Assertions

Within the above framework, the propositional contents of statements concerning matters of taste can be modeled in the familiar way as functions from worlds—worlds in which a substratum and a perspective jointly determine the general taste facts—to truth-values. The proposition expressed by 'Vegemite tastes great' is true relative to a world just in case Vegemite tastes great in that world. There is no need to further relativize propositional truth to a context of assessment as the relativizer has been built into the worlds. This allows the *contents* of statements representing subjective facts to abstract away from particular subjective perspectives and thereby capture the common feature of perspectival worlds which coincide as far as the subjective facts in question are concerned.⁶

⁶ I will look more closely at the representationality of such statements and the subjective–objective issue in Section 3.

On the proposed account, Moira and Clem do express contradictory propositions when they utter (1) and (2) respectively. However, we do need a further semantic notion to fully account for our intuitions about their exchange. For not only do they contradict each other. We also want to grant that—given Moira’s love for Vegemite and Clem’s hatred of it—each of them has, in a sense, said something true. Something “true for them” as it were. But what does that mean? None said something true about the world *simpliciter* because the world *simpliciter*, the substratum, does not determine any general taste facts. The sense in which Moira said something true when she uttered (1) is that she said something true about the world structured by her perspective. And Clem said something true about the world structured by *his* perspective. That might look problematic: If what Moira said is answerable to the facts in “her” world, while what Clem said is answerable to facts in “his” world, then how can it be that they said exactly the same? Does this not amount to a form of indexical relativism (‘In my world, Vegemite tastes great’)?

It does not. Here, we have to take seriously the distinction between the *propositional content* of an utterance as used on an occasion and the *assertion* a speaker makes by means of that utterance—the *judgment* he expresses. When a speaker judges that *p*, he in effect judges that *p* holds in the world actual for him. For example, when I judge that Mont Blanc is the Alps’ highest peak, I in effect judge that in the actual world, Mont Blanc is the Alps’ highest peak. That does not mean, however, that reference to the actual world figures in the propositional content of an actual utterance of ‘Mont Blanc is the highest peak in the Alps’. The statement ‘It is necessary that Mont Blanc is the highest peak in the Alps’ is, after all, not true. If the necessity-operator is understood as operating on the propositional content of the embedded sentence, then the proposition expressed by ‘Mont Blanc is the highest peak in the Alps’ had better not be world-bound. When my “twin” in a possible world in which the Matterhorn is the Alps’ highest peak utters ‘Mont Blanc is the highest peak in the Alps’, she expresses the same proposition I express when uttering that sentence, yet she asserts it of her world and thus she asserts something false.

Similarly with the exchange between Moira and Clem: They both express the same proposition when uttering ‘Vegemite tastes great’. But Moira expresses a judgment about her world while Clem expresses a judgment about his—they apply the same propositional content to different perspectival worlds. Thus, the factual relativist move can be seen as a natural extension of the standard account: Sentences without any context-sensitive elements express the same proposition regardless of the possible situation in which they are uttered, yet the truth-values of the propositions expressed vary across these situations.⁷

⁷ Here, I assume that the same semantic rules are at work in these counterfactual situations.

8.1.7 Accommodating Faultless Disagreement

Within the framework proposed here, we can distinguish between three degrees of disagreement. The first degree consists in mere syntactic contradiction: Two people disagree in this sense when they utter sentences one of which is a negation of the other. The second degree consists in asserting contradictory contents, that is contents which represent incompatible facts. The third and highest degree consists in asserting contradictory contents of the same world.

Indexical relativism accommodates faultless disagreement by capturing the first and weakest degree of disagreement: The speakers disagreeing utter contradictory sentences, but these sentences receive interpretations on which the contents expressed are not incompatible. Factual relativism does not, like indexical relativism, factor the relativizer into the proposition expressed (as in *Vegemite tastes good in my world*) but into the act of asserting a judgment: A speaker asserts a statement with a constant propositional content (such as *Vegemite tastes great*) of her world. This lets us capture disagreement of the second degree: The speakers disagreeing utter contradictory sentences, and these sentences receive interpretations on which the contents they express are incompatible. What allows for the disagreement to be faultless is the fact that the speakers might correctly judge the contradictory contents expressed to hold of their respective worlds. Propositional relativism seeks to accommodate the highest degree of faultless disagreement: The disagreeing speakers express contradictory propositions and judge them to hold of the same world. What allows for the disagreement to be faultless is that whether a proposition is true of a world depends on the perspective from which the proposition is assessed. In the last section of this paper, we will consider what assumptions about the relationship between a perspective, a proposition and what it represents we have to make in order to accommodate faultless disagreement in the way proposed by the propositional relativist.

8.1.8 Modalizing Matters of Taste

Building perspectives into possible worlds gives rise to a problem.^{8,9} After dinner, my friend complains about my—in her view—excessive use of garlic and remarks

(3) This dish might have tasted great.

What she wants to convey is that if the dish had had different physio-chemical properties (on account of having been prepared with a good deal less garlic), the dish would have tasted great to her. However, if the possibility-operator is understood in the usual way as an existential quantifier over all possible worlds, including those structured by other people's taste-perspectives, then (3) is true

⁸ This section draws heavily on section 4.1 of my [3].

⁹ Egan *et al.* [2] raise a similar difficulty for a related account.

on account of, say, *my* liking lots of garlic in my food. So witnesses to (3) include worlds in which the dish tastes just as overly garlicky to my friend as it actually does. Thus, (3) comes out true for the wrong reasons.

Our two-dimensional model of taste facts provides the means to overcome this problem. Once we have analyzed taste facts into two components, we see that there are two ways in which it is possible for a general taste-fact, such as the fact that Vegemite tastes great, to obtain or fail to obtain: on the one hand, the *substratum* of the world might differ while the same perspective is brought to bear. Suppose the dish was prepared with less garlic and so was very much to my friend's liking whose taste-perspective remains as it actually is. In this case, the dish would taste great in the standard sense (3) appears to aim at. On the other hand, the taste facts of a world might be determined by a different *perspective*. Suppose the dish had all the physio-chemical properties it actually has, but that the taste-perspective brought to bear was that of a garlic-lover. This taste-perspective determines together with the—from my friend's point of view bad-tasting—substratum that the dish tastes great.

We introduce two sets of modal notions to capture the two dimensions along which taste facts can vary: First, we have the notion of *s*-possibility $\Diamond_s$ which tracks possible variations of the substratum. Second, we have the notion of *p*-possibility $\Diamond_p$ which tracks possible variations of the perspective:¹⁰

$\Diamond_s \varphi$ is true at a world $w = \langle s, p \rangle$ if and only if there is a substratum s' such that φ is true at $w' = \langle s', p \rangle$ (i.e. if and only if φ is true at some world that differs from w only in virtue of its substratum).

$\Diamond_p \varphi$ is true at a world $w = \langle s, p \rangle$ if and only if there is a perspective p' such that φ is true at $w' = \langle s, p' \rangle$ (i.e. if and only if φ is true at some world that differs from w only in virtue of its perspective).

Now when my friend utters (3) with the intended meaning suggested above, she employs the notion of *s*-possibility, quantifying only over possible worlds whose taste facts are determined by her own taste-perspective: 'That dish might have tasted great—if only you had used less garlic'. This is, I believe, the more familiar kind of possibility in connection with statements about taste. But on occasion we might have in mind the notion of *p*-possibility when deliberating about how something might have tasted. Consider: As a child I was made to sit through a big raw oyster session. Now when faced with a raw oyster on a half-shell I say

(4) This might have tasted great

but given my conditioning, the taste of oysters is lost on me.

Note that the apparatus introduced allows for a straightforward semantics for the operator 'on any standard' which some people have discussed. The statement

(5) On any standard, x tastes great

¹⁰ For more on these two notions of modality see my [3].

is true if and only if $\ulcorner x \text{ tastes great} \urcorner$ is true relative to every standard of taste. The semantics of this operator is the same as that of the p -necessity operator:

$\Box_p \varphi$ is true at a world $w = \langle s, p \rangle$ if and only if for any perspective p' : φ is true at $w' = \langle s, p' \rangle$ (i.e. if and only if φ is true at every world that differs from w only in virtue of its perspective).

As expected, a statement claiming that a given item is tasty on any standard is true in a perspectival world iff it is true in every perspectival world with the same substratum.

8.1.9 Evaluating Taste-Beliefs and Assertions

In the previous section we saw how statements concerning subjective matters of fact embed under modal operators. Now I want to examine how they behave in belief-contexts and how that behavior is best modeled. In addition to the apparatus introduced above, the distinction between the propositional content of a belief and the judgment on the basis of which the belief is held helps shed light on the issues cross-perspective belief-evaluation give rise to.

The indexical relativist has to work around a difficulty with *ascribing* beliefs about subjective matters of fact across perspectives. Suppose Clem correctly asserts

(6) Moira believes that Vegemite tastes great.

On a simple-minded version of the contextualist account, Clem's utterance 'Vegemite tastes great' expresses the proposition that Vegemite tastes great to Clem. But this is not what Moira believes. To properly account for cross-perspective belief-ascription, the indexical relativist has to modify the semantics of propositional attitude constructions, accommodating possible shifts between the context in which the believer sincerely expresses her beliefs and the context in which the belief-ascriber sincerely attributes these beliefs.

Factual relativists do not face that problem as statements concerning subjective matters have the same content no matter who asserts them. However, on the factual relativist account we get an analogous problem, not with *ascribing* beliefs but with *evaluating* beliefs across perspectives. Suppose, again, Clem correctly asserts (6). While he correctly characterizes the *content* of Moira's belief, it is not clear how he should *evaluate* it. What is in question here is whose perspective is relevant when evaluating a belief held by someone else. There are two salient options: The relevant perspective may be that of the assessor's or that of the believer.

Suppose the relevant perspective is the believer's. Then Clem can infer

(7) Moira believes something true

from which, together with (6), he can further infer

(1) Vegemite tastes great.

But (1), from Clem's perspective, is not true. We run into a parallel problem if we assume that the perspective relevant for assessing a belief held by someone with a different perspective is the assessor's. In this case, Clem could infer

(8) Moira believes something false

from which, together with the fact that Moira has all the relevant evidence he can infer that Moira suffers from a cognitive shortcoming: She forms a belief which is false given all the objective evidence available to her. But that appears to be an unfair assessment of Moira's belief-forming mechanism. If anything, Clem might be justified in accusing Moira of poor taste—but certainly not for employing faulty doxastic procedures.

At the root of the problem is the failure to take into account the fact that Moira's belief is a belief about the world as structured by *her* taste-perspective. This, as argued above, is not reflected in the *content* of the embedded sentence. The evaluation of a perspectival belief should therefore be sensitive to more than just the content of the belief.

Consider again the ascription of beliefs to agents in counterfactual situations: My "twin" who lives in a possible world in which the Matterhorn is the Alps' highest peak believes that the Matterhorn is the Alps' highest peak. Although what she believes is actually false, it is true of her world. When evaluating my twin's belief, we could be interested in two different things: First, we might evaluate the propositional *content* of her belief (which is false, that is actually false). Or we might evaluate the *judgment* on the basis of which she holds the belief. Here, we have to take into consideration that she, in effect, judges the proposition the proposition in question to be true of her world. I suggest that we continue to evaluate a belief as true (false) if it is true (false) in the evaluator's world, and that we evaluate a belief as correct (incorrect) if its propositional content is true (false) in the believer's world—the world with respect to which the belief was formed.

Applying this distinction to beliefs about subjective matters, there are two things for us to evaluate: First, we can evaluate the propositional *content* of the belief. By default, we do so with respect to the world structured by our own perspective. Second, we can evaluate the *judgment* that underlies the belief. We do so with respect to the world structured by the believer's perspective. Thus, when Clem evaluates Moira's belief that Vegemite tastes great, he ought to say that Moira believes something false but that her belief is nonetheless correct.¹¹ The same distinction applies to the cross-perspective evaluation of *assertions* concerning subjective matters: We may evaluate someone else's taste-assertion as *true* when the assertion's propositional content is true of our—the evaluators'—world and as *correct* when it is true of the asserter's world.

Whether we evaluate with an eye to correctness or truth will depend on the purpose of the evaluation. If, for instance, we are interested in explaining the

¹¹ Moira judges that in "her" world Vegemite tastes great, which it does.

agent's behavior and beliefs, we will evaluate for correctness, that is in terms of the world about which they believe and assert certain things. If, on the other hand, we want the beliefs of others to provide guidance to what is the case in our own world, we will evaluate the contents of their beliefs for truth, that is with respect to our own perspectival worlds.

The analysis given can be generalized to other areas of discourse, such as discourse about what is funny, interesting, attractive, etc. and thus allows us to account for discourse about subjective matters without giving up the standard picture of propositional truth as sensitive only to what the facts are. Further, a structurally analogous account can be given of discourse involving epistemic modals—another area which some argue requires us to relativize propositional truth. I will briefly sketch how such an account might go.

8.2 MODELING EPISTEMIC MODALITY

The standard approach to the semantics of epistemic modals, *epistemic contextualism*, is a form of indexical relativism. An utterance of

(9) It might be that *p*

in which the 'might' is understood as expressing epistemic possibility is construed as having the truth conditions of

(10) What the currently relevant epistemic community knows or can come to know in relevant ways does not rule out that *p*

where what the relevant epistemic community is differs across variations of the indexical relativist account.¹²

This has recently been challenged.¹³ The data which calls into question the indexical relativist account consists mainly of cases in which, when assessing an utterance like (9), we don't appear to take into account the epistemic community relevant in the context of utterance. Consider:¹⁴ Sally says

(11) The Red Sox might win the game on Friday

to which Tom replies

(12) That's impossible! Their three best players have just been hospitalized with food poisoning

in response to which Sally retracts her earlier assertion:

(13) Then what I said was wrong. It is not the case that the Red Sox might win.

¹² See DeRose [1].

¹³ See MacFarlane [8] as well as Egan *et al.* [2].

¹⁴ See MacFarlane [8] for variations of this kind of counterexample to various kinds of indexical relativism.

On the indexical relativist account¹⁵ Sally has no reason to retract her assertion, for what the relevant epistemic community knew when she uttered (11) did not rule out the Red Sox' winning. If (11) was correctly interpreted along indexicalist lines, the proposition she expressed remains true, even though an utterance of (11) made *now* that she obtained additional information would express a false proposition. So either indexical relativism is false or it predicts that speakers of English are confused about the truth conditions of some of what they say.

The propositional relativist suggests that the proposition expressed by (11) remains constant across contexts of utterance,¹⁶ yet the proposition expressed receives a different truth-value depending on what context it is assessed from. Assessed relative to the context in which Sally originally uttered (11), the proposition is true. Assessed relative to the context in which Sally utters (13), and in which additional information is available, the same proposition is false.

8.2.1 Modeling Epistemic Possibility

What I take the propositional relativists' examples to show is that statements like (9) don't report *particular* epistemic facts such as the fact that as far as Sally or some other epistemic community knows, the Red Sox might win. Rather, they are intended to report *general* epistemic facts. I propose that such facts be construed as perspectival, and as jointly determined by two factors: the world as it is independently of what any particular group knows on the one hand and the justified beliefs of a particular epistemic community on the other. As before, call the contribution made by the world the *substratum*, and let the *epistemic perspective* (or just *perspective*) of an epistemic community (e.g. that of a single believer or a group of believers) be the collection of all propositions justifiably believed by that community.¹⁷

The world *simpliciter* determines, we may assume, all the non-epistemic facts, the particular epistemic facts as well as what justified beliefs of a given epistemic community amount to knowledge. But it does not determine the *general* epistemic facts because it does not privilege any particular epistemic community. Again, we reconceptualize objective particular epistemic facts as subjective general epistemic facts.

8.2.2 Representing Possible Worlds and Propositions

We model possible worlds as centered on epistemic perspectives. Let *S* be the collection of substrata, and let *P* be the collection of possible epistemic perspectives, collections of propositions which may jointly be justifiably believed.

¹⁵ . . . and supposing that the relevant epistemic community consists only of Sally; if the relevant community is taken to be more comprehensive, the example would have to be adjusted.

¹⁶ . . . modulo the reference of 'next Friday'.

¹⁷ I'm not concerned here with the question as to what exactly a *collective belief* (a belief held by a group of people) amounts to.

For $s \in S$ and $p \in P$, $\langle s, p \rangle$ represents the perspectival world in which the epistemic facts are determined by the substratum s and epistemic perspective p . Propositions involving epistemic modality can now be modeled as functions from possible worlds to truth-values: Such a proposition is true in a world just in case the world is as the proposition represents it as being. For instance, (11) is true in a world $\langle s, p \rangle$ just in case s and p jointly make it the case that the Red Sox might win, that is just in case the justified beliefs determined by perspective p which amount to knowledge relative to $\langle s, p \rangle$ leave it open whether or not the Red Sox will win.

8.2.3 Evaluating Epistemic Modal Assertions and Beliefs

As in the case of assertions about matters of taste, epistemic modal assertions can be evaluated either for truth—relative to the evaluator’s epistemic perspective—or correctness—relative to the asserter’s epistemic perspective. When we evaluate an epistemic modal assertions, including an earlier assertion of our own, in the context of deliberating about what to believe and how to act given our current situation, we evaluate the assertions for truth. When evaluating such assertions in the context of deliberating about whether the assertions were reasonable, we evaluate the assertion for correctness.

Against this background, let us analyze the above exchange. We may assume that when Sally utters (11), none of the relevant community’s justified beliefs which amount to knowledge rule out that the Red Sox might win. Thus the world relative to which she judges that the Red Sox might win, $\langle s_@, p_S \rangle$ (where $s_@$ is the actual substratum and p_S is Sally’s then current epistemic perspective) is indeed one in which the Red Sox might win. Her assertion, the application of the propositional content of (11) to $\langle s, p_S \rangle$, is correct. Then Sally receives new information. Her epistemic perspective changes, and so do the epistemic facts in “her” world (as well as some of the non-epistemic ones; e.g. it is now a fact that Sally has received a certain additional piece of information). When receiving the new information she judges that *what* she said before is false. This amounts to the judgment that what she said then is false of the world in which the epistemic facts are determined by her current perspective. Again, while the content of (11) remains constant throughout the exchange, the facts relevant for judging whether the content correctly represents the world change. The reason Sally retracts her claim after receiving additional information is that the epistemic facts are different from the perspective from which she re-evaluates her earlier claim.

8.3 RELATIVISM: INDEXICAL, PROPOSITIONAL OR FACTUAL?

Relativists of all three varieties agree that the truth of a sentence like (1) and (11) depends on something like a perspective, say a taste-perspective or an epistemic

perspective. They disagree about exactly what role the perspective plays in determining the sentence's truth-value. *Indexical relativists* factor the perspective into the determination of the proposition expressed. *Propositional relativists* put it alongside a world, as a further parameter to which the truth of the proposition expressed is sensitive. *Factual relativists* factor it right into the world.

The problem with indexical relativism with respect to the areas of discourse discussed is that it re-interprets overtly *general* taste- or epistemic modal statements as *particular* ones and therefore both understates the force these statements are meant to have in the mouth of the speaker and fails to account for the phenomenon of faultless disagreement.¹⁸

Propositional relativism is meant to overcome these problems. But it is one thing to introduce a further formal parameter and make the evaluation of propositions sensitive to it in one's semantics. It's another thing to explain how the additional parameter discharges its role. How are we to think of the perspective-relativity of propositional truth?¹⁹ Factual relativism offers an answer. It folds the perspective into the way the world is. Different perspectives, different facts, different truth-values. In defense of their proposal, propositional relativists sometimes point to an analogy between perspectives and possible worlds: The idea that propositional truth is relative to some parameter is not in itself objectionable since even on the standard account, propositions have truth-values relative to possible worlds.²⁰ However, the analogy between possible worlds and perspectives had better not be too close, for propositional truth is relative to possible worlds because the latter determine a range of facts. If perspectives merely do more of the same work possible worlds do, then propositional truth is relative to perspectives because the *facts* are relative to a perspective and so propositional relativism reduces to factual relativism. So for propositional relativism to be distinct from factual relativism, the role of perspectives has to be construed differently. But how?

8.3.1 What is Represented in Propositional Relativism

According to the standard account, propositions are the representational contents of indicative sentences on an occasion of use and so are intimately linked with the states of affairs they represent. I will consider two options for what role perspectives might play in the propositional relativist framework and argue that

¹⁸ See Kölbel [5] and MacFarlane [8].

¹⁹ Propositional relativists don't address that question. They focus attention on what it is to *assert* a perspective-relative proposition, arguing that assertions of perspective-relative propositions are governed by essentially the same standards as the assertion of absolute propositions and so are still recognizably assertions. See Kölbel [5] and MacFarlane [7].

²⁰ Note that *this* sort of relativity does not threaten the conception of propositions as absolute in the sense intended at the beginning of the paper: The truth-value of a proposition is fixed absolutely by the facts—that is the *actual* facts. Truth-values with respect to non-actual possibilities merely fix on the representational content of the proposition.

neither is entirely satisfactory: On the first, propositions don't represent any determinate states of affairs. On the second, they represent what appear to be the wrong states of affairs.

First, perspectives might select *what* state of affairs a proposition represents as obtaining. Then different facts are relevant for the evaluation of the proposition relative to different perspectives. In the case of propositions concerning matters of taste, the relevant facts are *particular* rather than *general* taste facts. For example, the fact on account of which *Vegemite tastes great* is true relative to Moira's perspective is the fact that Vegemite tastes great to Moira and the fact on account of which *Vegemite tastes great* is false relative to Clem's perspective is the fact that Vegemite tastes awful to Clem. Here, the proposition itself does not have any determinate content: perspectives, by selecting a particular taste fact relevant for the evaluation of the proposition, fix the representational content of the proposition. On this model, proposition and representational content are related, via perspectives, like Kaplanian character is to content via context.²¹ This is a significant departure from the standard account of the role of propositions. As a theoretical construct, a new notion of proposition could be introduced on which propositions are not tied to fixed representational contents. But while such a notion of proposition may do useful technical work in a formal semantic theory, it would be drained of much of its intuitive content.

The second possible role perspectives might play in determining propositional truth is best introduced by analogy. Consider the pictorial representation of objects and their shapes. Material objects have objective shapes.²² We may represent objects and their shapes by 2-dimensional pictures. Pictures always represent the shape of objects relative to a perspective. Suppose we want to represent the writing-surface of a particular desk, a rectangular slab of wood. What the picture looks like and whether it correctly represents the shape of the desk depends on the perspective from which the picture is supposed to represent the piece of wood. A given picture is a correct representation relative to some perspectives but an incorrect representation relative to some other perspectives. While there is an objective underlying shape which we aim at representing, we may have to distort the representing image to account for the perspective from which the shape is to be represented.

The propositional relativist may think of the facts to be represented, general taste facts or general epistemic facts, as analogous to objective shapes. They are objective features of the world and the fact that the world has these features can be represented from different perspectives. Here, perspectives are thought to provide some kind of distorting lens which maps the propositional representation of a state of affairs onto an objective fact. For instance, Moira's

²¹ See Kaplan [4].

²² Setting aside possible worries about relativity theory.

perspective maps the proposition expressed by ‘Vegemite tastes great’ onto some objective facts concerning the physio-chemical properties of Vegemite. Clem’s perspective maps the proposition expressed by ‘Vegemite does not taste great’ onto the same objective facts. The superficially contradictory propositions both represent the same underlying facts but they do so from different perspectives, just as differently shaped images represent the same objective shape from different perspectives.

The second possible role taste- and epistemic perspectives might play in determining propositional truth, then, is analogous to the role spacial perspectives play in determining whether a shape-representation is correct. On this construal, the state of affairs a general taste proposition intuitively *appears* to represent—*Vegemite tastes great*—is not the state of affairs it actually represents—*Vegemite has physio-chemical structure xyz*. Further, some contradictory propositions—*that Vegemite tastes great* and *that Vegemite does not taste great*—represent the same state of affairs. None of this shows that propositional relativism is untenable, but it reveals how much of a departure it requires from ordinary ways of thinking about how propositions behave.

8.3.2 What is Represented in Factual Relativism

The factual relativist views matters differently. He thinks of the represented facts as *perspectival*. The proposition that Vegemite tastes great does not represent the underlying physio-chemical facts concerning Vegemite. Rather, it represents that Vegemite tastes great. This perspectival fact is jointly determined by the objective physio-chemical facts and a perspective. In the spacial analogy, this corresponds to interpreting a rhomboidal picture not as representing that a given object has such and such objective shape, but as representing that the object *looks* rhomboidal. It doesn’t answer the question *what is the object’s shape?*, but *how does it look?* Similarly, the proposition that Vegemite tastes great is not meant to be an answer to the question *what is Vegemite’s physio-chemical make-up?* but rather to the question *how does it taste?* These *how*-facts go beyond the underlying objective facts in that they depend on particular perspectives.²³

8.3.3 Concluding Remarks on Factual Relativism

Factual relativism provides a model both for the *contents* of assertions about subjective matters and for the *states of affairs* that such assertions assert as obtaining. Both aspects offer advantages over propositional relativism. First,

²³ Propositions expressed by sentences containing epistemic modals can be understood along the same lines: The proposition that the Red Sox might win does not represent what the relevant portion of the world itself is like—physical condition of the players, trainers’ strategies, weather—but *how* it appears from a given epistemic perspective.

sentences reporting subjective matters mean what they appear to mean, and propositions play the role they were originally introduced to play. They can represent subjective facts while abstracting away from particular perspectives, representing a particular ‘look’ the ground-level facts might have from different perspectives. Second, subjective facts are treated as facts. On the propositional relativist account, general taste facts are eliminated and general taste propositions are made true, relative to a perspective, either by particular taste facts or by facts about the physical features of the world. But from a subjective point of view, that a certain item tastes good might be as much a ‘fact’ as the fact that the earth is round and the fact that two plus two equals four. Factual relativism makes available a level of analysis at which a wide range of what appear to be facts can be treated as facts, while at the same time providing the means to model the difference between objective and subjective facts in a way that traces the intuitive distinction.²⁴

References

- [1] DeRose, Keith (1991) ‘Epistemic Possibility’, *Philosophical Review*, 100, 581–605.
- [2] Egan, Andy, Hawthorne, John, and Weatherston, Brian (2005) ‘Epistemic Modals in Context’, in Georg Peter and Gerhard Preyer (eds.), *Contextualism in Philosophy*. Oxford: Oxford University Press.
- [3] Einheuser, Iris (2006) ‘Counterconventional Conditionals’, *Philosophical Studies*, 127(3), 459–82.
- [4] Kaplan, David (1989) ‘Demonstratives’, in Joseph Almog, John Perry, and Howard Wettstein (eds.), *Themes from Kaplan*. Oxford: Oxford University Press.
- [5] Kölbel, Max (2002) *Truth without Objectivity*. Routledge.
- [6] MacFarlane, John (2003) ‘Future Contingents and Relative Truth’, *Philosophical Quarterly*, 53.
- [7] ——— (2005a) ‘Making Sense of Relative Truth’, *Proceedings of the Aristotelian Society*, 105, 321–39.
- [8] ——— (2005b) ‘The Assessment Sensitivity of Knowledge Attributions’, in Tamar Szabó Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology* 1. Oxford: Oxford University Press.

²⁴ An earlier version of this paper was presented in September 2005 at the LOGOS-Conference *Relativizing Utterance Truth* in Barcelona. I would like to thank the organizers and the audience there for lively and helpful discussion. Richard Dietz and Max Kölbel have provided very detailed written comments which have helped putting the paper into its final form. Thanks also to two anonymous referees for Oxford University Press.

This page intentionally left blank

PART III

OBJECTIONS TO RELATIVISM

This page intentionally left blank

9

Assertion, Belief and Disagreement: A Problem for Truth-Relativism

Sebastiano Moruzzi

9.1 AN OLD DOCTRINE AND AN OLD CHALLENGE

Truth-relativism and the attempt to refute truth-relativism have a long tradition. Perhaps, the most famous argument against truth-relativism dates back to Plato's *Theaetetus*, where Plato tries to show that Protagoras' thesis is self-refuting since it cannot but be false if it is assumed to be true.¹ However, in this kind of critical discussion, truth-relativism has often been left undistinguished from subjectivism or conceptual relativism. This way of understanding truth-relativism has bound the validity of several arguments offered against it to depend on the acceptance, on the part of the truth-relativist, of doctrines which are not essentially connected to her theory, thus weakening the case against truth-relativism itself.² The discussion about truth-relativism has also suffered from two further problems.

¹ Protagoras' Thesis is the so-called Measure Doctrine:

there is nothing which is itself just one thing: nothing which you could rightly call anything or any kind of thing. If you call a thing large, it will reveal itself as small, and if you call it heavy, it is liable to appear as light, and so on with everything, because nothing is anything or any kind of thing. What is is really true, is this: the things of which we naturally say that they 'are', are in process of coming to be, as the result of movement and change and blending with one another. We are wrong when we say they 'are', since nothing ever is, but everything is coming to be. (*Theaetetus*, 152d–e)

Protagoras' Measure Doctrine is notoriously hard to interpret. It seems fair to say that it is a thesis about the status of perceptual ascriptions: nothing is cold in itself, but only relatively to someone; if we allow to use semantic ascent, one interpretation of this doctrine is that nothing is *truly* cold in itself, rather the truth of "It is cold" is always relative to someone. Burnyeat has offered an interesting reconstruction of Plato's argument. According to him the conclusion of the argument, which is divided in three steps, is that Protagoras' doctrine is false for everyone, and thus for Protagoras himself (see Burnyeat (1976b); see also Burnyeat (1976a) for an historical reconstruction of how the argument has been understood in later ancient Greek philosophy).

² Davidson's argument against relativism is based on the idea that the notion of conceptual scheme is untenable (see Davidson (1974); see also Newton-Smith (1982)). For a recent discussion of Davidson's argument see Bilgrami (2002).

The first problem is that the difficulty faced in properly stating the doctrine has generated the widespread suspicion that the doctrine is essentially unclear, if not incoherent.³ The second problem is that truth-relativism has been often understood as a thesis claiming that all sentences⁴ of a language are true relative to a parameter and false relative to another; such a strong formulation of relativism faces the self-refutation charge since it is natural to object that the truth-relativist thesis itself must then be false relative to some parameter.⁵

9.2 THE REVIVAL OF THE OLD DOCTRINE

Several proposals in the literature have recently revived truth-relativism trying to make sense of the notion of relative truth and attempting to apply the doctrine for solving different philosophical problems fruitfully.⁶

In this paper I will focus on John MacFarlane's truth-relativist proposal, since he has formulated a very detailed version of truth-relativism arguing that several areas of discourse of English such as future contingents, aesthetic judgements, judgements of taste, epistemic modals, and knowledge attributions can be fruitfully analysed with the help of a truth-relativist semantics (see MacFarlane 2003, 2005a). MacFarlane's formulations of truth-relativism have the merit of dispelling part of the objections which characterize the past literature on relativism. MacFarlane's relativism cannot be accused of unclarity, or even

³ See for example Meiland (1977), Newton-Smith (1982), Siegel's critique to it (Siegel 1986) and Swyer (1982).

⁴ Strictly speaking, because of uncontroversial phenomena of context-sensitivity due to, for example, indexical of expressions such as "I", "here" etc. . . ., we should consider *utterances* of sentences.

⁵ Of course, the attempt dates back to Plato. Recent discussions of this sort are, for example, Hales (1997) and Bennigson (1999); see also Baghramian (2004, ch. 4). For a recent critique along the traditional line see also Boghossian (2006) (especially ch. 4). In this latter work Boghossian tries to argue against the relativization of *facts*, however his argument is not forceful against the recent brand of truth-relativism that I will be considering in this paper. First, Boghossian argues against global relativism (i.e. relativism about all the domains) whereas the authors I will be considering have all argued for local relativism. Second, Boghossian seems to assume that truth-relativism is committed to the idea that an utterance of "p" expresses the claim "relatively to X, it is the case that p" (Boghossian (2006), 52), but this is rather what a contextualist would say. Notice also that his regress argument is based on this assumption plus the *further* assumption that from a global relativist point of view even the fact that "relatively to X, it is the case that p" expresses a claim that needs further relativization—i.e. that claim "relatively to X, it is the case that relatively to X, it is the case that p" which itself expresses a more complex proposition, and so on; the relativist is eventually committed to the idea that the utterance of "p" expresses an infinitary proposition "that we can neither express nor understand" (Boghossian (2006), 56). However, notice that even the standard contextualist is not committed to this further assumption: the context fixes what is said by the utterance and nothing can change this fact; nor is the truth-relativist committed to this idea (even though she could be committed to the idea that what is said is relative to the context of *assessment* of the utterance; see in this volume Cappelen 2008).

⁶ See Egan *et al.* (2004), Egan (2006), Kölbel (2002), Kölbel (2003), Lasersohn (2005), MacFarlane (2003), MacFarlane (2005b), and Richard (2004).

of using incoherent notions, since he has advanced his proposal providing a detailed compositional semantics and a definition of relative truth.⁷ According to MacFarlane, an appropriate semantics should not only use the relativization to the context in which the act of acceptance is performed (e.g. the context of use of a sentence in the case of assertion), but also to the context in which the truth-value of the sentence is assessed.⁸ The accuracy of acceptance (where acceptance can be a mental act such as a belief or a speech act such as assertion) of a proposition is relative to the assessor's context: the accuracy of my belief that cappuccino tastes better than espresso is relative to the standard of taste of whom is assessing my stance. More precisely:

(Relative Accuracy) An acceptance (rejection) of a proposition on matters of taste p relative to context of use C_U is accurate relatively to context of assessment C_A if and only if p is true relative to the circumstances $\langle W_{C_U}, S_{C_A} \rangle$ where W_{C_U} is the world of C_U and S_{C_A} is the standard of taste of the assessor at C_A .⁹

The doctrine of truth-relativism amounts to the claim that, in the language in question, there is at least one *assessment-sensitive* sentence, i.e. a sentence that, used in a certain context, expresses a proposition which can be correctly assessed in two different ways in two contexts of assessment.¹⁰

Since the accuracy of acceptance (or rejection) is relative due to the assessment-relativity of propositions accepted (or rejected) in certain contexts, the norms governing our practice of accepting (and rejecting) need a specific formulation. In fact, the traditional idea of conceiving, for example, truth as the aim of assertion must either be qualified (e.g. to assert "p" is to aim to its truth relatively to the asserter's assessment context) or given up. According to MacFarlane the relativist must choose the latter option.¹¹ Instead of conceiving acceptance as aiming at the truth, we must think of acceptance as a *commitment* to the truth of the proposition accepted. For example, consider the case of assertion: the

⁷ I will base the exposition of MacFarlane's theory on his latest statement of his theory to be found in MacFarlane (2007).

⁸ Relativism has to be distinguished from contextualism according to which the same sentence type uttered in two different contexts can have different truth-values since it comes to express two different propositions. Here I leave aside the subtle issues connected to the esoteric (but crucial) distinction between relativism and non-indexical contextualism (see MacFarlane (2007: 26) and, more extensively, MacFarlane (forthcoming)).

⁹ MacFarlane (2007: 26).

¹⁰ Since the notion of acceptance includes both the act of asserting and the act of believing, the existence of an assessment-sensitive belief is also sufficient for the correctness of truth-relativism. Whether or not there is a priority of the linguistic formulation of truth-relativism over the mental formulation or vice versa is an issue that can be left open here.

¹¹ There is no space in this paper for a discussion of reason for or against the "aiming conception" of acceptance (for MacFarlane's discussion see MacFarlane (2005b: 328–37), MacFarlane (2007: § 5). For critiques to the possibility of making sense of the aiming conception from a relativist standpoint see Evans (1979) and Percival (1994)). For the sake of the argument I will assume MacFarlane's characterization of the norms governing acceptance of a proposition.

commitment undertaken in asserting a proposition “is commitment to meeting all legitimate challenges to the accuracy of one’s assertion, and to withdrawing the assertion (disavowing the commitment) if one cannot do so” (MacFarlane (2007: 28)). The picture, borrowed by the works of Robert Brandom (Brandom 1994), is that acceptance is a move in the game of giving and asking for reasons, a game disciplined by several rules which define the content of the commitment undertaken by each of the moves. In particular, one of these rules dictates that a speaker of a language is committed to justify, if challenged, the proposition she has asserted:

(Relative Assertion 1) In asserting that p at the context C_1 , one commits oneself to withdrawing the assertion in any future context C_2 , if p is shown to be untrue relative to context of use C_1 and context of assessment C_2 .¹²

(Relative Assertion 2) In asserting that p at the context C_1 , one commits oneself to justifying the assertion when the assertion is appropriately challenged. To justify the assertion in a context C_2 is to provide grounds for the truth of p relative to context of use C_1 and context of assessment C_2 .

where C_2 is the context of assessment in which “the assertor is evaluating the putative refutation” (MacFarlane (2005b: 336)). In other words, if the assertor recognizes in a later context that what she asserted is untrue or unjustified, she is committed to retract her assertion on the pain of either being dishonest or of misunderstanding the practice of assertion.

Consider a dispute over the goodness in taste of cappuccino. It is well known that in accounting for the intuition that the dispute is faultless we are faced with a dilemma: either you go with the hard-core objectivist solution holding that one of the disputants is right and the other is wrong with the drawback that it is a mystery why both disputants think they are right and that you are committed to a dubious rampant realist metaphysics of the property of the goodness of taste of cappuccino; or you go with the contextualist holding that both disputants are right since each of them actually expresses something as “cappuccino tastes good (or bad) *for me*”, but the drawback is then that it leaves unexplained why they believe disagreeing (unless an error theory about the semantics of these expressions is postulated for the speakers of the language).¹³

MacFarlane’s solution to this puzzle is that since the notion of accuracy is relative to assessment-contexts—i.e. it is a perspectival notion—each of the disputants is right in his own acceptance and in believing that the other is wrong. Given that a necessary condition for disagreement is that not both of the parties can be accurate and that since from both of the contexts of assessment the respective opponent’s acceptance comes out as inaccurate, the perspectival

¹² MacFarlane also states another principle regarding assertion (MacFarlane (2005b: 336–37)). I will consider only Relative Assertion 1 and 2 because they are relevant for my argument.

¹³ For a statement of this dilemma see Wright (2006: 38–40) and MacFarlane (2007: 18–21).

notion of accuracy gives the right result even if the notion of disagreement is itself perspectival:

(Relative Disagreement) Two parties disagree relatively to the context of assessment C_A if (a) there is a proposition that one party accepts and the other rejects, and (b) acceptance and rejection cannot be both accurate (as assessed from C_A)

Does Relative Disagreement give also a *sufficient* condition for disagreement? MacFarlane himself has highlighted difficulties with regards to the sufficiency condition: given that there are temporal propositions, any acceptance or rejections of them would count as a disagreement (MacFarlane (2007), § 3). In Section 8.1 I will argue that the relativist proposal does not take into consideration the point of the disagreement which is to reach an agreement by ratiocinative means (e.g. by means of reasoning or by inspecting the relevant facts). Moreover in Section 8.2 I will argue that ignoring this feature commits the relativist to an ignorance theory for her proposal.

9.3 THE REVIVAL OF THE OLD CHALLENGE

Notice that this formulation of truth-relativism evades the classical charge of self-refutation since it does not imply that *every* utterance is assessment-sensitive. The truth-relativist is not thus committed to acknowledge that even the truth-relativist thesis is false in respect of some context of assessment.¹⁴

Still, it might be thought that the truth-relativist is in trouble anyway because of a different version of self-refutation. Suppose that two speakers take opposite stances on a given assessment-sensitive proposition p (one believes that it is true and the other believes that it is false). If the truth-relativist believes that her doctrine can be applied to the case in question—i.e. if she believes that the proposition in question is assessment-sensitive—how can she take a view on this proposition since she believes knowing that the proposition is true relatively to a context of assessment and false relatively to another context of assessment? How can she *rationally* privilege one of the two contexts once she believes that both can correctly determine the truth-value of the proposition in two incompatible ways? One might think that the relativist should either abstain from any judgement on the relevant proposition or she must give up the idea that the proposition is assessment-sensitive.

¹⁴ Moreover the distinction between object language and metalanguage can be used to deflect the charge of self-refutation arguing that the truth-relativist thesis is formulated in a metalanguage which is “devoid of assessment sensitivity” (MacFarlane (2005b): 338 n. 19)—i.e. a language where no proposition expressed by an utterance of its sentences has a truth-value which varies with the context of assessment.

Such a conclusion is certainly odd because it seems to imply that the attitude of belief towards a proposition is constrained by whether or not it is believed that the proposition is assessment-sensitive: either (i) it is believed that the proposition is assessment-sensitive without, at the same time, holding any opinion about the proposition in question, or (ii) it is a belief held about the proposition without accepting that the proposition is assessment-sensitive.

In the following I will develop this line of argument in detail. In Section 4 I will introduce a toy example in order to test the truth-relativist proposal in relation to disputes in matter of taste.¹⁵ In Sections 5, 6, and 7 I will develop three lines of argument which aim to show that acceptance of the truth-relativism in a context of a dispute has problematic consequences. Though I will argue that none of these arguments succeeds in showing that there is a real problem for truth-relativism, these arguments hint at a real problem for the truth-relativist. The problem will be stated in Section 8.

9.4 A SIMPLE SCENARIO

Scene 1 Ann: (after sipping cappuccino from a mug) This cappuccino tastes good! Mary: (after tasting the cappuccino from the same mug) No, it is horrible!

Scene 2 Linda, a truth-relativist about taste, joins the party and, listening to the conversation, explains the doctrine of truth-relativism to Ann and Mary. Ann and Mary start discussing about truth-relativism and they are eventually persuaded that truth-relativism is the best explanation of the intuition of faultless disagreement in matters of taste.

9.5 FIRST ATTEMPT

Consider now the scenario including Scenes 1 and 2 where Ann and Mary, after disputing about the taste of the cappuccino, are persuaded by Linda to go relativist. In such a scenario they believe that truth-relativism is the true doctrine in matters of taste, in other words they believe they grasp now the correct truth-conditions of their assertions. Let's call the proposition that cappuccino tastes good *CAPPUCCINO*. Given their new belief, they realize that the accuracy of their acceptance (rejection) of *CAPPUCCINO* is perspectival:

(Relative Ann) Ann's *acceptance* of *CAPPUCCINO* relative to her context of use (C_{Ann}) is *accurate relative to an assessor context* C_A if and only if *CAPPUCCINO* is

¹⁵ Kölbel (2002), Lasersohn (2005) and MacFarlane (2007) have argued that truth-relativism can be applied to these disputes.

true relative to the circumstances $\langle W_{C_{Ann}}, S_{C_{Ann}} \rangle$ where W_{C_U} is the world of C_{Ann} and S_{C_A} is the standard of taste of the assessor at C

(Relative Mary) Mary's *rejection* of *CAPPUCCINO* relative to her context of use (C_{Mary}) is *accurate relative to an assessor context* C_A if and only if *CAPPUCCINO* is true relative to the circumstances $\langle W_{C_{Mary}}, S_{C_A} \rangle$ where W_{C_U} is the world of C_{Mary} and S_{C_A} is the standard of taste of the assessor at C_A

Ann and Mary believe the truth-relativist doctrine, but can they coherently express that this dispute is faultless while at the same time being engaged in the dispute? Can Ann believe that the dispute between her and Mary is without fault even though she disagrees with Mary? To do this, she has to *claim* that from the context of assessment C_{Ann} , *CAPPUCCINO* is true and that from the context of assessment C_{Mary} *CAPPUCCINO* is false:

(Relative Cappuccino)

- *CAPPUCCINO*
is true relative to the circumstances $\langle W_{C_{Ann}}, S_{C_{Ann}} \rangle$ where W_{C_U} is the world of C_{Ann} and $S_{C_{Ann}}$ is the standard of taste of the assessor at C_{Ann}
- *CAPPUCCINO*
is false relative to the circumstances $\langle W_{C_{Mary}}, S_{C_{Mary}} \rangle$ where W_{C_U} is the world of C_{Mary} and $S_{C_{Mary}}$ is the standard of taste of the assessor at C_{Mary}

Relative Cappuccino is meant to express the thought that there is a genuine case of faultless disagreement: there are two assessment contexts where the same proposition is correctly evaluated with two contradictory semantic values—i.e. as true in one context and as false in the other context.

According to Relative Assertion 1, an assertion of Relative Cappuccino involves a commitment to withdraw the assertion in any future context, if in this context she recognizes that what she asserted is shown untrue. Can Ann and Mary keep the commitment undertaken by asserting and denying in their respective contexts *CAPPUCCINO* after asserting Relative Cappuccino?

The question is whether, for example, Ann is bound to reject *CAPPUCCINO* after having asserted Relative Cappuccino. After all, it might be suggested, since Ann accepts Relative Cappuccino she thereby *recognizes* that Mary has accurately rejected a false proposition from her point of view. And if this is the case, why can't this latter fact count as a piece of conclusive evidence against the commitment undertaken in her previous assertion of *CAPPUCCINO*?

The careful reader has probably already noted that this reasoning contains two errors: (1) remember that Ann believes that accuracy is perspectival, so she can recognize that Mary has accurately rejected *CAPPUCCINO* only if, for example, she would be capable of deferring to Mary's standard and so far

nothing has been said about the possibility of this mechanism of deference; (2) even granting that there is such a mechanism of deference, Relative Assertion 1 dictates to Ann rejecting her assertion of *CAPPUCCINO* only if *CAPPUCCINO* is untrue *relatively* to her standard of taste (selected by her context of assessment), but Ann's actual standard of taste is different from her *deferred* standard of taste.

In this sense, the charge of pragmatic self-refutation is formally avoided since the belief in the truth-relativist doctrine does not automatically cut the grounds to belief in *CAPPUCCINO*.

9.6 SECOND ATTEMPT

It might be thought that it is incoherent to believe that *CAPPUCCINO* is assessment-sensitive when we consider a neutral perspective on the matter. Let's suppose that Linda's context of assessment (C_{Linda}) is neutral on the dispute between Ann and Mary.¹⁶

Now, consider Scene 1 only where Ann and Mary still do not know anything about truth-relativism, and let's suppose that Linda is secretly listening to their discussion. If the truth-relativist believes that there is a genuine faultless disagreement between Mary and Ann, she would presumably hold that both are correct and that their assertions are perfectly justified from their respective contexts of assessment. Since Ann and Mary have, from Linda's perspective, a true justified belief it seems natural to say that Linda believes that Ann and Mary are knowledgeable with respect to *CAPPUCCINO*.¹⁷ Now, since the truth-relativist's perspective (i.e. Linda's perspective) is neutral with respect to *CAPPUCCINO*, Linda, the truth-relativist, should be accurate, relative to her context of assessment, in accepting the following propositions:

- (1) that Ann knows that *CAPPUCCINO*
- (2) that Mary knows that *NOT-CAPPUCCINO*¹⁸

Presumably, the truth-relativist version of factivity would be expressed saying that if a knowledge ascription is true relative to a context of assessment, the embedded sentence is true relative to the same context of assessment. Given the assumption that Linda's acceptance of previous knowledge ascriptions is accurate,

¹⁶ The following argument is an adaptation of an argument advanced by Jason Stanley in relation to the truth-relativist analysis of knowledge-ascriptions (see Stanley (2005), 146). Stanley believes that the argument, together with other two arguments, shows that the truth-relativist cannot give any coherent formulation of factivity. My use of the argument is different, since I start from the assumption that there is a natural formulation of factivity for the truth-relativist, but that this formulation of factivity jeopardizes the very idea of a neutral perspective.

¹⁷ Putting aside Gettier-style worries.

¹⁸ Where, of course, *NOT-CAPPUCCINO* is the negation of *CAPPUCCINO*.

it follows that she is committed to inconsistent propositions.^{19,20} However, it is unclear why the truth-relativist should be committed to the existence of a neutral perspective. On the contrary, assuming bivalence, there are reasons to think the contrary. What does it mean to be neutral about a proposition for a context of assessment? On an epistemic reading of neutrality, to be neutral with respect to a proposition means to be agnostic about the proposition. But since we are dealing here with truth-relativism, the idea is rather that a certain perspective does not commit to the truth or falsity of a certain proposition. However, while the absence of commitment in the epistemic sense amounts to the suspension of judgement, in the alethic case the analogous situation would be that the proposition is neither true nor false. But in this case it is the burden of the objector to justify the claim that for every assessment-sensitive proposition there is a context of assessment in relation to which the proposition is neither true nor false. Why should bivalence fail just because there must be such a neutral perspective?²¹ The truth-relativist might also add that she could make sense of the idea of a neutral perspective by means of a metalanguage which is devoid of assessment sensitivity.

¹⁹ To express more precisely the worry, the proposal amounts to say that acceptance of the relevant knowledge ascriptions is accurate relatively to an assessment context C_A just in case:

- (3) relatively to the context of assessment C_{Linda} , the proposition that Ann knows that *CAPPUCINO* is true relative to the circumstances $\langle W_{C_{Ann}}, S_{C_{Ann}} \rangle$ where W_{C_U} is the world of C_{Ann} and $S_{C_{Ann}}$ is the standard of taste of the assessor at C_{Ann}
- (4) relatively to the context of assessment C_{Linda} , the proposition that Mary knows that *NOT-CAPPUCINO* is true to the circumstances $\langle W_{C_{Mary}}, S_{C_{Mary}} \rangle$ where W_{C_U} is the world of C_{Mary} and $S_{C_{Mary}}$ is the standard of taste of the assessor at C_{Mary}

—in other words C_{Linda} features as an idle wheel. The aforementioned truth-conditions amount to the claim that knowledge ascriptions do not inherit assessment-sensitivity from the embedded propositions.

Given now the following rule of factivity of knowledge:

(Relative Factivity)

- from** relatively to the context of assessment C_A , the proposition that x knows *CAPPUCINO* is true to the circumstances $\langle W_{C_x}, S_{C_x} \rangle$ where W_{C_U} is the world of C_x and S_{C_x} is the standard of taste of the assessor at C_x
- infer** relatively to the context of assessment C_A , *CAPPUCINO* is true to the circumstances $\langle W_{C_x}, S_{C_A} \rangle$ where W_{C_U} is the world of C_x and S_{C_A} is the standard of taste of the assessor at C_A

Linda's context of assessment is bound to evaluate *CAPPUCINO* as both true and false.

²⁰ The previous footnote is symptomatic of the fact that a proposition p is assessment-sensitive if and only if the proposition that is known that p is assessment-sensitive. Suppose that p is true at C_1 and that not- p is true at C_2 , then $\neg K\neg p$ is true at C_1 and $\neg Kp$ is true at C_2 . If we assume that in both contexts there is a subject who has knowledge of p , then Kp is assessment-sensitive. The other direction is straightforward: suppose that Kp is assessment-sensitive, then two subjects know, respectively, that p and that $\neg p$. By Relative Factivity p is assessment-sensitive (I am indebted for this point to Dan López De Sa).

²¹ Assuming here that we are dealing only with propositions for which the truth-relativist analysis clearly does not require the failure of bivalence.

Finally, the truth-relativist might opt to state the rule of factivity by selecting the standard of assessment of the ascriber of the knowledge claim, hence avoiding that Linda contradicts herself.²²

9.7 THIRD ATTEMPT

It might be argued that *CAPPUCCINO* cannot be believed because from the point of view of the truth-relativist semantics it is an incomplete entity which cannot be the object of any belief.²³

A paradigmatic example of an incomplete semantic entity which cannot fulfil the role of a proposition is the propositional function expressed by “I am happy” when this sentence is just written on a blackboard in a class of English language. The sentence expresses something, but it is something we cannot accept or reject.²⁴

Consider now the following argument in relation to a scenario which includes Scenes 1 and 2.

- 1 (1) According to truth-relativism about matters of taste, it is the case that *CAPPUCCINO* is true relative to the circumstances $\langle W_{C_{Ann}}, S_{C_{Ann}} \rangle$ where $W_{C_{Ann}}$ is the world of C_{Ann} and $S_{C_{Ann}}$ is the standard of taste of the assessor at C_{Ann} ; (Ass.)
- (2) Define *CAPPUCCINO*+ as the proposition featuring in the that-clause of (1) (Def.)
- 1 (3) *CAPPUCCINO*+ has absolute truth-conditions (i.e. it is not assessment-sensitive): *CAPPUCCINO*+ is true if and only if *CAPPUCCINO* is true relative to the circumstances $\langle W_{C_{Ann}}, S_{C_{Ann}} \rangle$ where $W_{C_{Ann}}$ is the world of C_{Ann} and $S_{C_{Ann}}$ is the standard of taste of the assessor at C_{Ann} ; (1,2)
- 1 (4) *CAPPUCCINO*+ is a different entity from *CAPPUCCINO*; (1,3)

Given scenes 1 and 2, Ann believes *CAPPUCCINO* because of her personal taste and furthermore she believes *CAPPUCCINO*+ because she has been persuaded by Linda to believe truth-relativism, hence:

²² However, this new rule of factivity would fail to sanction, by conditional introduction, the truth relative to *any* standard of assessment $Kp \supset p$ which is plausibly an a priori truth.

²³ The following is an adaptation of Zimmerman’s argument (Zimmerman 2006), Egan raises a similar problem in relation to assertion (Egan 2006).

²⁴ Zimmerman says that such an incomplete semantic entity is something that we cannot “weigh the reasons for and against thinking that it is true” (Zimmerman 2006).

²⁵ This fact should not be given for granted. There are areas of discourse where it is arguable, once truth-relativism is assumed to be the correct semantic theory, that truth-relativism holds also for higher-orders (see in this volume Wright (2008: § 4) and Moruzzi and Wright (forthcoming)).

- 5 (5) Ann believes *CAPPUCCINO* and *CAPPUCCINO*+; (Ass.)
- 6 (6) If a proposition is assessment-sensitive, it cannot be con- (Ass.)
sidered independently from a context of assessment (i.e. it
cannot be entertained as having absolute truth-conditions;
analogy with 'I am happy');
- 1,6 (7) Ann cannot believe *CAPPUCCINO*, she can only believe (1,6)
CAPPUCCINO+;
- 1,5,6 (8) Contradiction. (5,7)

Thus the conclusion is that *CAPPUCCINO* fails to have the cognitive role that content should have, i.e. being the object of belief. If the above argument were sound, it would establish the disastrous result that semantics postulates contents which cannot be believed.

Of course the crucial step of the argument is line 6, but why should we assume that the truth-relativity of a proposition must compromise the cognitive role of being object of belief? The obvious problem with this argument is that it assumes, without further argument, that the truth-relativity of a semantic entity forbids that this entity can be the object of a belief *with the notable exception of the relativity to possible worlds*: only indexical propositions (characters), tensed propositions, and relativist propositions are propositional functions, and thus, being incomplete, they cannot be believed.²⁶

In the following section I will offer a better argument for the idea that an assessment-sensitive proposition cannot be the object of belief.

9.8 FOURTH ATTEMPT

Remember that the relativization of the accuracy of acceptance of a proposition (e.g. of the correctness of the assertion) does not leave room for the idea that the accuracy of the acceptance and rejection of *CAPPUCCINO* is absolute.²⁷

In the following I will argue that the perspectival notion of accuracy is unsatisfactory because: (1) it fails to explain why, in cases such as Ann's and Mary's dispute over *CAPPUCCINO*, the parties involved behave as if they really disagree over the truth of the proposition in question—i.e. as if they were genuinely engaged in a dispute; (2) if Mary and Ann believe the truth-relativist story, they are rationally compelled to abstain from disputing over *CAPPUCCINO*.

²⁶ The style of the argument reminds of Mark Richard's famous argument against temporalism (see Richard 1981).

²⁷ A different notion of accuracy where the context of assessment does not play a role in evaluating the accuracy of a certain acceptance could allow this (see the "use-centric notion of accuracy" in MacFarlane (2007: § 4.1)).

To control the factors involved in the following arguments, I will argue in two phases: first I will consider only Scene 1 of Simple Scenario where it is not assumed that the speakers know the truth-relativist semantics, then I will extend the scenario to Scene 2 where this assumption is in play.

To be clear about the example, let us put aside any possible difference between Ann and Mary in their expertise on the art of cappuccino-making,²⁸ and let us suppose also that the application of their standards does not involve, whatever that might mean, an error—i.e. no cognitive shortcoming is involved.

9.8.1 Argument Limited to Scene 1: A Problem of Explanation

Let us ask how Ann can recognize Mary's ground for the rejection of *CAPPUCCINO*. From Ann's perspective, *CAPPUCCINO* is true and Mary's rejection of *CAPPUCCINO* is not accurate.²⁹ Now, the problem is not that for this latter reason Mary is at fault. After all, if accuracy is perspectival so is the notion of fault. Alternatively, we could also decide to use "fault" just in case the proposition asserted is false relative to the standard of the assertor; in this sense Ann could assert that Mary is not at fault. But this is just a quibble. Rather, the problem arises in connection with *the point of the disagreement* between Mary and Ann.

According to Relative Assertion 2 (see above § 2) acceptance of a proposition involves a commitment to provide a defence if it is challenged from some context of assessment. To what extent can we *explain* the dispute over *CAPPUCCINO* on this base? If Mary's challenge to Ann cannot be met in principle unless Ann changes her standard of taste, how can this explain the behaviour of Ann and Mary?

Compare the truth-relativist way of representing the accuracy of Mary's rejection with a situation where Mary has rejected *CAPPUCCINO* because of a piece of evidence, unknown to everyone except to whom has decided to join her in the rejection.³⁰ In such a situation to what extent could Mary challenge Ann's belief about *CAPPUCCINO* if the only way to show the reasons for her disbelief would be to change her mind? Of course, Mary's challenge might be based on issues accessible also to Ann such as, for example, the fact that the cappuccino is bitter; but while for Mary this would count as a reason for disliking cappuccino, for Ann it would count as the opposite; after such a move they would find themselves in the same situation as before: Mary would say "But don't you feel how it is bitter?!" and Ann would reply "Exactly!" and here the game of giving

²⁸ See *contra* in this volume Iacona (2008) for the idea that there is no such thing as a dispute involving faultless disagreement when the disputants have the same level of expertise.

²⁹ Given that Ann and Mary's world is the same one and that no further parameter relative to circumstances of use is relevant for the relativist, I omit any reference to circumstances of evaluation in the following.

³⁰ The analogy is derived from a similar example in Rosenkranz (2008).

and criticizing reasons would have no point, there would simply be nothing more to do.³¹

9.8.1.1 *First Counter-objection: stalemate is the appropriate prediction*

It could be replied that nothing of what I have said is problematic, in fact this is exactly the kind of stalemate in which disputes of taste typically end. Hence, contrary to my allegations, the relativist predicts exactly what has to be predicted.

Reply: what is problematic is to explain these kind of conversations by means of the idea of challenges and defences: how can the appeal to challenges based on standards which are different from the standards of who is challenged describe such behaviours? What is the point of such challenges if in principle they cannot be met? Wouldn't it be more appropriate to describe these situations just as expressions of our own personal preferences instead of commitments which are in principle right and challenges which are unanswerable?

9.8.1.2 *Second Counter-objection: controversy just feels uncomfortable*

MacFarlane's reply to the latter questions is that our stubborn attitude to challenge in these extreme circumstances is useful also "to foster coordination of contexts [because we] have an interest in sharing standards of taste [...] with those around us" (MacFarlane (2007), 30). But why is it so? Because, MacFarlane suggests, controversy just feels uncomfortable in itself, and this is just a *datum* of human nature.³²

Reply: this move is explanatorily unsatisfactory: the perspectival notion of accuracy (Relative Accuracy) and the idea of characterizing assertion on the basis of a commitment governed by certain rules (Relative Assertion 1 and 2) were meant to provide an explanation of the puzzling phenomenon of faultless disagreement. The advertised advantage of the relativist proposal was that, whereas the hard-core objectivist and the contextualist are at odds in making sense, respectively, of faultlessness and disagreement, the relativist could explain both features. But it turns out now that the relativist package explains this phenomenon only as long as we assume as a brute fact that the behaviour of human agents exhibited in dispute of matters of taste is recalcitrant to any explanation on the basis of the game of asking and giving reasons—i.e. recalcitrant to the relativist's *explanans*!

9.8.1.3 *Third Counter-objection: a challenge is an invitation to experiment*

A way to make sense of the idea of challenges and defences in a truth-relativist scenario could be of conceiving a challenge to the assertion of "Cappuccino does not taste good" as an invitation to taste cappuccino in order to *get used* to its

³¹ A very similar point is made in Moltmann (forthcoming), § 2.2.1.

³² MacFarlane refers here to Gibbard (1990) according to whom there is an evolutionary explanation of this fact.

peculiar flavour in the long run; getting used to cappuccino would thus put you in a position to appreciate it properly.³³

Reply: this interpretation would leave anyway unintelligible a dispute between people who have been tasting cappuccino for a relevantly long period who nonetheless continue to disagree.

9.8.1.4 Fourth Counter-objection: disagreement is independent from its manifestation

It might be protested that the above line of reasoning ties the conditions for the existence of a disagreement with conditions relative to the manifestability of the disagreement: whether or not the point of a dispute is to reach agreement by ratiocinative means (e.g. by means of reasoning or by inspecting the relevant facts), a disagreement requires just the expression of opposite attitudes (i.e. acceptance and rejection) towards a certain proposition. There can be irresolvable disagreements. After all, *pace* verificationism, we can certainly dispute over unknowable propositions and hence we have cases of irresolvable disputes already. The fact that the conflicting attitudes of Ann and Mary are expression of an irresolvable conflict does not undermine the fact that they are disagreeing; if Ann or Mary change their standards, the dispute can be felicitously solved and there can be convergence of opinions; if they stick to their different standards, they attain to nothing. If the latter happens, the insolubility of the dispute sheds light on the nature of the attitudes involved: they are not beliefs which we are disposed to abandon by ratiocinative means, they are rather tokens of a different type of attitude, e.g. personal preferences, a form of acceptance which is different from, for example, beliefs about the physical properties of the world.³⁴

Reply: first, it is suspect to proliferate types of attitudes; how many types of acceptance are there? Can *CAPPUCCINO* be accepted in different ways in a dispute? If not, how can we describe the case of someone who, after disputing, decides to change her mind? Would she be in error—would she have misunderstood the nature of her own attitude? Or, if she were not in error, what would be the role of the ratiocinative means she has putatively employed in the change of her mind given that this attitude is by definition irresolvable by these means? If, on the other hand, *CAPPUCCINO* can be accepted in different ways in a dispute, what are the roles of these different attitudes in the dynamic of the dispute? Consider the case of someone believing in creationism for purely religious reasons who is involved in a debate with a biologist who believes, because of a complex scientific enquiry, in evolutionism; there is a sense in which they are talking past each other since the reasons presented by one are, in principle, completely irrelevant

³³ I owe this objection to Marco Santambrogio.

³⁴ I have elaborated this objection on the basis of remarks made by Annalisa Coliva. The underlying idea is that acceptance of a proposition can come in different sorts as Wittgenstein already hinted in *On Certainty*.

to the other. If a dispute over *CAPPUCCINO* can be of this latter sort, in what sense are they disputing given that they appeal to different epistemic norms? Can the very same person change her way of acceptance (e.g. from dogmatist mode of acceptance to non-dogmatist mode) and why and how can it happen? Until these questions are answered the proposal looks rather mysterious.

Second, the case of unknowable propositions is not relevant since what is at stake is not the possibility for the dispute to be solved, rather it is the normative requirement that the disputants should aim to cooperate to solve the problem they are disputing on.³⁵

9.8.1.5 Conclusion of the argument limited to Scene 1

In conclusion, it remains a mystery why Mary challenges Ann since any challenge cannot but be won by the challenger and, symmetrically, it is a mystery why Ann asserts *CAPPUCCINO* since her acceptance is systematically defeated by Mary's challenge. MacFarlane's reply is that this is so because it is a brute fact of human psychology, but I find this reply deeply unsatisfactory.

9.8.2 Argument Extended to Scene 2: Suspension of Judgement

Let us put this explanatory problem aside. After all, our doubts could be assuaged by the well-known (and philosophically frustrating) thought that we can't ask for reasons *ad infinitum*. At a certain point we reach the bedrock. Ann and Mary have simply reached, so to say, their personal bedrock for what counts as a good cappuccino.

Be that as it may, let us move to Scene 2 where Ann and Mary come to believe the truth-relativist story about *CAPPUCCINO*. I will argue that it is then impossible for the truth-relativist to make sense of the *rationality* of Ann's and Mary's respective opinions.

Ann and Mary know that acceptance of a proposition involves a commitment to provide a defence if challenged from some context of assessment. To what extent would they dispute about *CAPPUCCINO* once informed of this fact?

Mary *knows* that her challenge cannot be met in principle unless Ann changes her standard of taste, and since Ann knows this as well, there is no point in arguing about *CAPPUCCINO*.

9.8.2.1 First Counter-Objection: silence is appropriate

Again it could be replied that this kind of stalemate is the usual situation in the disputes of taste.

Reply: notice that once we assume knowledge of the truth-relativist doctrine, silence would *always* be the rational way to act. In fact, if assertion is understood

³⁵ For a similar point see Price (2003).

in the context of such a dispute as an action undertaking a commitment to face challenges from other perspectives, no challenge can make sense if grounded on a different assessment context, thus the only rational thing to do is to remain silent and not even start disputing. Once Ann and Mary know that there is no point in issuing challenges and replying to them in any other way than persuading the other to stop using her own standard, they would simply cease discussing.

9.8.2.2 Second Counter-Objection: truth-relativism does not require rationality

It might be objected that it is not part of the truth-relativist thesis that speakers have to be rational.³⁶

Reply: it is a startling and embarrassing consequence that the truth-relativist analysis of faultless disagreement is at odds with the idea that the disputing behaviour of speakers aware of the semantics of their language is rational.

Moreover notice that if belief, which is a type of acceptance of a proposition, amounts to being prepared to issue and face such challenges, then the fact that Mary and Ann see no point in issuing challenges and replying to them amounts to the fact that they simply would stop believing anything about *CAPPUCCINO*.

9.8.2.3 Conclusion of the argument extended to Scene 2

In conclusion, Ann and Mary should simply remain silent and unopinionated about *CAPPUCCINO*. Hence a dispute can take place only insofar as the truth-relativist story is ignored by the speakers. The truth-relativist is committed to an ignorance theory if she does not want to rule out a priori the rationality of disputes of taste. Thus the truth-relativist would fail to part company with contextualist approaches which analyse “It tastes good” as an elliptical expression meaning “It tastes good to me”, since also these latter theories are committed to an ignorance theory if they want to save the rationality of a dispute.

9.9 CONCLUSION

I have considered four arguments against the idea that assessment-sensitive propositions can represent the right entities that we believe in cases such as controversies in matters of taste. Though the first three arguments have been found to be faulty, they have collectively pointed to an idea developed in the fourth argument. The conclusion of this latter argument is that the truth-relativist story fails to *explain* why we behave as we actually behave when we are involved in disputes of matters of taste, and that this story entails that we *should* not

³⁶ I owe this objection to Richard Dietz.

behave as we actually behave if we were aware of the correct semantics of our language.

Let me close with two final notes. I have chosen the case of disputes of taste since it is a standard case for accounting for the puzzle of faultless disagreement, but I do not think that the points I have made are limited to this area of discourse. Secondly, the arguments presented in the preceding sections are specifically designed against MacFarlane's version of truth-relativism, as it is presently the most developed proposal available. My claim that truth-relativism actually risks losing its advantages over the other positions (mainly contextualism) in respecting the putative idea that there is a real disagreement between opposite parties in matters of taste is thus limited to MacFarlane's theory. However, it may still be the case that other versions of relativism—yet to be developed—would be better off. These other versions could for example develop different notions of acceptance and interplay between the normative rules that govern them. However, the details of these possible developments and changes will have to be investigated and assessed on another occasion.³⁷

References

- Baghramian, M. (2004). *Relativism*, London and New York: Routledge.
- Bennigson, T. (1999). 'Is Relativism Self Refuting?', *Philosophical Studies* 94(3): 211–36.
- Bilgrami, A. (2002). 'Realism and Relativism', *Philosophical Issues* 12: 1–25.
- Boghossian, P. (2006). *Fear of Knowledge*, Oxford: Oxford University Press.
- Brandom, R. (1994). *Making it Explicit*, Cambridge, Mass.: Harvard University Press.
- Burnyeat, M. (1976a). 'Protagoras and Self-Refutation in Plato's *Theaetetus*', *Philosophical Review* 85: 172–95.
- (1976b). 'Protagoras and Self-Refutation in Later Greek Philosophy', *Philosophical Review* 85: 44–69.
- Cappelen, H. (2007). 'Content Relativism and Semantic Blindness', this volume.
- Davidson, D. (1974). 'On the Very Idea of Conceptual Scheme', *Proceedings and Addresses of the American Philosophical Association*. repr. in *Inquiries into Truth and Interpretation*, Oxford: Oxford University Press (1984), 183–98.
- Egan, A. (2006). 'Epistemic Modals, Relativism, and Assertion', *Philosophical Studies* 133(1): 1–22.
- Egan, A., Hawthorne, J. and Weatherson, B. (2004). 'Epistemic Modals in Context', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*, Oxford: Oxford University Press.

³⁷ This paper has hugely benefited from the presentations made in the Arché AHRC Centre for the Philosophy of Logic, Language, Mathematics and Mind in St Andrews and the research seminar on analytic philosophy at the University of Bologna. Special thanks to Eline Busck, Manuel García-Carpintero, Annalisa Coliva, Richard Dietz, Paolo Leonardi, Dan López de Sa, Julien Murzi, Eva Picardi, Marco Santambrogio, and Crispin Wright.

- Evans, G. (1979). 'Does Tense Logic Rest on a Mistake?' in *Collected Papers*, Oxford: Oxford University Press, 343–63.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*, Cambridge, Mass.: Harvard University Press.
- Hales, S. (1997). 'A Consistent Relativism', *Mind* 421: 33–52.
- Iacona, A. (2008). 'Faultless or Disagreement', this volume.
- Kölbel, M. (2002). *Truth Without Objectivity*, London: Routledge.
- (2003). 'Faultless Disagreement', *Proceedings of the Aristotelian Society* 54: 53–73.
- Lasersohn, P. (2005). 'Context Dependence, Disagreement, and Predicates of Personal Taste', *Linguistics and Philosophy* 28(6): 643–86.
- MacFarlane, J. (2003). 'Future Contingents and Relative Truth', *The Philosophical Quarterly* 53: 321–36.
- (2005a). 'The Assessment Sensitivity of Knowledge Attributions', in T. Szabó Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, vol. i, Oxford: Oxford University Press, 197–223.
- (2005b). 'Making Sense of Relative Truth', *Proceedings of the Aristotelian Society* 105: 321–39.
- (2007). 'Relativism and Disagreement', *Philosophical Studies* 132: 17–31.
- (forthcoming). 'Semantic Minimalism and Nonindexical Contextualism', in G. Preyer and G. Peter (eds.), *Context-Sensitivity and Semantic Minimalism: Essays on Semantics and Pragmatics*, Oxford: Oxford University Press.
- Meiland, J. W. (1977). 'Concepts of Relative Truth', *The Monist* 60: 568–82.
- Moltmann, F. (forthcoming). 'First-Person-Oriented Genericity and Relative Truth', *Mind*.
- Moruzzi, S. and Wright, C. (forthcoming). 'Trumping Assessments and the Aristotelian Future', *Synthese*.
- Newton-Smith, W. (1982). 'Relativism and the Possibility of Interpretation', in M. Hollis and S. Lukes (eds.), *Rationality and Relativism*, Oxford: Oxford University Press, 106–22.
- Percival, P. (1994). 'Absolute Truth', *Proceedings of the Aristotelian Society*, 94: 189–213.
- Price, H. (2003). 'Truth as Convenient Friction', *Journal of Philosophy* 100: 167–90.
- Richard, M. (1981). 'Temporalism and Eternalism', *Philosophical Studies* 39(1): 1–13.
- (2004). Contextualism and Relativism, *Philosophical Studies* 119: 215–42.
- Rosenkranz, S. (2008). 'Frege, Relativism and Faultless Disagreement', this volume.
- Siegel, H. (1986). 'Relativism, Truth, and Incoherence', *Synthese* 68: 225–59.
- Stanley, J. (2005). *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Swyer, C. (1982). 'True for', in J. Meiland and M. Krausz (eds.), *Relativism: Cognitive and Moral*, Paris: University of Notre Dame Press.
- Wright, C. (2006). 'Intuitionism, Realism, Relativism, and Rhubarb', in P. Greenough and M. Lynch (eds.), *Truth and Realism*, Oxford: Oxford University Press, 38–70.
- (2008). 'Relativism about Truth Itself: Further Thoughts about the Very Idea', this volume.
- Zimmerman, A. (2006). 'Against Relativism', *Philosophical Studies* 133(3): 313–48.

Frege, Relativism and Faultless Disagreement

Sven Rosenkranz

10.1 FREGE ON TRUTH AND ASSERTION

As Frege famously argued, applying ‘true’ to a given thought or proposition is neither necessary nor sufficient for asserting it. It is not sufficient because we can express the thought that it is true that P without yet asserting P , for instance, when we assume its truth in a *reductio* proof or conditionalize on its truth in asserting a complex thought of which P is a constituent (Frege (1892), 49; (1897), 54; (1906), 75; (1918/19), 35–6). Applying ‘true’ to a thought or proposition is not necessary for asserting it, because in asserting P we already present or acknowledge P as true whether or not we use the term ‘true’ in order to do so (Frege (1892), 49; (1897), 39; (1918/19), 35–6). In the phrase ‘ A presents P as true’, the term ‘true’ is accordingly part of the description of A ’s assertoric performance, and not of the expression of the proposition which is its content.

According to Frege, the truth-value of a proposition P —its being true or its being false—never forms part of a proposition, be it P itself or the proposition that P is true. Otherwise we would be able to assure ourselves of a proposition’s truth-value simply by grasping that proposition or one of which it is a constituent; and this we cannot do (Frege (1892), 49). To assure ourselves of the truth of a proposition P we must ‘proceed from the proposition to its truth-value’; and this is what we do when we correctly judge or assert P (Frege (1892), 50; (1918/19), 35). Since presenting x as F will be correct just in case x is F , it follows that to present P as true is to present P as having the truth-value *True*. We can now straightforwardly derive from the truth-value of a proposition whether an assertion which has that proposition as its content is correct. In the same vein, we can derive from the correctness of the joint assertion of P and $P \rightarrow Q$ that an assertion of $\sim Q$ will be incorrect: since the joint assertion of P and $P \rightarrow Q$ is correct only if both P and $P \rightarrow Q$ have the truth-value *True*, and the latter requires that Q also have the truth-value *True*, and since Q and $\sim Q$ cannot

both have that truth-value, anyone who asserts $\sim Q$ will present $\sim Q$ as having a truth-value that it does not have. In this way, the laws of truth can be used to derive norms of judgement and assertion (cf. Frege (1897), 38–9, 64, 69; (1918/19), 30).

This part of Fregean orthodoxy has proved its longevity in that, to the present day, it is perceived to put considerable pressure on what in metaethics is known as ‘expressivism’. Expressivists deny that in asserting that murder is wrong, one presents the proposition that murder is wrong as true. For, according to expressivism, *that murder is wrong* is not truth-evaluable at all and so no proposition in Frege’s sense. Instead, expressivists hold that to assert that murder is wrong is to express an attitude, viz. one’s disapproval of murder (see Ayer (1946), 107). However, since in asserting $P \rightarrow Q$ we neither assert P nor conditionalize on its assertion, it now becomes difficult to *derive* from the correctness of the joint assertion of ‘Murder is wrong’ and ‘If murder is wrong, then promoting murder is wrong’, that asserting ‘Promoting murder is not wrong’ would be incorrect (cf. Geach (1960); (1965)). And yet, it is evident that if both my assertion of ‘Murder is wrong’ and my assertion of ‘If murder is wrong, then promoting murder is wrong’ are correct, it does not need any further consideration to conclude that were I to assert ‘Promoting murder is not wrong’ my assertion would be incorrect (Wright (2003), 187).

10.2 FAULTLESS DISAGREEMENT

In what follows, I argue that the same set of Fregean claims puts pressure on the view, recently advanced by Max Kölbel and others, that relativism about truth can make sense of cases of *faultless disagreement* (Kölbel (2003), Lasersohn (2005)). Given Kölbel’s definition, A and B faultlessly disagree just in case (i) A believes P and B believes $\sim P$, and yet (ii) neither A ’s belief nor B ’s belief is incorrect (Kölbel (2003), 53–4).¹ Here, (i) is meant to imply that A and B *disagree*, while (ii) is meant to imply that their disagreement is *faultless*. In what follows, I will continue to speak of assertions (as expressions of beliefs) where Kölbel speaks of beliefs instead. However, nothing will hinge on that.²

Cases of faultless disagreement seem to abound in areas of non-objective discourse that relate to matters of taste or aesthetic value (Wright (1992), 145–8; (2003), 445–7; Kölbel (2002), 99–100; (2003), 53–4). Thus, if you assert that Vincent Price was a horrid actor and I assert that he was not, then we seem to

¹ Kölbel’s own version of (ii) reads ‘neither A nor B has made a mistake (is at fault)’ (Kölbel (2003), 54). I here presume that if either A ’s or B ’s belief is incorrect, then either A or B has made a mistake or is at fault.

² In other words, nothing I say will hinge on the fact that assertions as public performances are subject to norms that go beyond the norms governing belief and judgement.

disagree, and yet it seems that neither of us needs to be at fault. Similarly, if I assert that crickets with chili and lemon are tasty and you assert that they are not, we seem to disagree, and yet it seems that neither assertion needs to be incorrect.

The trouble with taking these appearances at face value is that the propositions we respectively assert cannot both be true. Accordingly, the conclusion would appear to be inevitable that one of us has presented as true what is in fact false and so has performed an incorrect assertion (Wright (1992), 148–9; Kölbel (2003), 55–6; cf. also Shapiro and Taschek (1996), Wright (2003), 457–63):

1	(1)	<i>A</i> asserts <i>P</i>	ass.
2	(2)	<i>B</i> asserts $\sim P$	ass.
3	(3)	Both assertions are correct	ass.
4	(4)	<i>P</i>	ass.
2,4	(5)	<i>B</i> 's assertion is incorrect	2,4
2,3	(6)	$\sim P$	3,5, RAA
1,2,3	(7)	<i>A</i> 's assertion is incorrect	1,6
1,2	(8)	Not both assertions are correct	3,7, RAA.

10.3 RELATIVISM ABOUT TRUTH

In order to defuse this objection, and thereby to salvage belief in faultless disagreements from easy refutation, Kölbel suggests conceiving of truth as being relative (Kölbel (2002), 100; (2003), 69–72). The relativity in question is not the familiar one that explains why your assertion of 'I am tired' may be true while my assertion of 'I am tired' is not. Nor is it the world-relativity of truth familiar from possible world semantics that Kölbel has in mind. Rather, he suggests conceiving of truth as being relative even after we have fastened upon the relevant context of use and fixed the relevant possible world. It is the relativity of *propositional truth* (at a world) that Kölbel is after—just the kind of truth that Frege insisted is absolute (Kölbel (2003), 71–2; (2004), 305–7; cf. MacFarlane (2003), 328; Frege (1897), 43–6).

According to Kölbel's version of relativism, the truth of a proposition is relative to what he calls a 'perspective', where perspectives are functions from propositions to truth-values (Kölbel (2002), 100; (2003), 70). In the case of propositions about aesthetic value, such a perspective may be conceived of in terms of a subject's accepted aesthetic standards (say, for what counts as good acting in horrid films). In the case of propositions about matters of culinary taste, a perspective may be conceived of as being determined by the peculiarities of a subject's sense of taste (say, having increased acid flow when chewing on an item of spiky texture) (cf. Kölbel (2004), 306). Whatever the exact details, perspectives are meant to define points of assessment with respect to which a proposition is to

be evaluated as true or otherwise: which truth-value is to be assigned to a given proposition depends on which perspective is in question (Kölbel (2002), 100; (2003), 70; cf. MacFarlane (2003), 328–32; (2004), 322–8).

While the absolute truth of objective propositions—e.g. propositions of basic arithmetic—can unproblematically be recast in terms of truth relative to all perspectives (Kölbel (2002), 102; (2003), 70–1), the perspectives of subjects are likely to diverge when it comes to non-objective propositions about matters of taste or value. It is here, if anywhere, that relativism about truth will show its real potential.

The lead idea then is that, in the envisaged case, P may be true relative to A 's perspective while $\sim P$ is true relative to B 's perspective, even though $P \wedge \sim P$ is not true relative to any perspective. Once truth is relativized to perspectives in this way, one cannot simply infer from the fact that P and $\sim P$ cannot both be true, that A and B cannot both assert something true. So, it would appear that we can make sense of cases of faultless disagreement after all.

10.4 RELATIVE TRUTH AND THE CORRECTNESS OF ASSERTIONS

I shall now argue that this appearance is deceptive. Recall that Frege could derive the correctness of an assertion of P from P 's having the value *True* only because he identified asserting P with presenting P as having that value. In order to show that there are cases of faultless disagreement, Kölbel's relativist has to make good the claim that both A 's assertion and B 's assertion are correct. If one thinks that, to this end, it is sufficient to observe that P is true relative to A 's perspective, while $\sim P$ is true relative to B 's perspective, then one must also think that, in asserting P , A presents P as being true relative to A 's perspective, while in asserting $\sim P$, B presents $\sim P$ as being true relative to B 's perspective. For, if either of them presented a proposition as being true absolutely or true relative to all perspectives or true *simpliciter* or whatever else, one could not infer from the relative truth of that proposition that the respective assertion is correct³—just as one cannot infer from the fact that EVENTS E AND E* ARE SIMULTANEOUS is true relative to a chosen frame of reference, that it is correct for a physicist who uses that frame (rather than any other), to present EVENTS E AND E* ARE SIMULTANEOUS as true *simpliciter*.⁴

³ In general, if A presents P as being F , where P 's being true relative to A 's perspective does not entail that P is F , one cannot infer from P 's being true relative to A 's perspective that A 's assertion is correct.

⁴ Note that it is not a good objection to this analogy to say that while CRICKETS ARE TASTY is a (complete) proposition, modern science has shown that EVENTS E AND E* ARE SIMULTANEOUS is not a (complete) proposition. The individuation of propositions is not a matter for physicists to

Note that in order to avoid this conclusion, it will not help to suggest that the *correctness* of an assertion too needs to be relativized, viz. to the relevant speaker's perspective. If *A* does not merely present *P* as being true relative to *A*'s perspective, the fact that *P* is true relative to *A*'s perspective will not even allow us to derive the *relative* correctness of *A*'s assertion. As long as my standards of taste or aesthetic value are informed by the fact that others may not share them, they will not license my presenting VINCENT PRICE WAS A TERRIFIC ACTOR or CRICKETS ARE TASTY as being true absolutely, even if, relative to those standards, these propositions should be true. In the same vein, my standards of taste will not license my preparing crickets for Christmas dinner (as it happens, my family finds crickets disgusting). Indeed, even if my standards of taste or aesthetic value are not informed by the fact that others may not share them—after all, that they are not shared is not itself a matter of taste or aesthetic value—it still is quite unintelligible how they could possibly render correct my presenting VINCENT PRICE WAS A TERRIFIC ACTOR or CRICKETS ARE TASTY as being true relative to all perspectives, given that, relative to some perspectives (e.g. my brother's), these propositions are not true.

To be sure, Kölbel himself nowhere suggests that truth-relativism alone could establish the possibility of faultless disagreement. Quite the contrary, he is explicit that, to this end, truth-relativists would further have to reject the principle

(T) it is a mistake to believe (assert) a proposition that is not true.⁵

Although Kölbel agrees that '(T) is a fundamental principle governing belief' (Kölbel (2003), 66–7), he goes on to argue that (T) needs to be modified when those non-objective propositions are concerned disagreement about which we ordinarily think of as faultless (Kölbel (2003), 68–9). Relativism about propositional truth is then said to provide us with the right kind of modification, viz.

(TR) it is a mistake to believe (assert) a proposition that is not true in one's own perspective

(Kölbel (2003), 70).⁶ Evidently, (TR) does not suffice for running the argument against faultless disagreement mentioned towards the end of section 2, provided only that *P* is true relative to *A*'s perspective while $\sim P$ is true relative to *B*'s perspective. So, once unmodified (T) is replaced by (TR), it seems that relativism about propositional truth provides a means to block that argument.

However, as long as *A*'s and *B*'s assertoric performances are not relativized to their respective perspectives, unmodified (T) will remain in force, the validity of

decide. That is, there are no specifically scientific reasons for preferring a semantic theory according to which EVENTS *E* AND *E** ARE SIMULTANEOUS is not a proposition, but only part of one, to one according to which it is, yet one whose truth is relative to frames of reference.

⁵ Although Kölbel explicitly refers only to beliefs, it should be clear that assertions as expressions of beliefs at least are subject to the norms for correct belief.

⁶ See previous footnote.

(TR) and truth-relativism notwithstanding. If to believe (or assert) a proposition is to acknowledge (or present) it as true *simpliciter*, (T) will continue to be ‘a fundamental principle governing belief’ (and assertion), no matter which kind of proposition is concerned. So, if *A* acknowledges *P* as true *simpliciter*, although *P* is not true *simpliciter*, but merely true relative to *A*’s perspective while false relative to *B*’s perspective, *A* will be at fault, no matter what *P* is. Any argument to the effect that unmodified (T) cannot be upheld must therefore be at least as strong as an argument for a relativistic conception of what *A* and *B* are doing when performing their respective assertions (or adopting the corresponding doxastic attitudes).

10.5 THE RELATIVIST’S DILEMMA

But now, once *what A and B are doing* is conceived in relativist terms so as to allow for the correctness of their doings and hence their faultlessness, the finding that *A* asserts *P* and *B* asserts $\sim P$ will no longer legitimize description of the envisaged case as a case of *disagreement*. We would have no more right to call it that than we would have to describe a case in which *A* presents *P* as worthy of consideration and *B* presents $\sim P$ as an example of what is believed by Catholics, as a case of disagreement just because *P* and $\sim P$ cannot both be true.

Elsewhere Kölbel argues that it is a necessary condition for *A* and *B* to disagree that *A* ‘cannot rationally accept what [*B*] has asserted without changing [his mind]’ (Kölbel (2004), 305). This condition is clearly met even if *A*’s and *B*’s assertions should be relativized in the way suggested.⁷ But equally clearly this condition is not sufficient for disagreement as long as, in asserting $\sim P$, *B* does something other than what *A* would do, did *he* assert $\sim P$. This becomes clear once we consider a case in which the difference in kind between *A*’s and *B*’s performances is not obscured by uniform, but systematically ambiguous, description: if *A* presents *P* as true relative to his perspective and *B* presents $\sim P$ as an example of what is believed by Catholics, *A* cannot rationally accept what *B* has thus presented without changing his mind, and yet *A* and *B* do not disagree.

What is missing in such cases is the possibility for *A* to perceive *B*’s attitude towards $\sim P$ as being incompatible with the attitude he himself adopts towards *P*—incompatible in the sense that, were *A* to adopt *B*’s attitude towards $\sim P$, he would have to change his mind. Whatever reasons the relativist may have for saying that *A* cannot present or acknowledge $\sim P$ as true *relative to B’s perspective*, this alleged impossibility would do nothing to rule out that if *per impossibile* *A* were to adopt *B*’s attitude towards $\sim P$, he would not have to change his mind. For, in that case *P* could still be true relative to *A*’s perspective.

⁷ In this respect, Kölbel’s relativism about truth crucially differs from those versions of relativism that relativize the propositions expressed (see Kölbel (2004); 303–8; cf. as well Lasersohn (2005)).

So, recasting Kölbel's definition of faultless disagreement in terms of relativized belief has the untoward effect that we can no longer view what is being defined as a case of disagreement.⁸ Thus, although there is a way of describing what *A* and *B* are doing such that *A* and *B* can still be said to do the same—either presents or acknowledges a proposition as true relative to his or her own perspective—that uniform description does nothing to alleviate the fact that *A* and *B* can no longer perceive each other as making incompatible moves in the same game, individuated by the same set of norms. In this respect, (TR) is more like a schema for such norms, whose instances vary depending on which attitude is designated by 'believe' (or 'assert'), than it is the expression of any particular such norm.

In other words then, Kölbel's relativist faces a dilemma: either *A* and *B* are said merely to present *P* and $\sim P$, respectively, as being true relative to their own perspective, in which case the relative truth of *P* and $\sim P$ ensures that their assertions are correct but there is no longer any genuine disagreement between them. Or else, they are said to present these propositions as being true *simpliciter* (true absolutely, true relative to all perspectives), in which case they do indeed disagree but relativism fails to show that their disagreement is faultless. For, the relative truth of the propositions asserted no longer ensures the correctness of their assertions given that (T) remains in force.⁹ Thus, how there could be cases of faultless disagreement remains a mystery, even if relativism about propositional truth should be correct.¹⁰

⁸ Kölbel has responded to this charge by claiming that there is a perfectly valid sense of 'disagreement' in which *A* and *B* already disagree if *A* presents *P* as true relative to his perspective and *B* presents *Q* as true relative to her perspective and *P* and *Q* cannot both correctly be presented as true relative to any perspective (private communication; cf. as well Lasersohn (2005), 683–4). The issue would then seem to be a merely verbal one, viz. about what to call 'a disagreement'. But for the reasons outlined above, I do not think that this diagnosis is apt.

⁹ One might wonder why these two conceptions of assertions should exhaust the options. But in light of Kölbel's remark that 'it seems wrong to say that truth is not a norm of belief in the discretionary area', no other option comes to mind (Kölbel (2003), 70).

¹⁰ MacFarlane seems to face a structurally similar dilemma when he argues that relativism about truth for future contingents allows us to reconcile indeterminism about the future with the retrospective assessment of assertions about the future as determinately true (correct) or false (incorrect). Addressing an objection by Evans (1985) akin to the one here raised for Kölbel, he writes: 'we can still think of assertions as commitments to the truth of the sentences asserted (at their context of utterance), for to be committed to the truth of a sentence (at a context of utterance *u*) is simply to be obliged, if challenged at any context of assessment *a*, to give adequate reasons for thinking that the sentence asserted is true (with respect to *u* and *a*), and to withdraw the assertion if the challenge cannot be met' (MacFarlane (2003), 336). Now, it should be clear that the relevant challenges need not be actual, otherwise one could fulfill the said obligation by simply avoiding such challenges, e.g. by leaving the room. So the characterization of this obligation must be taken to involve a subjunctive conditional rather than a material one. But then, in so far as the future is objectively indeterminate, that is itself a challenge that the asserter of a future contingent could not meet at the time of utterance, whether it be put to him or not (cf. MacFarlane (2003), 335). Accordingly, if the commitment undertaken by making the assertion at *u* engenders the obligation to answer challenges available at *any a*, then the fact that the future contingent is true relative to some later context of assessment does nothing to change the fact that the relevant assertion is incorrect and hence should be withdrawn even later. On the other hand, if there is no such general

10.6 A RELATIVIST ACCOUNT OF AGREEMENT AND DISAGREEMENT

Our discussion so far suggests that on a relativistic conception of assertion (and belief), we cannot make sense of genuine disagreement at all, be it faultless or not. This seems unpalatable because it would now appear that, for the very same reasons, we cannot make sense of genuine *agreement* on that conception either. In order for *A* to genuinely agree with *B*, it must at least be possible for *A* to perceive *B* as adopting the same attitude towards the same proposition. Yet, if *B* presents *P* as true relative to *B*'s perspective, while *A* presents *P* as true relative to *A*'s perspective, the assumption that *A* and *B* thereby adopt the same attitude towards *P* is entirely unwarranted. To see why, let us assume that *A*'s and *B*'s perspectives differ so that *P* is true relative to *B*'s perspective just in case *P* is false relative to *A*'s perspective. In this case, *A* and *B* evidently do not agree as not both assertions can be correct.

The problem arises because of a certain asymmetry between truth and correctness: whereas propositions are said to be true only relative to a perspective, the correctness of assertions appears to be determined, fully and absolutely, by two interlocking factors: the nature of the relevant assertion as an acknowledgement of *P*'s being *F*, and *P*'s being *F*. We saw in section 4 that even if *P* is true relative to *B*'s perspective, this will not yet allow us to speak of the relative correctness of *B*'s assertion of *P* as long as that relativity is not reflected in what *B* presents *P* as being. But once *B* is understood to present *P* as being true relative to his own perspective, his assertion would seem to be correct or incorrect absolutely, depending on whether or not *P* is true relative to *B*'s perspective. So even where *F*-ness is taken to be truth relative to a certain perspective, there seems to be no space left for further relativizing the correctness of an assertion with content *P* (cf. Evans (1985), 349).

If it was a matter of perspective *what it takes for P to be what B presents it as being*, then there would after all be a sense in which the correctness of *B*'s assertion was not absolute. However, such a further relativization was ruled out by the idea that both *A*'s assertion and *B*'s assertion should be faultless and hence faultless absolutely. Since such faultlessness could only be achieved by denying that there is genuine disagreement between *A* and *B*, we should accordingly abandon that idea. We are then free to conceive of the correctness of assertions as relative to perspectives in the sense just suggested. It may then prove possible to recover a

commitment or obligation undertaken by making the assertion, but rather a series of commitments or obligations undertaken by letting the assertion stand as the contexts of assessment change (time passes), then the fact that the future contingent is true relative to some later context of assessment does nothing to show that *the original commitment* which it is the aim of *retrospective* assessments to evaluate was correct, unless indeterminism about the future is false after all.

sense in which *A* can perceive *B* as adopting the same or an incompatible attitude towards *P*.

The task accordingly becomes one of explaining how both of the following two inferences can be understood to be valid if premisses and conclusions are evaluated in terms of truth relative to *A*'s perspective, or *A*-truth for short:

(1)	<i>B</i> asserts <i>P</i> <i>P</i>	(2)	<i>B</i> asserts <i>P</i> $\sim P$
	—————		—————
	<i>B</i> correctly asserts <i>P</i>		<i>B</i> incorrectly asserts <i>P</i> .

We already know that, on a relativist conception of assertion, the following principles are true relative to all perspectives:

- (3) *B* (correctly) asserts *P* $\leftrightarrow$ *B* (correctly) presents *P* as *B*-true
 (4) *B* presents *P* as *B*-true $\rightarrow$ (*B* correctly presents *P* as *B*-true $\leftrightarrow$ *P* is *B*-true).

A fortiori, (3) and (4) are both *A*-true. In order to ensure that, in (1) and (2), the *A*-truth of the premisses implies the *A*-truth of the conclusion, what is still needed is the *A*-truth of

- (5) *P* is *B*-true $\leftrightarrow$ *P*.

With the *A*-truth of (5) being in place, both (1) and (2) prove valid if premisses and conclusions are evaluated in terms of *A*-truth. I here merely show this for (2):

(2')	(a) <i>B</i> asserts <i>P</i>	ass.
	(b) <i>B</i> presents <i>P</i> as <i>B</i> -true	(a), (3)
	(c) $\sim P$	ass.
	(d) $\sim(P$ is <i>B</i> -true)	(c), (5)
	(e) <i>B</i> incorrectly presents <i>P</i> as <i>B</i> -true	(b), (d), (4)
	(f) <i>B</i> incorrectly asserts <i>P</i>	(e), (3).

The question that now needs addressing is whether, and if so how, the idea that (5) is true relative to *A*'s perspective can be motivated.

10.7 MORAL RELATIVISM AND MORAL DISAGREEMENT

To begin with, it should be clear that in the case of those *P* that concern matters of taste or aesthetic value, it is hard to make sense of subjects whose perspectives map *P*'s truth-value relative to someone else's perspective onto the truth-value *P* has relative to theirs. But contrary to what Kölbel and others suggest, the appearance of genuine disagreement in such areas is faint. With respect to matters of taste or aesthetic value, we are prone to pleading faultlessness at the expense

of genuinely disagreeing: *de gustibus non est disputandum*. In the moral case, by contrast, we have strong intuitions that there is genuine disagreement, while the appearance of faultlessness is rather faint.

The contrast is best brought out by a comparison of practical conflicts concerning matters of taste or aesthetic value with those arising in the moral sphere. Suppose *A* and *B* have a practical conflict when it comes to deciding which movie to watch, because *A* denies, while *B* affirms, VINCENT PRICE IS A TERRIFIC ACTOR. If *A* loses, he may afterwards concede that although watching the movie was a horrid experience for him, watching it was nonetheless a good thing for *B* because, from *B*'s perspective, Vincent Price is a terrific actor. By contrast, if *A* and *B* have a practical conflict when it comes to averting a law against abortion, because *A* denies, while *B* affirms, ABORTION IS WRONG, no corresponding concession can be made. If *A* loses, he cannot afterwards concede that while the law's being passed was a bad thing for him, it was nonetheless a good thing for *B* (or anyone else, for that matter). And *A* may reaffirm ABORTION IS NOT WRONG as a way of giving his reasons for his intransigence. Yet, it is hard to make sense of all this unless *A* can be said to view the correctness of his assertion as undermining any reason *B* might have for passing a law against abortion, including the reasons *B* gives when she asserts ABORTION IS WRONG. The most natural way of interpreting this is by saying that, from *A*'s perspective, not both *A*'s assertion and *B*'s assertion can be correct.

Thus, in light of practical conflicts of this kind, simply to apply the dictum *de gustibus non est disputandum* to the moral sphere offends against our conception of what is at stake. According to our standard practice, what *A*'s claim must be taken to imply in such cases of conflict is that *B* erroneously believes abortion to be wrong and so is at fault. So the moral sphere is the best place to look for in order to motivate a type of relativism that invokes (5).

Why, though, should it be rational to adopt the policy of evaluating the assertions of others on the basis of our own view of the matter rather than theirs, as (5) suggests? Why should we be so uncharitable? The short answer is: because we have no choice. Even if there should be no objective fact of the matter that would privilege one perspective over any other, we cannot help but privilege our own. We have no other way of understanding what makes an assertion a claim specifically about moral truth (as we conceive it) save through reference to what our moral code identifies as true. Yet, neither can we simply interpret others as making claims about truth in some other sense of 'truth', viz. one we think inappropriate for moral matters. For, the context in which we find ourselves (according to our perception of it) would make any such claim deeply misplaced. Thus, in the context of what we think of as a moral debate, promoting a law against abortion for any but moral reasons not only strikes us as morally wrong, but also as at odds with the very point of the debate. On such occasions, taking others to express claims about truth in some non-moral sense would frustrate any attempt to make sense of their contributions to that debate. So, in this respect,

we are charitable interpreters after all. While ascribing to them a mistaken belief about moral truth (as we conceive it) has a chance of making sense of their contributions, ascribing to them a correct belief about non-moral truth (as we conceive it) has not.

10.8 RELATIVISM AND PARTIALLY DEFINED FUNCTIONS

In order to defend this kind of moral relativism, much more needs to be said than what has been said here. But instead of pursuing this matter any further, I will close by addressing the more pressing worry that any type of relativism that invokes (5) is ultimately incoherent. The worry is that relativism as thus conceived will inevitably collapse into absolutism about truth. Evidently, this collapse can be avoided only if (6) can be true:

(6) $\sim(P \text{ is } B\text{-true} \leftrightarrow P \text{ is } A\text{-true})$.

The problem then is that if (5) is true relative to any subject A 's perspective, there would seem to be no perspective relative to which (6) is true. Let (6) be true relative to C 's perspective, then we get:

(7) $\sim(P \text{ is } B\text{-true} \leftrightarrow P \text{ is } A\text{-true})$ is C -true.

Given that A and B were chosen arbitrarily, it would seem that what holds good for A should *mutatis mutandis* hold good for B and vice versa, and that what holds good for either should *mutatis mutandis* hold good for C . However, the C -truth of instances of (5) would yield

(5') $(P \text{ is } A\text{-true} \leftrightarrow P)$ is C -true

(5'') $(P \text{ is } B\text{-true} \leftrightarrow P)$ is C -true,

which jointly imply

(8) $(P \text{ is } B\text{-true} \leftrightarrow P \text{ is } A\text{-true})$ is C -true.

But (8) contradicts (7); and so it would seem that (6) cannot be C -true, for any C —contrary to what was assumed.

This proof crucially depends on the assumption that C 's perspective is a function from propositions to truth-values which is defined for P as argument. Otherwise, (5') and (5'') could not be derived from the A -truth of (5) and the fact that A and B were arbitrarily chosen from among those subjects whose perspectives evaluate either P or $\sim P$ as true. Thus, insofar as C 's perspective is a partially defined function which is not defined for non-objective propositions as arguments, neither (5') nor (5'') is available. Accordingly, the perspective relative to which (6) holds should be understood to be a perspective that fails to evaluate non-objective propositions. The theoretician's perspective is in this sense neutral

on non-objective matters.¹¹ With respect to moral matters, this is precisely the position of someone who concedes that moral propositions are truth-evaluable relative to some perspectives, yet whose own worldview ‘asks no more from the world than what we know is there—the ordinary features of things on the basis of which we make decisions about them, like or dislike them, fear them and avoid them, desire them and seek them out. It asks no more than this: a natural world, and patterns of reaction to it’ (Blackburn (1984), 182). In other words then, moral relativism of the intended kind turns out to be a type of non-cognitivism in cognitivist clothing.

However, it is noteworthy that this relativist account does not face the Fregean problem that besets traditional expressivism. Relativists of this stripe can accept and explain that arguments such as (9) are valid:

- (9) Murder is wrong
 If murder is wrong, promoting murder is wrong
-
- Promoting murder is wrong.

Any perspective relative to which both premisses are true will be a perspective relative to which the conclusion is true; and anyone who acknowledges both premisses as true relative to his own perspective, but acknowledges the conclusion as false relative to that same perspective, will be in error. What relativists of this stripe cannot do is accept arguments like (9) as sound, where the soundness of an argument requires that its premisses be true. But neither are they committed to holding that these premisses are, relative to their own perspective, false. As long as truth and falsity are relative to perspectives and perspectives are functions from propositions to truth-values, no inconsistency looms: the possibility of partially defined functions ensures this.

Relativism of the kind suggested here may in the end be insufficiently well motivated—in fact, it may even be deemed immoral for its reluctance to countenance any moral truths. But it is at least coherent and can make clear sense of agreement and disagreement without collapsing into absolutism about truth. Faultless disagreement, however, is anathema to relativists of this kind.¹²

¹¹ Insofar as (6) is objective, the idea, mooted in section 3, that objective propositions are, if true, true relative to all perspectives must of course be rejected: otherwise (7) would imply that (6) is *A*-true, which would be inconsistent with the *A*-truth of (5). Accordingly, the moral relativism suggested here implies an error-theory about the participants’ conception of moral discourse. Thanks to Elia Zardini and Martin Smith for alerting me to this.

¹² Work on this paper was made possible by a DFG Heisenberg fellowship. I have benefited from discussions with Herman Cappelen, Fabrice Correia, Manuel García-Carpintero, Patrick Greenough, Jörg Hardy, Philipp Keller, Max Kölbel, Sebastiano Moruzzi, Olaf Müller, Thomas Schmidt, Martin Smith, Crispin Wright and Elia Zardini.

References

- Ayer, Alfred J. (1946). *Language, Truth and Logic*, 2nd edn., London: Victor Gollancz.
- Blackburn, Simon (1984). *Spreading the Word*, Oxford: Oxford University Press.
- Evans, Gareth (1985). 'Does Tense Logic Rest upon a Mistake?', in G. Evans, *Collected Papers*, Oxford: Oxford University Press, 341–63.
- Frege, Gottlob (1892). 'Über Sinn und Bedeutung', in G. Frege, *Funktion, Begriff, Bedeutung*, Günther Patzig (ed.) (1986), Göttingen: Vandenhoeck and Ruprecht, 40–65.
- (1897). 'Logik', in G. Frege, *Schriften zur Logik und Sprachphilosophie*, Gottfried Gabriel (ed.) (1990), Hamburg: Meiner Verlag, 35–73.
- (1906). 'Einleitung in die Logik', in G. Frege, *Schriften zur Logik und Sprachphilosophie*, Gottfried Gabriel (ed.) (1990), Hamburg: Meiner Verlag, 74–91.
- (1918/19). 'Der Gedanke. Eine logische Untersuchung', in G. Frege, *Logische Untersuchungen*, Günther Patzig (ed.) (1986), Göttingen: Vandenhoeck and Ruprecht, 30–53.
- Geach, Peter (1960). 'Ascriptivism', *Philosophical Review* 69: 221–5.
- (1965). 'Assertion', *Philosophical Review* 74: 449–65.
- Kölbel, Max (2002). *Truth Without Objectivity*, London: Routledge.
- (2003). 'Faultless Disagreement', *Proceedings of the Aristotelian Society* 104: 53–73.
- (2004). 'Indexical Relativism versus Genuine Relativism', *International Journal of Philosophical Studies* 12: 297–313.
- Laserson, Peter (2005). 'Context Dependence, Disagreement, and Predicates of Personal Taste', *Linguistics and Philosophy* 28: 643–86.
- MacFarlane, John (2003). 'Future Contingents and Relative Truth', *Philosophical Quarterly* 53: 321–36.
- (2005). 'Making Sense of Relative Truth', *Proceedings of the Aristotelian Society* 105: 321–39.
- Shapiro, Stewart and Taschek, William W. (1996). 'Intuitionism, Pluralism, and Cognitive Command', *Journal of Philosophy* 93: 74–88.
- Wright, Crispin (1992). *Truth and Objectivity*, Cambridge, Mass.: Harvard University Press.
- (2003). *Saving the Differences*, Cambridge, Mass.: Harvard University Press.

This page intentionally left blank

Epistemic Modals and Correct Disagreement

Richard Dietz

Epistemic modals are devices of marking the epistemic possibility/necessity of an underlying proposition. For example, an utterance of ‘It might be raining now in Sydney’ is true just in case the proposition that it is raining in Sydney at the utterance time is possible in view of what is known in the relevant epistemic situation—or so the standard truth-conditional approach to epistemic modals suggests. According to relativists about epistemic modals, the epistemic situation that is relevant to the truth-valuation of a given epistemic modal statement may vary with occasions of its assessment. I aim to show that the standard relativist account of epistemic modals is wrong, not only in letter but also in spirit. Furthermore, it is suggested that a puzzle which has been invoked in support of relativism about epistemic modals can be dissolved in non-relativist terms.

11.1 INTRODUCTION

Epistemic modals are expressions like ‘might’, ‘possibly’, ‘there is a possibility that’ or ‘it is necessary that’ as used in an *epistemic* sense. For instance, the epistemic *might* is a device of marking a relevant proposition as *epistemically possible*, in other words as possible in view of what is known. More precisely, the *relevant* proposition is the proposition expressed by the embedded sentence, henceforth *the underlying proposition*.¹ For example, in saying

Earlier versions of this paper were presented at the Arché Research Centre in St. Andrews (2005 and 2006), at the Logos workshop *Relativizing Utterance Truth* in Barcelona (2005), at the Ockham Society in Oxford (2005), at the University of Helsinki (2006), at the Joint Session Meeting of the Aristotelian Society and the Mind Association in Southampton (2006), and at the University of Münster (2006). I would like to thank all audiences for stimulating questions. In particular, I would like to thank Herman Cappelen, Iris Einheuser, John MacFarlane, Sebastiano Moruzzi, Julien Murzi, Mark Sainsbury, Isidora Stojanovic, Timothy Williamson and Crispin Wright. My work on this paper was funded by the Arché Research Centre for Logic, Language, Metaphysics and Epistemology, to which also thanks are due.

¹ Or the *prejacent*, to use the technical term which is common in linguistic literature.

(Rain in Sydney) It might be now raining in Sydney,

I mark the proposition that it is raining in Sydney at the time of utterance as epistemically possible. Whether a proposition is epistemically possible depends on the relevant *epistemic situation*—e.g. as I do not have any evidence concerning the current weather in Sydney, it is in view of what I know possible that it is raining now in Sydney, though it may not be possible in view of what some people in Sydney know (suppose that sunny weather conditions are just observed by some people in Sydney). On the standard *truth-conditional* approach, epistemic modals are taken as devices of making *truth-valuable* statements of epistemic possibility/necessity.² For the standard examples *might* and *cannot*, truth-conditional accounts standardly follow the schemas:

might An utterance of a *might* sentence is true just in case the underlying proposition is an *epistemic possibility* relative to the relevant epistemic situation

and

cannot An utterance of a *cannot* sentence is true just in case the underlying proposition is *not an epistemic possibility* relative to the relevant epistemic situation

respectively. Different accounts of the informal term ‘the relevant epistemic situation’ within the standard account schema **might** may suggest different truth-valuations for the same epistemic *might* statement. The *received view* has it that the relevant epistemic situation is fixed by *contexts of utterance*. *Relativists* about epistemic modals, by contrast, contend that the relevant epistemic situation may vary with (if it is not fixed alone by) respective *contexts of assessment*. Before discussing the argument in favour of relativism, I need to set the stage first. I explain the fundamental notion of epistemic possibility in more detail and compare the standard contextualist account of epistemic modals with its relativist rival. For simplicity, as usual, the epistemic *might* is referred to as the standard illustrative example.

11.1.1 Epistemic Possibility

We say that a proposition *p* is epistemically possible relative to what is known in an epistemic situation *S* just in case *p* is consistent with what is known in (or ‘possible in view of what is known in’) *S*. Let us cash out the basic notion of *epistemic possibility* in more exact set-theoretic terms. *Propositions* are modelled

² For a systematic survey article on truth-conditional formal semantics for various sorts of modality, see von Stechow (2005). I do not discuss gradable modifiers like ‘probably’ or ‘it is highly probable that’, which allow for more differentiated evaluations of the relevant epistemic situation.

as sets of tuples of relevant parameters.³ On the received view, only possible circumstances will be relevant, but certain views suggest that propositional truth-values may vary with yet further parameters such as time or standards of taste—our discussion will be neutral regarding this issue. A proposition p entails a proposition q just in case p is a subset of q ; a set of propositions Γ entails a proposition p just in case the intersection of all members of Γ entails p . A proposition p is consistent with a proposition q just in case p and q are not disjoint; a proposition p is consistent with a set of propositions Γ just in case p is consistent with the intersection of the members of Γ . *Epistemic situations* are represented by the intersection of all propositions known (in the situation). The epistemic situation of an individual i at circumstances w and a time t is the intersection of all propositions known by i at w at t . As regards the more general case that we speak of the epistemic situation of a community of epistemic subjects at respectively relevant worlds and times $\{ \langle i_1, w_1, t_1 \rangle, \dots, \langle i_n, w_n, t_n \rangle \}$, where n may be greater than 1, it is not necessary here to give a more exact account. There is more than one serious option, e.g. plausibly one may say that what is known in $\{ \langle i_1, w_1, t_1 \rangle, \dots, \langle i_n, w_n, t_n \rangle \}$ amounts to the intersection of all propositions p where there is a d such that $1 \leq d \leq n$ and p is known by i_d at world w_d at time t_d (*Partial Knowledge*); or one may say that it amounts to the intersection of all propositions p where p is known by i_d at w_d at t_d , for all d such that $1 \leq d \leq n$ (*Distributed Knowledge*).⁴ For my purposes, I need to assume only that the notion of *epistemic situation* is used in a way no weaker than the latter, distributive notion. More formally:

Epistemic Situation

The epistemic situation of $\{ \langle i_1, w_1, t_1 \rangle, \dots, \langle i_n, w_n, t_n \rangle \}$ entails $\bigcap_{1 \leq d \leq n} \{ p : p \text{ is known by } i_d \text{ at } w_d \text{ at } t_d \}$.

Epistemic possibility is then constrained as follows:

Epistemic Possibility

Proposition p is epistemically possible relative to an epistemic situation S iff $S \cap p$ is not empty.

11.1.2 Contextualism

On a standard *contextualist* account of epistemic modals, the respectively relevant epistemic situations are thought to be functions of utterance contexts. That

³ One may conceive of more fine-grained (e.g. Russellian or Fregean) accounts, according to which propositions are entities that determine sets of tuples of relevant parameters. None of the points which will be made will hinge on this issue.

⁴ For the suggestion of bringing distributed knowledge to bear on epistemic modality, compare Teller (1972) and more recently Gillies and von Fintel (2004). For still other options, see MacFarlane (forthcoming).

is to say, the truth-conditions of *might* are thought to be constrained as follows:

might/contextualism

‘*might*: P ’ is true at a context of utterance c iff the proposition expressed by ‘ P ’ at c is consistent with C_c , otherwise ‘*might*(P)’ is false at c ,

where for context, the function C takes c as value the intersection of all propositions known in the epistemic community relevant to c at the time of c .⁵

The account schema **might/contextualism** leaves it still entirely open to what extent the relevant epistemic situation is supposed to vary with utterance contexts. That it may vary both with the circumstances and the time of utterances, seems undeniable: If I utter (Rain in Sydney), I do not mean to make a statement as to a merely metaphysically possible epistemic situation, or as to the actual epistemic situation at some time in the past or in the future. Whether my statement is correct depends on whether the proposition that it is raining in Sydney at the time of utterance is possible in view of what is known in a relevant community under the *actual* circumstances and at the *present* time. Or so one may reasonably argue. John MacFarlane refers to the view that the relevant epistemic situation is invariably set by the world and the time of utterance as *universalism*.⁶ Here is an argument against universalism:⁷ Suppose I know a proposition p . Any sentence of the form ‘*might*: $\sim P$ ’, where P expresses p relative to my utterance context, is then false as uttered by me. The most natural way of accommodating this intuition is to follow the received line and to adopt the following constraint:

Speaker Inclusion

For any context of utterance c : C_c entails $\{p$: p is known by s_c at w_c at t_c , where s_c , w_c and t_c are the speaker, world and time of c respectively}.⁸

Now if the relevant epistemic situation were not to vary with speakers for fixed circumstances and times of utterance, it could meet Speaker Inclusion only at the price of entailing for any world-time pair $\langle w, t \rangle$, any proposition which is known by at least some potential epistemic subject at w at t . In other words, the epistemic situation is to be thought of as what is *partially* known in the class of all potential epistemic subjects under the relevant circumstances at the relevant time.⁹ As a consequence, if someone knows p (at a world w and time t), any

⁵ For semantic accounts of epistemic modals in the contextualist tradition, see Kratzer (1977), DeRose (1991), and Gillies and von Fintel (2004).

⁶ MacFarlane (2003: § 4). Compare Egan *et al.* (2005: § 3).

⁷ For arguments in the same spirit, see MacFarlane (2003: § 4) and Egan *et al.* (2005: § 3).

⁸ DeRose (1991: 596).

⁹ MacFarlane (2003: § 4) suggests that on universalism, for a proposition to be epistemically possible, it needs not only to be compatible with what is *partially known* in the universal class of all potential epistemic subjects, but also with what is *partially knowable* in this class by any way

utterance of any sentence of the form $\ulcorner \text{might} : \sim P \urcorner$ (at w and t), where p is the underlying proposition, would be false. But this result would in effect rule out any correct uses of the epistemic *might* such as in:

(Brown in Town) Brown might be in town. He might not be. I don't know.
But surely Brown's wife knows where he is.

It seems incontrovertible that speakers use sometimes the epistemic *might* in an *egocentric* way, in the sense that they mean to make a statement only as to their personal epistemic situation.¹⁰ And it seems as incontrovertible that such uses are correct, insofar as they mirror adequately the epistemic situation of the speaker. Consequently, universalism cannot be considered as a viable option. There is good reason not only for rejecting universalism but also for letting the relevant epistemic situation vary with *speakers*.

11.1.3 Relativism

The previous considerations seem to tell in favour of a standard contextualist approach to epistemic modals. In recent time though a number of authors have argued that this approach is to be replaced by a different, *relativist* approach. According to relativism about epistemic modals, the relevant epistemic situation for a given epistemic modal statement may vary with assessments. This suggests the following account schema:

might/relativism

$\ulcorner \text{might} : P \urcorner$ is true at a context of utterance c and an assessment context d iff the proposition expressed by $\ulcorner P \urcorner$ at c is consistent with C_d , otherwise $\ulcorner \text{might}(P) \urcorner$ is false at c and d .

We may distinguish different versions of relativism, according to the parameters with respect to which the relevant epistemic situation is supposed to be assessment-relative. Insofar as we can ignore assessments of statements across worlds, the room for a genuine relativism effectively narrows down to the assessment time and the assessor. Tamina Stephenson (2005) argues for a relativism on which the class of relevant epistemic situation may vary with assessors.¹¹ Most relativists about epistemic modals have taken a more radical line. For example, John MacFarlane (2003) and (forthcoming) defends a relativism, according to which the relevant epistemic situation may also vary with the assessment time.

of obtaining knowledge. My argument against universalism does not need consideration of this stronger version of universalism.

¹⁰ MacFarlane (forthcoming) and Gillies and von Fintel (2006) speak of 'solipsistic' uses of epistemic modals. I prefer the term 'egocentric'. Thanks to Crispin Wright for the suggestion.

¹¹ Basically this is an application of Lasersohn's account of predicates of personal taste, in (2005), to epistemic modal discourse.

Egan *et al.* (2005), Andy Egan (2007), and Iris Einheuser (2008) seem to agree with MacFarlane in this point.¹² Specifically, the idea is that the relevant epistemic situation function is *egocentric* in the sense that it takes for any assessment context as value the epistemic situation of the assessor at the assessment time (under the circumstances of assessment):

Community/assessment-sensitive/egocentric

For any context of assessment d : $C_d = \{p: p \text{ is known by } a_d \text{ at } w_d \text{ at } t_d, \text{ where } a_d, w_d \text{ and } t_d \text{ are the assessor, world and time of } d \text{ respectively}\}$.

11.1.4 The Semantics and Metasemantics of Epistemic Modal Discourse

According to standard contextualism, all relevant parameters to the truth-value of epistemic sentences are provided by the utterance context.¹³ Relativism, on the other hand, strongly suggests that with the epistemic situation, some relevant parameters are rather fixed by the assessment context. Depending on whether we follow a contextualist or a relativist approach to epistemic modals, we are hence led to different positions about the semantics for epistemic modal statements. It needs to be stressed though that the distinction contextualism/relativism is neutral regarding the *metasemantic* question of what sort of role the particular relevant parameters play in the fixation of the truth-value.

For one, a parameter x may be held to be relevant to the truth-value of a sentence on the ground that it is relevant to the fixation of its propositional content. In that case, we hold an *indexical* view about the role of x . For another, x may be held to be relevant to the truth-value of a sentence on the ground that it is relevant to the truth-value of the proposition expressed by the sentence (in other words, x is treated like the world parameter, with respect to which propositions are to be evaluated as true/false). In this case, we hold a *non-indexical* view about the role of x . To illustrate the indexical view, the received line on sentences involving the indexical 'I' is to say that the propositional content varies with utterance contexts depending on the speaker, which is given by utterance contexts—e.g. if Didi utters 'I am going to Sydney', his utterance says something *about* Didi, whereas as uttered by Naomi, the sentence says something *about*

¹² Einheuser is less concerned with the semantic account of epistemic modals than with the metaphysics of epistemic possibility. But she agrees with a semantic account along MacFarlane's line (p.c.).

¹³ Insofar as the propositional content of embedded sentences may vary with assessments, there is, of course, room for the assessment-sensitivity of epistemic modal sentences, even on a contextualist account of epistemic modal sentences. What is at issue here though is the question of whether the contribution of *epistemic modals* to the truth-conditional content of sentences is assessment-sensitive. We are therefore free to ignore any potential assessment-sensitivity of the truth-conditions for the embedded sentence.

Naomi.¹⁴ A relativist about epistemic modals may take an analogue indexical line about the assessor parameter. This would suggest that as assessed by Naomi, an utterance of (Rain in Sydney) says something about *her* epistemic situation, whereas as assessed by Didi, the same utterance says something about *his* epistemic situation (analogously for assessment times, on a radical version of relativism). The previously suggested versions of a relativism about epistemic modals all take a *non-indexical* view about the supposed relevant parameters. According to this, an utterance of (Rain in Sydney) has a fixed propositional content relative to varying assessments, but the content may vary in truth-value with respect to assessors (analogously for assessment times, on a radical version of relativism).

To a large part, my discussion will be neutral regarding the indexicalism/non-indexicalism distinction.¹⁵ The distinction will be however important in my reply to a particular possible objection (see § 6.3).¹⁶

11.1.5 Disclaimers

At this point, a couple of disclaimers are in order:

- (i) *Epistemic modals embedded*. The question of what epistemic modals contribute to the truth-conditions of compound sentences in which they do not take the widest scope has been only partially a subject of serious investigation. For the purposes of this discussion, it is not essential to take a stand on controversial issues such as how epistemic modals interact with tense, propositional attitude verbs or quantification.¹⁷
- (ii) *Norms for knowledge*. One may argue that insofar as norms for knowledge vary with contexts of utterance (or assessment), the underlying notion of an epistemic situation is to be refined, by making it sensitive to norms for knowledge. Indeed insofar as there is room for a relativism about knowledge ascriptions, there seems to be also room for another version of relativism about epistemic modals.¹⁸ However, I follow the standard line and focus on the relativist claim that there are types of assessment-sensitivity—viz.

¹⁴ Kaplan (1989).

¹⁵ The distinction between indexical and non-indexical views about parameters relevant to sentential truth goes back to Cresswell (1973). For further discussion, see MacFarlane (2005b), Glanzberg (2007) and Cappelen (2008).

¹⁶ Wright (2007) offers other arguments against a truth-relativism about epistemic modals. Discussion of Wright's points would go beyond the scope of this paper.

¹⁷ There is a vast linguistic literature on the question of how epistemic modals interact with tense. Condoravdi (2001) offers a critical survey. Dietz (2005) and Gillies and von Stechow (2008a: § 5) argue that a relativist account of epistemic modals has absurd consequences for epistemic modals in the scope of the past tense. As for quantifiers, see von Stechow and Iatridou (2003). For discussion of embeddings of epistemic modals in various other types of contexts, see von Stechow and Iatridou (2002).

¹⁸ Compare n. 31. For relativism about knowledge ascriptions, see MacFarlane (2005a).

as to the relevant epistemic subjects and/or the relevant time—that are characteristic of epistemic modal discourse. As for the further relativist claim that norms for knowledge are assessment-sensitive, it seems more natural to discuss it in connection with knowledge-ascriptions.

- (iii) *Evidentiality*. It is a familiar (if obscure) idea that at least some epistemic modals convey information about the type of available evidence as to the underlying proposition. The standard illustrative example are modals like ‘must’, or ‘cannot’, which are used in apodeictic statements. Modals like these seem to indicate that there is inferential evidence, but not direct observational evidence, in favour of the underlying proposition (e.g. to say ‘It must be raining’ in a context where you directly observe that it is raining seems odd).¹⁹ Insofar as some epistemic modals do incorporate a kind of evidential meaning component, this semantic component is not captured by the meaning component that is at issue in standard truth-conditional semantics. I focus on the latter component, since this is the very component intended by relativists about epistemic modals.
- (iv) *Non-factuality*. It is common ground between standard contextualists and relativists that epistemic modals contribute to the truth-conditional content of sentences, in the sense that they make a truth-valuable comment as to the underlying proposition. In the linguistic literature, it is often claimed that this assumption is wrong-headed—according to most suggestions along this line, epistemic modals serve just as means of expressing partial or full commitment to the underlying proposition.²⁰ More recently, various proposals have been made to cash out this idea in more exact formal semantic terms.²¹ Both a discussion of the alleged evidence for non-factuality and a discussion of the particular proposals would go beyond the scope of this paper. For the sake of argument, I start from the factuality assumption.

11.1.6 Preview

To begin with, I set out what may be considered as the master-argument in support of a relativism about epistemic modals (§ 2). Standard relativism not only supplies means of accommodating certain puzzling scenarios of apparent disagreement involving epistemic modals, where either part seems to make a semantically correct statement. I argue that it also predicts scenarios of correct disagreement involving epistemic modals which are utterly bizarre (§ 3).

¹⁹ For discussion of this aspect of epistemic modal statements, see Gillies and von Fintel (2008b: §3).

²⁰ For critical discussion of the literature, see von Fintel and Iatridou (2002). In the philosophical and also in formal semantics literature, this view has been less influential. But see Price (1983) for an early exception from the rule.

²¹ Compare Veltman (1996), Yalcin (2005), Swanson (2005), Yablo (2006), and Gillies and von Fintel (2008b).

My discussion of a possible escape route will suggest that standard relativism is wrong not only in letter but also in spirit (§ 4). I aim to show that there is a more powerful standard contextualist account of the puzzling evidence that has been produced in favour of a relativism about epistemic modals (§ 5). Finally, I consider some possible objections against the suggested contextualist account (§ 6).

11.2 A PUZZLE ABOUT EPISTEMIC MODALS

Why should we adopt a relativism about epistemic modals? Insofar as there is a master-argument for the relativist case, it has been formulated most clearly in MacFarlane (2003: Section 2) and (forthcoming).²² The argument invokes possible dialogue scenarios which are puzzling, since they appear to suggest that there is room for *correct disagreement* over the truth-value of a given epistemic modal statement.²³

In general, the dialogue scenarios concerned have the following structure: *Knowing Assessor*. Someone, call her *Ignorant*, makes an epistemic *might* statement (e.g. of the form ‘Goldbach’s conjecture might be true’). For *Ignorant*, it is epistemically possible that the underlying proposition holds (e.g. in the concerned example, *Ignorant* does not know whether Goldbach’s conjecture is true). Some other subject, call her *Knower*, is in a better epistemic situation as to the underlying proposition in the sense that she knows whether it holds. *Ignorant*’s epistemic *might* statement seems then correct, as tested against *Ignorant*’s epistemic situation. On the other hand, it seems that *Knower* can correctly disagree with *Ignorant*’s statement—plausibly, either in form of an epistemic modal statement (‘Goldbach’s conjecture cannot be true’), or in form of a statement that seems to be equivalent to a predication of falsehood of *Ignorant*’s statement (e.g. ‘That’s false’).

Let us take stock of what is suggested for Knowing Assessor scenarios. We have:²⁴

(CD1) *Correctness of Ignorant’s statement by her standards:*

Ignorant’s epistemic modal statement meets the test for correctness as tested against her perspective—if both you knew that you were in an epistemic situation like *Ignorant*’s one and made a statement of the same form as *Ignorant*’s one, then your knowledge would provide conclusive evidence for the correctness of your statement.

²² For similar lines of argument, compare Egan *et al.* (2005) and Stephenson (2005).

²³ MacFarlane (forthcoming) refers to Hawthorne (2004: 27, n. 68), who expresses the suspicion that “[. . .] there is a real puzzle here”.

²⁴ This is my way of reconstructing the puzzle. MacFarlane acknowledges that the reconstruction is adequate (p.c.).

(CD2) *Disagreement between Ignorant and Knower:*

Knower makes a statement which disagrees with *Ignorant's* statement, in the following *minimal* sense: against whatever perspective they are assessed, they cannot jointly meet the test for correctness.²⁵

(CD3) *Correctness of Knower's statement by her standards:*

Knower's statement meets the test for correctness as tested against her perspective—if both you knew that you were in an epistemic situation like *Knower's* one and you made a statement of the same form as *Knower's* statement of disapproval, then your knowledge would provide conclusive evidence for the correctness of your statement.

Insofar as (CD1–3) are not only individually but also jointly plausible, they indeed generate a real puzzle: If both *Knower's* and *Ignorant's* statement are correct as tested against *Knower's* and *Ignorant's* standards respectively, how can they disagree?²⁶ Here is an example of the *Knowing Assessor* scenario type. Time is assumed to be irrelevant, in the sense that insofar as the two relevant statements are made at different times, there is no change in the epistemic situation as to the underlying proposition on either part. In this sense, we can qualify it as a *synchronic* example.

Goldbach's-conjecture_{synchronic}/Knowing Assessor. Didi, being asked whether there is a counterexample to Goldbach's conjecture, responds:

(1) There might be a counterexample to Goldbach's conjecture.

Naomi, who has (unknown to others) found a way of proving the conjecture and is eavesdropping, says with reference to Didi's statement (*sotto voce*):

(2) That's wrong (or: Didi is wrong).

And she adds:

(3) There cannot be any counterexample to Goldbach's conjecture.

On a face-value reading of Naomi's statement of (2), it disagrees with Didi's epistemic modal statement of (1). Furthermore (2) seems correct, in view of what she knows. The same point can be made regarding Naomi's epistemic *cannot* statement of (3). But also Didi's *might* statement of (1) seems correct,

²⁵ I do not claim that the given characterization of *disagreement* can be considered as a full-fledged account of the informal notion. But it seems unproblematic to take it at least as a necessary condition of *disagreement* in the informal sense. Cf. MacFarlane (2007).

²⁶ In earlier drafts of this paper, I used the label *faultless disagreement* for the type of scenarios invoked in relativist arguments. This label may be regarded as rather misleading. MacFarlane suggests that *faultlessness* primarily means merely that the speaker is not blameworthy for making a statement for which he has no good reason to believe that it is true (p.c.). I have replaced here *faultlessness* by the term *correctness*.

considering that he is ignorant as to the underlying proposition (i.e. Goldbach's conjecture).²⁷

Here is a *diachronic* counterpart-example. Speaker and assessor are here the same persons, but the epistemic situations at the utterance and the assessment time are essentially different.

Goldbach's-conjecture_{diachronic}/Knowing Assessor. It's Monday. Didi wonders whether there is a counterexample to Goldbach's conjecture. Naomi to Didi:

(4) There might be a counterexample to Goldbach's conjecture.

On Tuesday, Naomi finds a proof. Didi confronts her with her earlier statement. Naomi replies:

(5) That's wrong (or: I was wrong).

And she adds:

(6) There cannot be any counterexample to Goldbach's conjecture.

On a face-value reading of Naomi's statement of (5), it expresses retraction of her former epistemic *might* statement of (4)—which seems correct, in view of what she has found out meanwhile. But (4) seems correct as well, considering that she was ignorant as to the underlying proposition (i.e. Goldbach's conjecture).

Assuming that truth-values attach to epistemic modal statements absolutely, we cannot coherently keep all the assumptions (CD1–3). By (CD1) and (CD3), it follows that both statements are correct *as assessed against the epistemic situations of the respective speakers at the respective utterance times*. If utterance truth is absolute, however, the relativization of correctness to epistemic situations is semantically vacuous: if a statement is correct as assessed against some epistemic situation, so it is for any. But then both statements are correct, from whatever epistemic situation they are assessed—which contradicts (CD2), the thesis of disagreement. A relativism about epistemic modals, on the other hand, provides a straightforward way of accommodating *Knowing Assessor* scenarios without abandoning any of the prima facie plausible theses (CD1–3): *Correctness* (i.e. *truth*) of a statement is relative to assessment contexts. *Disagreement* is characterized as the case where for whatever assessment context, two statements are not jointly correct (true) relative to this assessment context.²⁸ As a result, there is no longer a tension between disagreement and correctness by standards of the speaker on either part. E.g. consider the *Knowing Assessor* scenario **Goldbach's-conjecture_{synchronic}**: Didi's and

²⁷ Perhaps Goldbach's conjecture is in fact false, and it is to be ruled out in view of what we already know about number theory—in which case the above scenario is metaphysically impossible. Of course, I could make the same point by adjusting the example appropriately, by letting Naomi know that the conjecture is false.

²⁸ Analogously for other types of discourse that have been invoked by relativists. For the general strategy, compare MacFarlane (2007).

Naomi's statements disagree, in the sense that whatever our epistemic situation is, we cannot assent correctly to both statements. But either statement is correct as assessed by the respective speaker, as the relevant epistemic situation is sensitive to *assessors*: Didi's *might* statement is true relative to his assessment, because the possibility that there is a counterexample to Goldbach's conjecture is not to be ruled out in view of what is known by *him*. Naomi's disagreeing statement, on the other hand, is true relative to her statement, because this possibility is to be ruled out in view of what is known by *her*. As for *Knowing Assessor* scenarios like **Goldbach's-conjecture_{diachronic}**, the suggestion is to resolve the puzzle analogously, by appeal to the sensitivity of the relevant epistemic situation to *assessment times*.

11.3 IGNORANT ASSESSORS

Standard relativism provides means of accommodating the possibility of correct disagreement scenarios which have the structure of *Knowing Assessor* scenarios. Let us turn to the flip-side of the coin. Standard relativism also predicts correct disagreement scenarios which are utterly bizarre. Specifically we can convert any *Knowing Assessor* scenario into such a bizarre scenario. In *Knowing Assessor* scenarios, *Ignorant* is the speaker, and *Knower* is the assessor. Suppose we put things upside down and let *Knower* and *Ignorant* swap their roles: i.e. *Knower* is now making an apodeictic epistemic modal statement, which is assessed by *Ignorant*. By hypothesis, a certain proposition which is epistemically possible for *Ignorant* is ruled out in view of what *Knower* knows. That is, by **Community/assessment-sensitive/egocentric**, an associated epistemic *cannot* statement of the form 'cannot: *P*', where *P* expresses the proposition concerned, will be correct relative to *Knower's* assessment context. Now *Knower's* epistemic situation is not relevant to *Ignorant's* assessment context—recall that for any assessment context, the relevant epistemic situation is the epistemic situation of the assessor at the assessment time. As a consequence, *Knower's* apodeictic statement will be false as assessed by *Ignorant*. Call the result an *Ignorant Assessor* scenario.

It is easy to illustrate the oddity of the *Ignorant Assessor* scenario-type by way of examples. Consider the following scenario that is obtainable from the *Knowing Assessor* scenario **Goldbach's-conjecture_{synchronic}**:

Goldbach's-conjecture_{synchronic}/Ignorant Assessor. Let the scenario be as in the *Knowing Assessor* scenario **Goldbach's-conjecture_{synchronic}**, and assess Naomi's apodeictic statement

- (7) There cannot be any counterexample to Goldbach's conjecture

against Didi's worse epistemic situation. On standard relativism, Naomi's statement of (7) is false as assessed by Didi. That is, it is granted that Didi would make a correct statement in saying

(8) That's false (or You are wrong)

with reference to Naomi's statement of (7).²⁹

The suggested relativist account of the scenario seems oddly strong. Indeed, if we tested Naomi's statement of (7) against an epistemic situation, where it is known that Goldbach's conjecture is false, or where it is known that Naomi just does not know better than us, then it would clearly fail the test for correctness. But imagine yourself in Didi's epistemic situation, where it is not known whether Naomi really has a proof, nor that there is no proof. In this situation, contrary to what standard relativism suggests, it is odd to reject Naomi's statement of (7) categorically. Rather it seems that what we know leaves it open whether Naomi's statement of (7) is true. Specifically, from Didi's perspective, it seems correct to say rather:

(9) I don't know whether Naomi is right. Perhaps she is. Perhaps she is not. It depends on whether she has really a proof.

Here is another example, this time it is obtained from the *Knowing Assessor* scenario **Goldbach's-conjecture_{diachronic}**:

Goldbach's-conjecture_{diachronic}/Ignorant Assessor. On the standard relativist account of the *Knowing Assessor* scenario **Goldbach's-conjecture_{diachronic}**, the relevant epistemic situation to her assessment context on Monday would be her epistemic situation *on Monday* (not on Tuesday). Hence Naomi's statement on Tuesday of the form

(10) There cannot be any counterexample to Goldbach's conjecture

is false as assessed by Naomi on Monday.

This result seems as odd as the result in the synchronic counterpart scenario, for analogous reasons. To spell out the oddity, suppose Naomi makes the following hypothetical retraction of her future statement on Tuesday:

(11) If I said tomorrow 'There cannot be any counterexample to Goldbach's conjecture', I would speak falsely.

²⁹ Consider a case where Didi eventually gets to know that Goldbach's conjecture is true, by trusting Naomi's testimony—in which case, we have no longer an *Ignorant Assessor scenario*. Even then, a relativist account leads to bizarre results. For example, once Didi knows that Goldbach's conjecture is true, he can correctly say with hindsight *Before I learnt from Naomi that there is no counterexample to Goldbach's conjecture, it would have been correct to say 'You are wrong' in reply to Naomi's statement 'There cannot be any counterexample to Goldbach's conjecture'*—which is bizarre.

Standard relativism suggests that this conditional statement is in any case semantically correct, as assessed by Naomi on Monday. But this account seems too strong. Rather it seems that what we know does not rule out the possibility that a future statement of the form ‘There cannot be any counterexample to Goldbach’s conjecture’ as made by Naomi would be true. It depends on whether it would accord with Naomi’s *then* epistemic situation, we would say.³⁰

Let us take stock of where we are now. Our point of departure were *Knowing Assessor* scenarios, which *prima facie* seem to be scenarios of correct disagreement over a given epistemic modal statement. Insofar as we want to accommodate this *prima facie* account of *Knowing Assessor* scenarios, a relativism about epistemic modals is indeed strongly suggested. For insofar as we do not leave a standard contextualist framework, there is no room for something like correct disagreement. On closer inspection though, a standard relativism has absurd consequences. It suggests that an ignorant assessor is in any case free to disqualify a given apodeictic epistemic statement as false, *even if* the speaker is better informed as to the proposition concerned.³¹ MacFarlane (forthcoming) himself acknowledges that the objection from *Ignorant Assessor* cases reveals a serious problem with standard relativism about epistemic modals. It seems hence fair to conclude that relativism is odd, at least in the letters in which it has been standardly spelled out. But it may still turn out a viable option, if considered in an alternative version. In what follows, I turn to an escape route that has been suggested by MacFarlane.

11.4 AN ESCAPE ROUTE?

Standard relativism has it that for any assessment context, only the epistemic situation of the assessor at the assessment time counts as relevant—even if the

³⁰ The same point can be made also for diachronic *Ignorant Assessor* scenarios with backward looking categorical assessments, where we can dispense with conditionals.

³¹ One may argue that there are also possible situations of correct disagreement over a given *apodeictic* epistemic modal statement: Suppose, for example, Didi has seen Brown in town just two minutes ago, and that he refers to normal standards for knowledge. As assessed by his standards, his statement of the form *Brown must be in town* seems then correct. Suppose furthermore that Naomi, by contrast, is in a state of sceptical neurosis. As assessed by her standards, her statement of the form *You are wrong. He might have been meanwhile kidnapped and flown out in a helicopter* seems also correct—though she disagrees with Didi. Or so one may argue. I take this point (thanks to Isidora Stojanovic here (p.c.)). But it needs to be noted that this type of scenario does not have the structure of an *Ignorant Assessor* scenario, where we refer to a fixed standard for knowledge. In the considered scenario, Naomi’s evidential situation is no worse than Didi’s evidential situation. The two subjects only disagree with respect to the norms for knowledge. As I noted (in § 1.5), potential assessment-sensitivity in norms for knowledge does not play any role in the examples invoked in the relativist literature, and there is good reason to ignore norms for knowledge in the context of epistemic modals.

statement concerned is made by another subject who is in a better epistemic situation as to the underlying proposition. As *Ignorant Assessor* scenarios show, this result is reducible ad absurdum. In reply to my objection, MacFarlane (forthcoming) suggests making the relevant epistemic situation function not only sensitive to assessment but also to utterance contexts. Specifically, the idea is to take into account both the epistemic situation of the assessor and the speaker as relevant to the assessment of epistemic modal statements. More formally, the suggested revision is to replace **Community/assessment-sensitive/egocentric** by:

Community/assessment+utterance-sensitive

For any context of utterance c and any context of assessment d : $C_{\langle c, d \rangle}$ takes as value the epistemic situation in $\{ \langle s_c, w_c, t_c \rangle, \langle a_d, w_d, t_d \rangle \}$, where for contexts x : s_x, a_x, w_x and t_x are the speaker, assessor, world and time at x respectively.

As for the epistemic situation associated with epistemic communities, we just need to assume the minimal constraint **Epistemic Situation** (compare Section 1.1). In effect, the relativist account schema **might/relativism** is revised by replacing mentioning of a one-place epistemic situation function by a two-place epistemic situation function.

It is easy to see that the resulting *contextualized relativism* is immune against the possibility of correct disagreement cases with the structure of *Ignorant Assessor*. Consider the example **Goldbach's-conjecture_{synchronic}/Ignorant Assessor**: For Didi's assessment of Naomi's utterance of

(3) There cannot be any counterexample to Goldbach's conjecture,

not only Didi's but also Naomi's knowledge is relevant. Hence also relative to Didi's assessment context, Naomi's utterance of (3) comes out true. And hence any utterance that disagrees with Naomi's utterance is not true relative to Didi's assessment context.

A contextualized relativism looks much less elegant than the original proposal, but it may be considered as a serious alternative, insofar as it does not give rise to other serious problems. So can the relativist lay back now? I contend that a contextualized relativism is objectionable in that it undermines the evidential basis for relativism. Here is an argument: On a contextualized relativism, the argument from correct disagreement seems to lose its intuitive force. Recall that the argument invokes an account of possible *Knowing Assessor* scenarios, on which there is correct disagreement over a given epistemic *might* statement. For example, as for **Goldbach's-conjecture_{synchronic}/Knowing Assessor**, it is suggested that Naomi's apparently correct statements of the form (3) and of the form

(2) That's wrong (or: Didi is wrong)

are *both* expressions of disagreement with Didi's utterance of

- (1) There might be a counterexample to Goldbach's conjecture.

On a contextualized relativism, however, this account of the scenario is no longer sustainable. Specifically, the epistemic *cannot* statement of (3) can no longer be evaluated as an expression of disagreement with Didi's *might* statement of (1). For there is now an assessment context—namely Didi's context—relative to which the statements of (3) and (1) are jointly true. If disagreement with a given *might* statement is no longer expressible by way of an associated apodeictic epistemic modal statement, the class of natural candidates for ordinary language evidence for correct disagreement seems to shrink to apparent predications of falsehood such as (2). To put the same point in another way, the account of statements like (2) as predications of falsehood seems no longer supported by other evidence for disagreement in the form of epistemic modal statements. Insofar as there are plausible alternative accounts of such phrases, as I will argue (in § 5), the evidential support for the thesis that there are cases of correct disagreement over given epistemic modal statements hence seems to fade away entirely.

A contextualized relativism may provide an effective strategy of immunizing relativism against the possibility of bizarre cases of correct disagreement. But by the same token, it seems to undermine also the intuitive motivation for adopting any relativist account about epistemic modals.

11.5 DISSOLVING THE PUZZLE

Recall how the puzzle of correct disagreement cases was generated: There are possible *Knowing Assessor* scenarios, i.e. scenarios for which it is *prima facie* plausible to give the following account: (CD1) *Ignorant* makes an epistemic *might* statement that meets the test for correctness as tested against her epistemic situation; (CD2) *Knower* makes a statement that disagrees with *Ignorant's* statement; (CD3) *Knower's* statement meets the test for correctness as tested against her epistemic situation. Starting from (CD1–3), a standard contextualist about epistemic modals indeed seems in deep trouble, for on a non-assessment-sensitive account of utterance truth for epistemic modal discourse, the assumption of correct disagreement is contradictory.

I do not deny that for some possible scenarios, a correct disagreement account has some intuitive force. I contend though that for any such scenario, a standard contextualist may give an alternative account that is no less plausible. I do not aim to provide a universal resolution strategy, which tells for any given *Knowing Assessor* scenario where the correct disagreement fails. Rather I suggest a case-by-case resolution, according to which the correct disagreement account fails on different grounds, depending on the particular scenario concerned. Let

me illustrate the idea by way of the two examples of *Knowing assessor* scenario. Take first the example **Goldbach's-conjecture_{synchronic}/Knowing Assessor**. Didi (i.e. *Ignorant*) states

(1) There might be a counterexample to Goldbach's conjecture.

Does he (a) use the epistemic *might* in an egocentric way? This would be the natural reading if Didi made further comments of the form 'I don't know, but Naomi surely knows'. Does he (b) mean to make a statement with respect to what is presently known in the scientific community? This would be the natural reading if Didi made further comments of the form 'At present, nobody knows a proof, nor a counterexample'. Or does he (c) mean to make a statement with respect to what is within the 'epistemic reach' (whatever that vague phrase means more specifically) of the scientific community? This would be the natural reading if Didi made further comments such as 'This problem is just too complex ever to be solved by a human mind'.³² There may be yet further plausible interpretations of Didi's statement, depending on what sort of intentions Didi is supposed to have in making his statement.³³

Now here is the crucial point. Whatever the speaker's intentions are like, it seems to be only fair to let them be the decisive standard for the correctness of epistemic modal statements. If a speaker means to make a statement with respect to a certain epistemic situation, then her statement is correct (in other words true) just in case it accords with that very epistemic situation. Since the question of what a speaker's intentions are is hardly assessment-sensitive, this strongly suggests that truth-values attach to epistemic modal statements *absolutely*. If in uttering (1), for example, Didi means to make a statement as to his individual state of knowledge, his statement is *absolutely* correct. Specifically, if Naomi rejects his statement on the ground that it does not reflect her epistemic situation, it seems fair to say that she does not do justice to Didi's statement.

For illustration of my case-by-case approach, consider the synchronic example: Suppose we understand Didi's statement in the (a)-sense. On that reading, we have to accept (CD1) for Didi's statement. But then (1) cannot fail the test for correctness as tested against Naomi's epistemic situation either—otherwise, Naomi would need to know that Didi knows in fact that there is no counterexample. That is, if Naomi's reply (2) is meant to be a predication of falsehood of Didi's statement of (1), (CD3) fails then for (2). If it is not meant to be a predication of falsehood of (1), (CD2) fails for (2). As for Naomi's epistemic

³² Hacking (1967) urges considerations for the conclusion that (at least sometimes) the epistemic *might* is used in the (c)-sense. For a standard contextualist account of epistemic modals in Hacking's spirit, see DeRose (1991).

³³ On non-egocentric uses of epistemic modals, intentions may be vague. That is, they may leave it for some epistemic subjects open as to whether they are members of the relevant epistemic community. For simplicity, I do not discuss here any potential vagueness of speaker intentions and any associated vagueness of epistemic modal statements.

modal reply (3), on the (a)-reading of Didi's statement, it in fact does not disagree with Didi's statement—for Naomi intends to make a statement as to an epistemic situation in which her knowledge is relevant, whereas Didi does not do so in his statement. Thus (CD2) fails for (3). We have hence reached a point where the puzzle can be dissolved on the (a)-reading of (1). Suppose we understand Didi's statement in the (b)-sense. On that reading, though Didi makes his statement in the best knowledge, it is incorrect, for someone (viz. Naomi) in fact knows that there is no counterexample. The respective instance of (CD1) for Didi's statement hence fails to hold. Again, the puzzle is dissolved. The same resolution strategy may be applied to the (c)-reading. As for the diachronic counterpart example **Goldbach's-conjecture_{diachronic}/Knowing Assessor**, we leave it at mentioning that no new points would need to be made for dissolving the puzzle. As argued, the correct disagreement account is challengeable in all its essential premisses, and depending on what further assumptions we make about the speakers' intentions, we have to target different premisses. It is hard to see that there are possible scenarios for which this case-by-case dissolution strategy may be blocked effectively.

11.6 OBJECTIONS AND REPLIES

It is time to consider some possible objections to my account:

11.6.1 Expressions of disapproval in the context of statements that accord with the epistemic situation intended by the speaker

Objection: It was suggested that Naomi's statement of (2) is incorrect if both (i) it is meant to be a predication of falsehood of Didi's statement and (ii) Didi's statement accords with the epistemic situation intended by him. This evaluation seems to clash with the intuition that we can correctly express disapproval of epistemic *might* statements, in case we know that the underlying proposition does not hold—even if we know that the speaker intends to make a statement as to an epistemic situation in which our knowledge does not count. Consider, for example, the egocentric (a)-reading of Didi's *might* statement of (1) in the synchronic Goldbach's-conjecture example. Imagine yourself in Naomi's epistemic situation. It seems that there is nothing incorrect about her use of expressions of disapproval of the form (2) with respect to Didi's statement—and this intuition is not undermined by the further information that Didi just intends to make a statement as to his own epistemic situation. For another example, consider the diachronic counterpart example and interpret Naomi's *might* statement on Monday, (4), as a statement as to the *then* state of research. Now imagine yourself in Naomi's epistemic situation on Tuesday. It would be correct to say (5) with respect to the earlier *might* statement, even if we were

reminded that it was meant to be a statement as to the *then* epistemic situation. The standard contextualist account of these examples seems inadequate, for it suggests that Naomi should dismiss her statements (2) (in the synchronic case) and (5) (in the diachronic case) as incorrect, if she learnt about the epistemic situations intended in the statements of (1) and (4) respectively.

Reply: The objection takes for granted that expressions of disapproval such as ‘That’s wrong’, as used with respect to a given epistemic modal statement, can only be analysed as *predications of falsehood of a statement*. I contend, by contrast, that ‘That’s wrong’-type expressions may be used in still other ways in the context of epistemic modal statements. Take the type of scenarios that the objection refers to. Why does Naomi’s statement of disapproval seem to be correct, even if in his *might* statement, Didi means to make just a statement as to his individual epistemic situation? If we interpret Naomi’s statement as targeting Didi’s claim of correctness for his *might* statement, then, as argued, Naomi’s statement is to be dismissed as incorrect. But we do not need to interpret it that way. Here is an alternative account: In making epistemic *might* statements, we do not only make an assertoric statement as to a relevant epistemic situation. We also express partial commitment to the underlying proposition. Analogously for apodeictic statements of the form ‘cannot: P ’: in making them, we do not only make a statement that qualifies a relevant epistemic situation, we also express full commitment to the proposition that not- P (or, in the case of ‘must: P ’, we express full commitment to the proposition that P). This suggests that expressions of disapproval made in the context of a given epistemic modal statement may refer to different speech acts associated with such statement, depending on the utterance context. They may be used for rejecting the claim of correctness for an assertoric statement, or what is asserted by the statement. They may alternatively (or additionally) be used for rejecting a proposition partial or full commitment to which is expressed by the same token.^{34,35} For example, in making his statement, Didi not only makes an assertoric statement, he also expresses partial commitment to the underlying proposition, viz. that there is a counterexample to Goldbach’s conjecture. If we interpret Naomi’s statement

³⁴ Thanks to Mark Sainsbury for drawing my attention to this idea (p.c.). Compare Gillies and von Fintel (2008a: § 3) and (2008b: § 5). MacFarlane (forthcoming) suggests that, regardless of what epistemic situation a speaker may have in mind in making a *might* statement, it is correct to assess the *whole* statement as incorrect in a context where we know the underlying proposition to be false. I do not think that this suggestion has much intuitive force.

³⁵ One may argue for yet further alternative accounts. Swanson (2005: § 2.5) interprets apparent truth-valuations such as ‘That’s false’ in the context of epistemic modal statements as expressing in fact evaluations of illocutionary speech acts associated with such statements. For example, Didi’s *might* statement may not serve only as an assertion but also a piece of advice not to overlook the possibility that Goldbach’s conjecture is false. According to this, Naomi’s disapproving reply may be just a correct comment on the appropriateness of Didi’s advice—since Naomi’s epistemic situation is better than Didi’s, it is only correct to refrain from following Didi’s advice. Or so one may argue.

of disapproval as a rejection of this underlying proposition, her statement *is* correct. For as Naomi knows this proposition to be false, she is in a position to reject it.

This *multiple-target* account of ‘That’s wrong’-type statements not only provides a non-relativist explanation of why they seem appropriate in the context of epistemic modal statements even if the latter are made in accordance with the intended epistemic situation. The multiple-target account also supplies sufficient means of accommodating the intuitive asymmetry between *Knowing Assessor* and *Ignorant Assessor* scenarios of apparent disagreement: If *Ignorant* makes a correct epistemic *might* statement, *Knower* can still correctly reject the proposition partial commitment which is expressed by *Ignorant’s might* statement, for *Knower* knows this proposition to be false (though he cannot correctly reject the claim of correctness for *Ignorant’s* statement). On the other hand, if *Knower* makes a correct *cannot* statement, *Ignorant* cannot correctly express disapproval with respect to *Knower’s* statement—whether her disapproval targets the whole assertoric statement (note: it is correct) or the proposition that Goldbach’s conjecture has no counterexample, full commitment to which is expressed by the same statement (note: by assumption, *Knower* knows this proposition to hold).³⁶

11.6.2 Disagreement with what is asserted by statements that accord with the epistemic situation intended by the speaker

Objection: The problem is deeper than suggested in the first objection. Naomi’s statement of disagreement (2) seems not only correct if it is interpreted as a comment on the underlying proposition that Goldbach’s conjecture is false, but also if we interpret it as a comment on Didi’s whole assertoric statement of (1). Imagine yourself in Naomi’s situation, where it is known that there is no counterexample to Goldbach’s conjecture. Hence it is true to say that there cannot be a counterexample to Goldbach’s conjecture. But Didi’s statement expresses the proposition that there might be one. Hence what is expressed by Didi’s statement is false. However, if what Didi’s statement expresses is false, how can the statement be considered as semantically correct after all?

Reply: The objection suggests that it makes no difference for what is expressed by *There might be no counterexample to Goldbach’s conjecture* whether the sentence is uttered by Didi or by Naomi—otherwise, we could not take it for granted that the content of the sentence as uttered in Didi’s context may be disquotationally

³⁶ It is easy to see that also Swanson’s alternative account of expressions of disapproval may supply sufficient means of accommodating the asymmetry between *Knowing* and *Ignorant Assessor* scenarios in non-relativist terms. Swanson in fact dismisses the factuality assumption for epistemic modality discourse. I do not pursue this line here any further.

expressible in Naomi's context by means of the same sentence. Speaking more generally, it is suggested that the contents of epistemic modal sentences are not sensitive to utterance contexts.³⁷ I am not convinced that intuitions uniformly support this view. But to the extent to which this view has intuitive force, it has to be taken seriously. Here is a way of accommodating it in purely contextualist terms: Recall that a contextualism about epistemic modals is compatible both with an indexical and a non-indexical view about the role relevant contextual parameters play in the fixation of the truth-value of epistemic modal sentences (cf. § 1.3). An *indexical* contextualist about epistemic modals could not make room for the intuition that Naomi correctly rejects what is expressed by Didi's statement—for according to the indexical contextualist account, Didi's utterance of (1) says something about his epistemic situation at the time of his utterance, namely that what he knows does not rule out the possibility that Goldbach's conjecture is false. But then what is expressed by Didi's statement is not rejectable in any context—for Naomi cannot correctly say *In view of what Didi knows, it is to be ruled out that Goldbach's conjecture has a counterexample* either. On the other hand, on a *non-indexical* contextualist account of epistemic modals, Naomi can correctly reject what is expressed by Didi's statement. For with respect to her utterance context, what is expressed by *There might be a counterexample to Goldbach's conjecture* is true. But what is expressed by the sentence relative to her context is the same as what is expressed by Didi's utterance of the sentence. Hence Naomi can correctly reject what is expressed by Didi's utterance. From this, however, it does not follow that Naomi can correctly reject Didi's utterance itself as incorrect. Whether Didi's utterance is correct depends on whether its propositional content (i.e. the unqualified proposition that it is epistemically possible that Goldbach's conjecture is false) holds with respect to the epistemic situation in Didi's utterance context—which, by hypothesis, it does. *Upshot*: A non-indexical contextualism supplies sufficient means of accommodating the view that there is room for correct disagreement over the truth-value of the propositional content of a given epistemic modal statement.³⁸ There is no need for a relativism about utterance truth.

11.6.3 Disagreement in *Ignorant Assessor* scenarios

Objection: Consider any *Ignorant Assessor* scenario, e.g. Didi's assessment of Naomi's utterance of (7), *There cannot be any counterexample to Goldbach's conjecture*. In Didi's epistemic situation, it is true to say *There might be a counterexample*. But this contradicts what is expressed by Naomi's utterance. Thus Didi can correctly say *That's false* with respect to what is expressed by

³⁷ I set aside here any potential indexicality of the embedded sentence.

³⁸ For further discussion of this point, with reference to future contingents, see MacFarlane (2008).

Naomi's statement. So where is the alleged asymmetry between *Knowing* and *Ignorant Assessor* scenarios?³⁹

Reply: Suppose, for the sake of argument, that (i) we interpret, as suggested in the objection, *Knowing* and *Ignorant Assessor* scenarios as cases where the assessor disagrees with what is asserted by a given epistemic modal statement. On the further supposition of (ii) an *indexical* contextualism about epistemic modals, then there is no asymmetry between both types of scenarios. Specifically, statements of disapproval of a correct epistemic modal statement are by no means correct: If Didi's *might* statement says something about his epistemic situation, Naomi cannot correctly disagree with what is said by his statement—analogously for the converse case that Didi disagrees with what is said by Naomi's correct *cannot* statement. On the other hand, if we replace (ii) by (iii) a *non-indexical* contextualism about epistemic modals, then again there is no asymmetry, but now in the converse sense that statements of disapproval of a correct epistemic modal statement are not only in *Knowing Assessor* scenarios correct (as argued in the foregoing reply) but also in *Ignorant Assessor* scenarios. But these considerations all hinge essentially on premise (i), which was challenged already in § 6.1. Here is another point: even if we start from (i) and (iii), the resulting account of *Knowing* and *Ignorant Assessor* scenarios provides no good reason for adopting a relativist view about epistemic modals. Even if there is room for correct disagreement over *the truth-value of the propositional content* of a given epistemic modal statement, it is odd on any account to say that there is room for correct disagreement over *the truth-value of a given epistemic modal statement*. The genuinely relativist claim, however, is that there is also evidence for the latter type of correct disagreement. Naomi may plausibly argue *There cannot be a counterexample. Didi says that there can be one. Hence what his statement says is wrong*. But the further conclusion *Hence Didi is wrong in making his statement* (or *He is making an incorrect statement*) would be on any account bizarre.

11.7 CONCLUSION

There are possible dialogue scenarios in which there is apparently correct disagreement over a given epistemic modal statement. If we take appearance for reality, we have to abandon the standard contextualist approach to epistemic modals, and a relativist approach is strongly suggested. But then there is a serious problem with immunizing a relativism against the possibility of bizarre

³⁹ Herman Cappelen (p.c.).

correct disagreement cases over epistemic modal statements without making also the alleged evidence for relativism disappear. Relativism about epistemic modals is odd in its standard version, and it seems hard to escape the problems of standard relativism without taking other serious problems on board. The move to relativism seems even more ill-motivated, considering that a standard contextualist framework supplies sufficient means of giving an alternative account of the alleged evidence for relativism. The suggested account is more elegant, since it avoids the relativist commitment to bizarre cases of correct disagreement without losing its intuitive motivation. The account is also more powerful, since it explains what makes bizarre cases of disagreement appear bizarre. Relativism about epistemic modals is of no use, and fortunately, it is also of no need. And it is *even* of no need, *if* we accept the view that *Ignorant Assessor* scenarios are no more bizarre than *Knowing Assessor* scenarios.

References

- Cappelen, Herman (2008). 'Content Relativism and Semantic Blindness', in this volume.
- Condoravdi, Cleo (2001). 'Temporal Interpretation of Modals: Modals for the Present and for the Past', in D. Beaver *et al.* (eds.), *Stanford Papers on Semantics*, Stanford: CSLI Publications.
- Cresswell, Max (1973). *Logics and Languages*, London: Methuen.
- DeRose, Keith (1991). 'Epistemic Possibilities', *Philosophical Review* 100: 581–605.
- Dietz, Richard (2005). 'Truth-Relativism and Faultless Self-Disagreement Over Time', MS, Arché, University of St Andrews.
- Egan, Andy (2005): 'Epistemic Modals, Relativism and Assertion', *Philosophical Studies* 133: 1–22.
- *et al.* (2005). 'Epistemic Modals in Context', in G. Preyer and P. Peters (eds.), *Contextualism in Philosophy*, Oxford: Oxford University Press, 131–69.
- Einheuser, Iris (2008). 'Three Forms of Truth-Relativism', in this volume.
- Gajewski, Jon *et al.* (eds.), (2005): *New Work on Modality* (= *MIT Working Papers in Linguistics*, v. 51).
- Gillies, Anthony S. and von Fintel, Kai (2004). "'Might" Made Right', Hand-out for Cognitive Science Colloquium Osnabrück, Dec. 16, 2004. URL <http://mit.edu/fintel/www/osnabruck.pdf>
- (2008a): 'CIA Leaks', *Philosophical Review* 117: 77–98.
- (2008b): 'An Opinionated Guide to Epistemic Modality', in T. Szabó Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology* 2, Oxford: Clarendon Press, 32–63.
- Glanzberg, Michael (2007). 'Context, Content, and Relativism', *Philosophical Studies* 136: 1–29.
- Hacking, Ian (1967): 'Possibility', *Philosophical Review* 76: 143–68.

- Hawthorne, John (2004). *Knowledge and Lotteries*, Oxford: Clarendon Press.
- Kaplan, David (1989). 'Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals', in J. Almog *et al.* (eds.), *Themes from Kaplan*, New York: Oxford University Press, 481–566.
- Kratzer, Angelika (1977). 'What *Must* and *Can* Must and Can Mean', *Linguistics and Philosophy* 1: 337–55.
- Lasersohn, Peter (2005). 'Context Dependence, Disagreement, and Predicates of Personal Taste', *Linguistics and Philosophy* 28: 643–86.
- Lewis, David (1980). 'Index, Context, and Content', in S. Kanger and S. Öhman (eds.), *Philosophy and Grammar*, Dordrecht: Reidel.
- MacFarlane, John (2003). 'Epistemic Modalities and Relative Truth', draft of Nov. 13, 2006, University of Berkeley. URL <http://johnmacfarlane.net/epistmod-2003.pdf>
- (2005a). 'The Assessment Sensitivity of Knowledge Attributions', *Oxford Studies in Epistemology*, vol. 1: 197–233.
- (2005b). 'Semantic Minimalism and Nonindexical Contextualism', in G. Preyer and G. Peter (eds.), *Content and Context: Essays on Semantics and Pragmatics*, Oxford: Oxford University Press.
- (2007). 'Relativism and Disagreement', *Philosophical Studies* 132: 1–16.
- (2008). 'Truth in the Garden of Forking Paths', in this volume.
- (forthcoming). 'Epistemic Modals are Assessment-Sensitive', draft of Aug. 23, 2006, University of Berkeley, forthcoming in an OUP volume on epistemic modals, ed. by B. Weatherston and A. Egan. URL <http://johnmacfarlane.net/epistmod.pdf>
- Price, Huw (1983). 'Does "Probably" Modify Sense?', *Australasian Journal of Philosophy* 61: 396–408.
- Stephenson, Tamina (2005). 'Assessor Sensitivity: Epistemic Modals and Predicates of Personal Taste', in J. Gajewski *et al.*, 179–206.
- Swanson, Eric (2005). 'Something "Might" Might Mean', draft of Dec. 14, 2005, MIT.
- Teller, Paul (1972). 'Epistemic Possibility', *Philosophia* 2: 303–20.
- Veltman, Frank (1996). 'Defaults in Update Semantics', *Journal of Philosophical Logic* 25, 221–61.
- von Fintel, Kai (2005). 'Modality and Language', in D. M. Borchert (ed.), *Encyclopedia of Philosophy*, 2nd. edn., Detroit: MacMillan Reference. URL <http://mit.edu/fintel/www/modality.pdf>
- and Iatridou, Sabine (2002). 'The Meanings of Epistemic Modality', MS, MIT. URL <http://web.mit.edu/fintel/www/konstanz-ho.pdf>
- (2003). 'Epistemic Containment', *Linguistic Inquiry* 34: 173–98.
- Wright, Crispin (2007). 'New Age Relativism and Epistemic Possibility', *Philosophical Issues* 17: 262–83.
- Yablo, Stephen (2006). 'That May Be', MS, MIT, URL <http://www.mit.edu/~yablo/thatmaybe.pdf>
- Yalcin, Seth (2005). 'Epistemic Modals', *Mind* 116 (2007): 988–1026.

PART IV

ALTERNATIVES TO RELATIVISM

This page intentionally left blank

Content Relativism and Semantic Blindness

Herman Cappelen

For some relativists some of the time the evidence for their view is a puzzling data pattern. On the one hand, there's evidence that the terms in question exhibit some kind of content stability across contexts. On the other hand, there's evidence that their contents vary from one context of use to another. The challenge is to reconcile these two sets of data. Truth-relativists claim that their theory can do so better than contextualism and invariantism. Truth-relativists, in effect, use an argument to the best explanation: they present data that they claim to be able to handle better than any competing theory.¹

My interest is in how semanticists should react to this allegedly puzzling data pattern. I argue that what generates the appearance of a puzzle is a mistaken assumption about the relationship between semantic content and speech act content (i.e. the relationship between semantic content and what speakers assert, say, and claim²). When this mistaken assumption is corrected for, *any* semantics can deal with this data pattern. It doesn't cut either way with respect to the debate between the contextualist, the invariantist and the truth-relativist.

I show this by first presenting what I take to be data (some might want to call it a theory) about speech act content and semantic content. I call this collection of data *Pluralistic Content Relativism* (PCR). I show that when PCR is added to contextualism *or* invariantism, those theories can easily deal with the (allegedly) puzzling data pattern relativists use to motivate their theories. Not only can they deal with it—PCR predicts it. PCR is, moreover, independently

Earlier versions of this paper were given as talks at the Universities of Oslo, Rutgers, and Birmingham. Thanks to the audience for helpful comments on these occasions. Two presentations in the Arché Relativism seminar were immensely helpful. I'm particularly grateful to Jessica Brown, John Hawthorne, Jeff King, Ernie Lepore, Elia Zardini, Jason Stanley, and Crispin Wright for conversations about these issues. Ernie Lepore, Max Kölbel and John MacFarlane all provided extensive and extremely useful written comments.

¹ This is obviously not the *only* argument that can be or has been used in favor of relativism. The claim is only that this is the central argument in the papers I discuss (by John MacFarlane and Peter Lasersohn).

² There are a number of terms in the family of 'say' and 'assert'. What follows is supposed to apply to the whole range of such terms.

motivated, i.e., it is not a view introduced specifically to reply to the objections from truth-relativists.

Warning: those looking for an argument in favor of a particular semantic theory as an alternative to relativism will be disappointed. This paper is about methodology: it is about how *not* to do semantics. The kind of data truth relativists focus on will *not* help you adjudicate between competing semantic theories. The only positive theory defended in this paper is PCR—but it is not a semantic theory and doesn't tell you anything about the semantics of any particular set of terms.

My focus is on the work of John MacFarlane and Peter Lasersohn,³ not because their arguments are particularly guilty of the kind of mistake I will outline. On the contrary, their presentations are framed in a way I find particularly useful. They present the data more or less exactly as I see it⁴ and the clarity of their presentations makes it particularly easy to show the relevance of PCR.

The paper has two parts. In Part One, I present PCR and sketch the kinds of argument used by Cappelen and Lepore (C&L, for short, 1997, 1998, 2001, 2003, 2004, 2006) and Scott Soames (2002, 2005, forthcoming) in support of it. In Part Two I use PCR to evaluate four arguments for truth-relativism (three from MacFarlane and one from Lasersohn).

Part I: Pluralistic Content Relativism (PCR)

PCR can be presented as the conjunction of three theses: P1, P1.1, and P2.

P1. An utterance *u* of a sentence *S* in a context *C* will (literally⁵) assert (and *say* and *claim*) a plurality of propositions.⁶

I call this thesis *Pluralism*.

P2. What's said by an utterance *u* of *S* in a context of utterance *C* varies across contexts of interpretation.⁷

I call this thesis *Content Relativism*.

Finally:

P1.1. At most one of the propositions asserted by an utterance *u* of a sentence *S* relative to a context of interpretation *CI* is the proposition semantically expressed by *u*.

I discuss these in turn.

³ For related work see Egan, Hawthorne, and Weatherson (2005), Egan (forthcoming), Kölbel (2002).

⁴ For a related presentation of this kind of data, see Cappelen and Lepore (2006).

⁵ As opposed to *roughly*, or *metaphorically*, or *indirectly*, or *almost*, or any such qualifier. For a discussion of this point see C&L (1997). I'll leave out this qualifier from now on.

⁶ If you prefer to talk about speakers making assertions by uttering sentences rather than utterances making assertions, that's fine. Nothing in what follows will depend on the choice of terminology here.

⁷ A *context of interpretation* is just what you would think it is: a context from which an utterance is interpreted. An utterance *u* of *S* in *C* can be interpreted from infinitely many contexts of interpretation.

12.1 PLURALISM (P1)

Jeff King in conversation pointed out to me that this paper does no more than defend a conditional: *if* PCR is correct, *then* one line of argument in favor of truth-relativism fails. I'm not trying to defend the antecedent here. I might have written too much about that already, so I do no more than refer readers to earlier work and sketch the kind of argument that can be used to establish P1. That of course makes the exercise somewhat less interesting if you are not convinced by PCR (and P1 in particular). However, even if you're unconvinced by the antecedent, there might be some interest in seeing whether the conditional can be defended: if you really don't like relativism about truth and I can convince you that if PCR is true, it can account for the data advanced in support of truth-relativism, this might make you more sympathetic towards PCR.

I think of P1 as a characterization of our considered judgments about the nature of assertion as these are reflected in our practice of indirectly reporting each other. Here's a brief sketch of the kind of data that used as evidence for P1:⁸

Illustration #1: The Dresser (C&L (1997)): Imagine an utterance of (1), by A.

1. At around 11 p.m., I put on a white shirt, a blue suit, dark socks, and my brown Bruno Magli shoes, I then got into a waiting limousine and drove off into heavy traffic to the airport, where I just made my midnight flight to Chicago.

According to Pluralism, (1.1–1.7) are all true descriptions of what was said by an utterance of (1) (Note: all of these are obviously different propositions, and determine different truth conditions):

- 1.1. A said that he put on a white shirt.
- 1.2. A said that he put on shoes.
- 1.3. A said that he was wearing shoes when he went to the airport.
- 1.4. A said that he was stuck in traffic on his way to the airport.
- 1.5. A said that he that dressed around 11 p.m., went to the airport and took the midnight flight to Chicago.
- 1.6. A said that he that dressed before he went to the airport.
- 1.7. A said that he put on some really fancy shoes before he went to the airport.

The extent to which (1.1–1.7) seem natural will depend on the circumstances of the report; so arguments for Pluralism are usually accompanied by vignettes that

⁸ For many other examples of this kind and more elaborate discussion of them, see C&L 1997, 1998, 2004 and Soames 2002, 2003, and (forthcoming).

describe the context for the report. Having argued that these reports (in those contexts) are literally true (and not just appropriate or warranted), Pluralists conclude that in uttering (1), A said the complement clauses of (1.1–1.7). That's only a tiny sample of what was said in uttering (1).

Illustration #2: The Terrorist (from Soames (2002)):

A terrorist has planted a small nuclear device in a crowded stadium downtown. There is no time to evacuate the building or the surrounding area. In speaking to the negotiator, he utters (2)

(2) 'I will detonate the bomb if my demands are not met,'

knowing that it is obvious that if he does so, thousands of people will die, and intending to communicate precisely that. The negotiator reports to his superior that the terrorist said that he will kill thousands of people if his demands are not met.

All of (2.1.–2.4) are true:

- 2.1. He says that he will kill thousands of people if his demands are not met.
- 2.2. He says that he will detonate the bomb if his demands are not met.
- 2.3. He says that he will create mayhem downtown if his demands are not met.
- 2.4. He says that he will inflict great damage on our community if we don't do as he says.

Cappelen and Lepore (1997) summarize the lessons from such examples as follows:

...indirect reports are sensitive to innumerable non-semantic features of reported utterances and even on the context of the report itself. As a result, typically there will be indefinitely many correct indirect reports of any particular utterance. (p. 291)

Soames (2002) draws a related conclusion:

[The phenomenon of many propositions being expressed by an utterance of a sentence]... is an extremely general one that has nothing special to do with proper names, indexicals or any of the semantically contentious issues that are of special concern here. On the contrary, the phenomenon of asserting more than the semantic content of the sentence one utters in a context is all but ubiquitous... what an assertive utterance of a sentence *s* counts as asserting depends not only on the semantic content of *s*, but also on the obvious background assumptions in the conversation and the speaker's intention about how the speaker's remark is to be interpreted in the light of them. (pp. 77–9)

These are the kind of arguments I take as evidence of P1. A more convincing case for P1 would require much more data and responses to alternative interpretations of it. For some efforts in this direction see references in note 8 above.

12.2 PLURALISM AND SEMANTIC CONTENT (P1.1)

In what follows I assume that an utterance of a sentence, *S* in a context of utterance *C* has at most one proposition as its semantic content relative to a context of interpretation *CI*.⁹ If so, at most one of the propositions expressed by an utterance *u* of *S* relative to *CI* can be its semantic content, i.e. P1.1 (repeated here) is true:

P1.1. At most one of the propositions asserted by an utterance *u* of a sentence *S* relative to a context of interpretation *CI* is the proposition semantically expressed by *S*.

We could introduce a notion of semantic content that allows an utterance to have more than one semantic content relative to a context of interpretation, but that's not an option I will explore here.

P1.1 leaves us with two options for how to think about the relationship between asserted content and semantic content. To describe these (and much else in this paper), I have to talk about sentences, utterances, contexts of utterance, and contexts of interpretation. To make this talk more perspicuous, I use the following simplifying device: '*u_{s/c}*' denotes an utterance *u* of sentence *S* in context of utterance *C*. The two versions of P1.1 are P1.1.1 and P1.1.2:

P1.1.1. The semantic content of *u_{s/c}* is among the propositions asserted relative to every context of interpretation.

P1.1.2. The semantic content of *u_{s/c}* is not among the propositions asserted relative to every context of interpretation.

The choice between P1.1.1 and P1.1.2 depends, in part, on how one thinks about semantic content. For someone who thinks that the semantic content of *u_{s/c}* relative to *CI* is always a proposition, P1.1.1 is attractive: it is hard to see how, if speakers know the meaning of their words and those meanings determine that *u_{s/c}* expresses *p* and the speaker knows that her audience knows this (etc.), she isn't asserting *p* (among other propositions she is asserting) relative to every context of interpretation. If, on the other hand, you think there are cases where the semantic content is some kind of sub-propositional object, e.g. a propositional skeleton or some such thing,¹⁰ P1.1.2 is the natural position to endorse: if the semantic content isn't a proposition, it isn't assertable. If so, it's not assertable relative to any context of interpretation. Other issues can affect the choice between P1.1.1 and P1.1.2, too.

In what follows we'll not choose between P1.1.1 and P1.1.2—in earlier work I have moved back and forth between them; C&L (1997) defends P1.1.1 and

⁹ I say 'at most' because I want to keep open the possibility that semantic contents of sentences (relative to contexts of utterance) are non-propositional objects.

¹⁰ See Bach (1994), Soames (forthcoming) and Cappelen (forthcoming) for a discussion of skeletons.

C&L (2004) is agnostic. Soames (2002) defends P1.1.1, but moves on to P1.1.2 in Soames (2005) and (forthcoming). More generally, this paper is neutral on how pluralists should think about semantic content in specific cases. The goal in this paper is to focus on the implications of PCR for certain kinds of arguments (the arguments adduced in favor of relativism), not to push a particular theory of semantic content.

12.3 CONTENT RELATIVISM (P2)

Two versions of P2 will be discussed in what follows. The first follows from P1 combined with some minimal additional assumptions about human psychology. Recall that, according to P1, each utterance u of a sentence S expresses a set of propositions R . The following seems obvious: for an interpreter, some members of R will be more salient than others. Suppose all of (1.1–1.7) are said by the utterance of (1). It is still unlikely that all of these are equally salient to all interpreters (including the speaker). Consider (2.4) as a description of what is said by an utterance of (2). An interpreter who thinks the terrorist is helping the community by blowing up an ugly building (without causing any casualties), might not find (2.4) to be a salient component of what was said. In this regard, P2.1 follows almost immediately from Pluralism (given mundane facts about salience):

P2.1: What's *saliently* asserted by $u_{s/c}$ varies across contexts of interpretation.

P2.1 is a *weak* version of content relativism. According to P2.1, what is *saliently said* by $u_{s/c}$ will vary across contexts of interpretation. The picture is this: A utters S in C . In so doing she expresses a range of propositions $p_1 \dots p_n$. Contingent on an interpreter's interests and beliefs, one (or more) of these propositions will be salient. Which one this is will vary between contexts of interpretation.

According to P2.1 the propositions asserted by $u_{s/c}$ are the same relative to all contexts of interpretation. According to the second version of P2 the set of asserted propositions varies across contexts of interpretation:

P2.2: The set of propositions asserted by $u_{s/c}$ varies across contexts of interpretation.

Think of this as *strong content relativism*. There are various ways one might defend this stronger version of content relativism. Here are two:

- (a) If the locution 'A said that p ' is context sensitive, i.e., if it is true to utter "A said that p " in C , but not in C' , there's a short step to P2.2 (if we assume P1, as I will in what follows). The context sensitivity of 'said that' leads to P2.2 once we add a couple of innocuous assumptions: if what speakers say is closely related to true indirect reports of what they say (and how

could it not be?) and if the latter varies across contexts of interpretation, then what speakers say varies across contexts of interpretation as well, i.e. P2.2 is true.^{11,12}

- (b) Alternatively, we could just assume P2.2 and see what work that assumption will do for us in theorizing about language and communication. If that assumption turns out to help us account for a great deal of data, we have a reason for believing it to be true.

In what follows I rely on the weaker version of content relativism, i.e. P2.1, not P2.2. Though I think P2.2 is true, it is harder to establish than P2.1. Since the arguments below only require P2.1, there's no need to rely on the stronger principle. (For one possible exception to this, see note 28.)

There's a third version of content relativism, which we can call *monistic content relativism*. This is P2.2 and P1.1 without P1, i.e. content relativism without the pluralism. On this view, there's one proposition asserted by any utterance of a sentence, that proposition is its semantic content, and that content can vary across contexts of interpretation. Monistic content relativism will not play a role in any of what follows, but if you endorse monistic content relativism, the line of argument here used against truth relativism could be taken over almost directly. All you need to do is change the terminology a little.

Three final points about PCR, before moving on to the discussion of arguments for truth relativism:

- PCR implies that speakers don't have first-person authority over what they say. By uttering a sentence, they might end up saying things they are not aware of having said—they might even end up saying things they would deny

¹¹ Cappelen and Lepore (2001) say:

In effect, our practice of reporting others treats what is said as a four-place relation between a sentence and its context of utterance and a reporting sentence and its context of utterance. In determining what is said we obviously draw upon information about specific intentions, knowledge, and history of the speaker in C and, not so obviously, we can also draw upon like features of C*, the context in which we report what is said. (p. 43)

¹² Soames (2002) considers, but does not endorse, the following view:

Roughly speaking, *X* says that *S* is true relative to a context C iff (i) the referent of *X* assertively uttered a sentence *S** in a context C*, and (ii) the proposition semantically expressed by *S* in C is something reasonably intended by *X* to be a potentially direct, immediate, and relevant inference (on the part of the conversational participants in C*) from *S** together with the background assumptions presupposed in C*. Suppose further that what counts as direct, immediate, and relevant is inferable from the speaker's utterance plus background information is something that varies to some extent from context to context; . . . If this view is correct, then there should be cases in which the proposition expressed by *Soames said that S* in C is compatible with the proposition expressed by *Soames didn't say that S* in C'. In a case like this, what at first might seem to be a factual dispute turns out to be nothing more than the adoption of different discretionary standards regarding how close a proposition must be to the semantic content of the sentence uttered by the agent in order to count as something the agent said. (Soames (2002), p. 336 n. 24)

having said. For elaboration on this possibility, see C&L (2004), chapter 13 (some such cases will also arise in the discussion below).

I doubt that what I have said so far will convince someone not already convinced of PCR, but one methodological point should be clear: it is not an *a priori* truth that a speaker who utters a sentence asserts only one proposition and that this one proposition is asserted relative to all context of interpretation. It is an empirical claim about assertion and it is itself in need of justification. Given the data presented briefly above, my sense is that the burden of proof here is on someone who denies PCR. (For some attempts, see Richard (1998), Reimer (1998), and the replies in C&L (1998).)

- According to PCR, many propositions are asserted by an utterance *u* relative to a context of interpretation. So far I have said nothing about *how* this content is generated. I have not presented a theory of speech act content. I also haven't given you a theory of the relationship between the semantic content and the rest of the speech act content. I have not said anything about *how* a particular part of the speech act content becomes salient in a context of interpretation. For reasons I've given elsewhere (see C&L (2004), chapter 13), I'm not convinced this is the kind of thing that lends itself to a theory, but nothing depends on that assumption in what follows.¹³

Part II: PCR and Four Arguments For Relativism

In the rest of this paper I look at some arguments for truth-relativism and evaluate those arguments on the assumption that some version of PCR is true. I argue that if PCR is correct, these arguments fail. More generally: the *kinds* of argument used in support of relativism falter if PCR is correct. The arguments (i.e. the data patterns appealed to) provide no evidence for truth-relativism over PCR. Since PCR is compatible with both contextualism and invariantism, no evidence has been provided that favors truth-relativism over any of these alternatives.

¹³ It is worth mentioning three additional features of PCR that would need more extensive discussion if my goal were a full-fledged presentation of the view:

1. *Local vs. Global PCR*: PCR can apply either globally or locally. A global version applies to all utterances of all sentences. A local version is less ambitious; PCR would have to justify why some sentences (or utterances of sentences) are not subject to the kind of relativism that other sentences (utterances) are.
2. *PCR and Negotiation*: A speech act pluralist need not say that for every proposition *p* it is settled whether $u_{s/c}$ said that *p* relative to context of interpretation *CI*. A possible view is that it can be indeterminate whether $u_{s/c}$ said that *p* relative to *CI*. It might, for example, be up for negotiation in *CI* whether $u_{s/c}$ said that *p* relative to *CI*.
3. *PCR and Error*: Interpreters could be wrong about what's said by $u_{s/c}$ relative to *CI*. It is certainly no part of PCR that interpreters are infallible. A lot of the explanations provided in the text below could be presented as cases where interpreters make mistakes about speech act content. To keep the discussion simple, I don't pursue that strategy, but a fuller presentation of PCR might rely heavily on it.

12.4 ARGUMENT #1: MACFARLANE'S 'BIG PICTURE ARGUMENT' FOR RELATIVISM ABOUT KNOWLEDGE

John MacFarlane has argued for a relativistic semantics for a number of terms and here I will focus on his argument for relativism about 'know' (in MacFarlane 2005). The kind of data MacFarlane focuses on is a peculiar pattern of content stability and variability between contexts. On the one hand, there's strong evidence that sentences containing 'know' exhibit variability in content across contexts. This provides evidence for contextualism about 'know', i.e. the view that its semantic value shifts between contexts of utterance. On the other hand, these sentences exhibit certain kinds of content stability across contexts. This provides evidence that 'know' is semantically stable. Here is MacFarlane's description of the variability data:

D1—Variability Data: Normally, I am perfectly happy to say that I know that my car is parked in my driveway. I will say this even when I'm at work, several miles away. But if someone asks me how I know that my car has not been stolen (and driven away), I will admit that I do not know this. And then I will have to concede that I do not know that my car is in my driveway: after all, if I knew this, then I would be able to deduce, and so come to know, that it has not been stolen. (MacFarlane (2005), p. 200)

Here is MacFarlane's description of the stability data:¹⁴

D2—Stability Data: When standards have been raised, I will say not only that I don't know that my car is in my driveway, and that I didn't know this earlier, but that my earlier assertion of "I know that my car is in the driveway" was false. In part, this is because we tend to report knowledge claims homophonically, even when they were made in very different epistemic contexts. . . . I won't just say that it was false; I will treat it as false. If challenged, I will retract my earlier claim, rather than reformulating it in a way that shows it to be consistent with my current claim. (pp. 202–3)

According to MacFarlane, no *standard* semantics for 'know' can account for all this data. About the invariantist he says:

The apparent variability of standards suggests that the truth of sentences containing "know" depends somehow on varying epistemic standards. That would rule out strict invariantism. (p. 204)

About the contextualist, he says:

. . . the facts about truth ascriptions and retraction suggest that the semantics of "know" is use-invariant. That would rule out contextualism. (p. 204)

¹⁴ MacFarlane also discusses data about how 'know' behaves when embedded within the scope of temporal and modal operators. Since this is used primarily in the discussion of 'subject sensitive invariantism', a discussion I ignore for the purposes of this paper, I leave that out. See C&L (2006) for some further discussion of that view.

The solution, according to MacFarlane, is truth-relativism:

How could there be a semantics for “know” that was use-invariant and circumstance-invariant, but still in some way sensitive to changing epistemic standards? . . . Here is my proposal. The epistemic standards relevant to determining the extension of “know” are not those in play at the context of use or those in play at the circumstance of evaluation, but those in play at the *context of assessment*. (p. 207)

On this view, the proposition asserted by an utterance in a context can vary between contexts of assessment:

A sentence is context-sensitive in the usual way, or *use-sensitive*, if its truth value varies with the context of use (keeping the context of assessment fixed). A sentence is context-sensitive in the new way, or *assessment-sensitive*, if its truth value varies with the context of assessment (keeping the context of use fixed). (p. 207)

The view explored in this paper, PCR, is an alternative to truth-relativism that’s compatible with both invariantism and contextualism. According to PCR, what varies between contexts of assessment (what I call contexts of interpretation) is not truth-value, but asserted content. I’ll claim several advantages for this view over truth-relativism: it is independently motivated and given the overwhelming independent data for PCR, it requires no radical rethinking of the nature of truth or content. In particular, it avoids the extreme view that a proposition can be true for A and false for B.

According to PCR, invariantism is compatible with the kind of variability described in D1 and contextualism is compatible with the kind of stability described in D2. To see why, first notice that both (i) and (ii) are predicted by PCR:

- i. *How utterances with the same semantic content can make different assertions.* If PCR is true, two utterances, $u_{s/c}$ and $u'_{s/c'}$ (where $s = \text{‘A knows that } p \text{ (at } t\text{)}\text{’}$) could have the same semantic content, and still (saliently) say different things relative to a context of interpretation CI. The propositions (saliently) asserted by u and u' relative to CI could include reference to different epistemic standards, even though their semantic contents are identical. So you could have shared semantic content, combined with the intuition that $u_{s/c}$ and $u'_{s/c'}$ made different assertions relative to CI.
- ii. *How utterances with different semantic contents can make the same assertion.* If PCR is true, two utterances, $u_{s/c}$ and $u'_{s/c'}$ (where $s = \text{‘A knows that } p\text{’}$) could have *different semantic contents*, even though *what they say is the same* relative to a context of interpretation CI. Remember, in addition to their semantic contents, $u_{s/c}$ and $u'_{s/c'}$ assert a plurality of propositions relative to CI. So, assume P is the semantic content of $u_{s/c}$ and P' the semantic content of $u'_{s/c'}$. Relative to CI, $u_{s/c}$ might also say P_2 , P_3 , and P_4 and $u'_{s/c'}$ might say P_4 , P_5 and P_6 (in addition to P'). If so, they both assert that P_4 (relative to CI). If P_4 happens to be salient in CI, an interpreter will have the intuition that u

and *u'* asserted the same proposition (i.e. that the speaker of *u* said what the speaker of *u'* said.)

With (i) on the table, variability of the kind MacFarlane focuses on doesn't count against invariantism. Semantic invariantism is compatible with variability in (saliently) asserted content (and variability in what the speaker (saliently) said). The observation that there is such variability is neither here nor there with respect to the semantics of 'know'.

With (ii) in play, the kind of stability described in D2 provides no argument against contextualism. There can be stability in some part of the speech act content even though the semantic content varies with the contextually salient epistemic standards in the context of utterance (for a less schematic development of this possibility see Section 6 below).¹⁵

This is the big-picture reply to MacFarlane's 'big picture' argument against contextualism and invariantism. In the body of the paper he considers various ways in which invariantists and contextualists can try to account for the recalcitrant data. I now look at two of those proposals and show how adding PCR to invariantism and contextualism provides replies to MacFarlane's objections.

12.5 ARGUMENT #2: MACFARLANE AGAINST THE SKEPTIC

The skeptic is someone who thinks 'know' is semantically invariant and that its semantic value is such that it is very hard, if not impossible, to stand in the knowledge relation to a proposition. MacFarlane describes it as the view that 'know' is *invariant with fixed high standards*, i.e., just one standard is relevant, in all contexts of utterance and that standard is so high that it is (almost) impossible to stand in the knowledge relation to any proposition.

As we have seen, the alleged problem for such a position is that it can't account for the variability data (presented as D1 above). MacFarlane considers four ways in which a skeptic can try to explain why we are inclined to make positive knowledge attributions (and treat them as true) when we are in low-standard contexts. The first three presuppose that the speaker knows that the skeptic is right about the semantics for 'know' and the fourth assumes that the speaker is mistaken in that respect.

1. *Helpfulness*. The enlightened skeptic is trying to avoid misleading those who are not aware of the fact that skepticism is true, i.e., who are not aware that the

¹⁵ As Jessica Brown emphasized to me, those who rely on (i) in explaining recalcitrant data will depend on P1. Those who rely on (ii) will depend on P2. So if you find the latter less plausible than the former, you would find the use of PCR in (i) more convincing than its use in (ii). Since the semantic invariantists rely on (i) and the contextualists on (ii), this would provide reason for favoring semantic invariantism over contextualism.

standards for knowledge are high (and hence would misinterpret utterances were she to speak as if skepticism is true).

2. *Hyperbole*. The enlightened skeptic says more than what she has good reason to assert; it is a form of hyperbole.
3. *Inconvenience*. MacFarlane says, “A third approach would appeal to the inconvenience of adding all the pedantic hedges and qualifications that would be needed to make our ordinary knowledge claims strictly true.” (MacFarlane (2005), p. 207)
4. *Error*. MacFarlane also considers the possibility that speakers are wrong about the semantics for ‘know’, i.e., speakers don’t know that skepticism correctly describes the semantics for ‘know’. He says: “A sincere speaker who wants to speak the literal truth and avoid literal falsity may fail to do so if she has false beliefs, either about the facts or about the literal meanings of the words she usesSceptical invariantists . . . must argue that speakers systematically overestimate their success in meeting the standards for knowledge and as a result claim to know when in fact they do not.” (pp. 210–11)

MacFarlane thinks all these moves fail. I won’t discuss his counter-arguments here. I’m interested in possibilities left out—possibilities opened up by adding PCR to the picture. If the skeptic takes on PCR, she can argue as follows: speakers utter sentences of the form ‘A knows that p’ because in so doing they manage to (saliently) assert (and *say* and *claim*) a true proposition.^{16,17} MacFarlane simply assumes that Skeptical Invariantism is incompatible with the speaker literally asserting some proposition that’s true when she utters ‘A knows that p’. He assumes this because he also assumes that speakers in uttering a sentence assert only one proposition, the semantic content. That is to say, MacFarlane excludes the possibility that in addition to asserting the semantic content (which is false, according to the skeptic) the speaker asserts a true proposition to the effect that *A knows that p by standard S*, where S is some non-skeptical epistemic standard (one A can stand in with respect to a proposition without ruling out bizarre skeptical hypothesis). PCR provides the skeptical invariants with a reply to the argument from variability. In response to the question, “If the skeptic were right, why would someone assert ‘A knows that p’ given that its semantic content is always false?” she can answer: “Because in so doing, she saliently asserts something true.”¹⁸

¹⁶ For an elaboration of this defense of skepticism, see Cappelen (2005).

¹⁷ For a ‘PCRist’, all talk about speech act content should be relativized to contexts of interpretation. To simplify writing, I will exclude the explicit relativization when it’s obvious or superfluous.

¹⁸ Note that there need be no error involved here. In particular, there need be no error about the what the semantic content is. *Semantic content* is the theorist’s notion—ordinary speakers are only interested in what speakers say/assert. And in this case, they are right about what was said.

12.6 ARGUMENT #3: MACFARLANE AGAINST THE CONTEXTUALIST

Contextualism about ‘know’ is the view that ‘know’ is semantically context sensitive. This view is motivated by the variability data outlined in D1. The problem for contextualists is, as we have seen, the stability data brought out in D2. MacFarlane considers three ways in which the contextualist can try to account for this data:

1. Pragmatic Explanation: It’s just too complicated to re-express what’s asserted by an utterance of ‘A knows that p’ in a context other than the one we are in—so instead we just withdraw our earlier knowledge claims.

The two other imagined replies from the contextualist are versions of the error theory, i.e. the view according to which speakers do not fully understand that (or the way in which) ‘know’ is context sensitive. MacFarlane considers two versions of this view:

2. Error Theory Version (I): They don’t realize ‘know’ is context sensitive: “the contextualist can suppose . . . that ordinary speakers are wrong about the semantics of ‘know’—treating it as use-invariant when it is not—” (MacFarlane (2005), p. 215)
3. Error Theory Version (II): Speakers are mistaken about what standards are in play: “[Speakers] make systematic errors about what standards are in play in contexts other than their own.” (p. 215)

MacFarlane argues that all of 1–3 fail. I won’t evaluate those arguments here—my goal is to show that PCR opens up another option for the contextualist. PCR allows the contextualist to argue as follows: suppose the semantic content of an utterance u1 of “A knows that q” in context of utterance C1 is the proposition *that A knows (by low standards) that q*. Call this proposition P1. Suppose the semantic content of an utterance u2 of “A knows that q” in a context of utterance C2 is the proposition *that A knows (by high standards) that q*. Call this proposition P2. So far this is just standard contextualism: the view that the semantic content of ‘know’ varies between context of utterance depending on the ‘standards that are in play’ (as contextualists like to say). Suppose the speaker of u1 finds herself in C2 (the high-standards context). She’s thinking back on her utterance in C1 (the low-standards context) and relative to C2, what is saliently asserted by u1 is P2, i.e., the proposition *that A knows (by high standards) that q*. Since this is what’s saliently asserted by u1 relative to C2, she’ll think and say (in C2) that what she said in C1 is false (assuming that A doesn’t satisfy the high standards for knowledge with respect to q). This is compatible with the semantic content of u1 being P1. It is even compatible with the semantic content

being true (if we assume that A *does* satisfy the low-standards requirement for knowledge with respect to q).

Natural question: why think that relative to C2 (i.e. the high-standards context) u1 asserts P2? The answer, for a proponent of PCR, is simple: because that's what the data indicates. Speakers in C2 treat the speaker in C1 as if she has said something false, i.e. as if she has expressed a proposition that's false. Assuming that u1 (saliently) asserted P2 relative to C2 explains this data.¹⁹ Of course, there's no denying that when you first encounter it, this kind of view sounds strange. I don't think that should be too much of a worry. Here's a hypothesis: we are trying to explain exceedingly strange data patterns and any attempted explanation will contain surprises from a pre-theoretical point of view. We will, at the end of the day, be comparing strength and weakness of strange explanations. The truth-relativist is at one extreme of this strangeness: according to MacFarlane, the proposition asserted by the speaker in C1 is false from C2 and true from C1. The truth or falsity varies even though what was said is the very same proposition. PCR has the advantage of preserving the intuitive idea that a proposition has a stable truth-value no matter who you are, where you are, what you care about etc. PCR is, moreover, independently motivated. It is not, as truth-relativism is, introduced specifically to deal with these data patterns.

12.6.1 Point of Clarification: How Speech Act Content is generated

So far I have said nothing about the mechanism that makes it the case that an utterance of S ends up (saliently) asserting a proposition p relative to a context of interpretation. Notice that in this respect, any PCRist is free to piggyback on the relativist (who, as MacFarlane points out, can piggyback on the contextualist's story). The relativist owes us a story about how a particular standard becomes salient in a particular context of assessment; some story about why, in a context of assessment, it is standard S1, out of all the infinitely many possible standards, that determines the truth of a proposition p. Of course, at this point, no such theory is available, but we should not hold that against the relativist. Nor should we hold it against PCR that it does not come with a theory of how speech act content is generated and becomes salient. It's common ground between relativist and PCR that the salient standards vary between contexts of interpretation/assessment. For the relativist, this contributes to fixing a truth-value. According to PCR it determines the set of propositions asserted (relative to a context of interpretation).

¹⁹ In § 3 above I pointed out that PCR implies that speakers might end up saying things they would be surprised to be told they have said. This is an illustration of how that can happen.

12.7 ARGUMENT #4: LASERSOHN ON FAULTLESS DISAGREEMENT

Lasersohn characterizes the puzzle that leads him to endorse truth-relativism as follows:

Our basic problem is that if John says ‘This is fun’ and Mary says ‘This is not fun’, it seems possible for both sentences simultaneously to be true (relative to their respective speakers), but we also want to claim that John and Mary are overtly contradicting or disagreeing with each other . . . (Lasersohn (2005), p. 662)

So consider utterance *u*₁ by A of ‘Roller coasters are fun’ and *u*₂ by B of ‘Roller coasters are not fun’. It is Lasersohn’s intuition that:

(I1): Both A and B are right; both said something true (assuming that roller coasters are fun for A, but not for B).

(I2): A said that roller coasters are fun and B denied it, so they disagree.

Lasersohn’s strategy for reconciling these intuitions is to invoke truth-relativism—the same proposition (or proposition-like object²⁰) *P* is affirmed by A and is denied by B, but that proposition is true relative to A’s circumstance of evaluation, and false relative to B’s. Lasersohn says:

What I would like to suggest is that we refine the notion of disagreement so that two people can overtly disagree—we might even go so far as to say they contradict each other—even if both their utterances are true. (p. 662).

This suggestion is implemented as a variation on Kaplan’s view in ‘Demonstratives’ by adding people to the circumstances of evaluation, in addition to the three parameters Kaplan had (worlds, times, and locations):

All we have to do is assign words like fun and tasty the same content relative to different individuals, but contextually relativize the assignment of truth values to contents, so that the same content may be assigned different truth values relative to different individuals. This will allow for the possibility that two utterances express identical semantic content, but with one of them true and the other one false. (Lasersohn 2005, 662).

This is Lasersohn’s explanation of how I1–I2 can both be true.²¹

²⁰ Though Lasersohn characterizes these objects as propositions they are not propositions in the standard sense, i.e. they are not functions from worlds to truth-values.

²¹ There’s an important difference between Lasersohn and MacFarlane. Lasersohn does not relativize truth to contexts of assessment; he includes people (and their standards of taste) as parameters in circumstances of evaluation. MacFarlane does not consider that a full-blooded form of relativism. I don’t think the difference between those two versions of relativism has any bearing on the arguments in this paper.

Again, I'll argue that PCR renders this move to relativism superfluous. Some of what I say here will be familiar and (hopefully) predictable by now—what it adds is a focus on the notions of agreement and disagreement.

12.7.1 The PCR Explanation of (I1) and (I2) and the Illusion of Faultless Disagreement

For those who share Lasersohn's intuitions about these cases (I'm not one of those²²), here's how PCR can help explain them:

Explanation of (I1): According to PCR, both *u*₁ and *u*₂ express many propositions. Among these, assume *u*₁ expresses the proposition *Q* (e.g. the proposition *that roller coasters are fun for A*) and *u*₂ expresses proposition *R* (e.g. the proposition *that roller coasters are not fun for B*). Both *Q* and *R* can be true; and that explains the intuition that both *u*₁ and *u*₂ are true.

Explanation of (I2): Relative to the theorist's context of interpretation, there is a proposition *P* that *u*₁ asserts and *u*₂ denies. This explains the intuition that *u*₁ and *u*₂ contradict one another. For now, let's say that *P* is the proposition *that roller coasters are fun*. I say more about *P* below.

PCR explains how there can be an apparent contradiction between *u*₁ and *u*₂, even though both can be true. The apparently contradictory intuitions occur because we shift our attention from one part of the total speech act content to another. When the focus is on *Q* and *R*, we can say truly that *A* and *B* spoke truly. When the focus is on *P*, we can say truly that they disagree and that only one of them is right.²³ On this view, faultless disagreement is an illusion. There is no *one* proposition such that *A* commits to it, *B* commits to its negation, and both *A* and *B* are right. That's impossible. We get the illusion that something like this is going on by shifting our attention (in ways described above) from one part of asserted content to another.

12.7.2 More On the PCR Explanation of I2

In the explanation of (I2) I described the shared (semantic) content of the sentence *A* and *B* both uttered as the proposition *that roller coasters are fun*. I didn't say whether, in this case, I am opting for contextualism or invariantism

²² I very much doubt that the intuition of disagreement is solid in cases involving predicates of taste, but I will not argue for that view here. Lots of people seem to share Lasersohn's intuitions about these cases, and my goal is to show how PCR can help account for them.

²³ There are interesting questions about whether an interpreter changes context of interpretation as she shifts her attention from one part of the speech act content to another. More generally, there are interesting questions about just how to individuate contexts of interpretation—at this point I'm not convinced I need to take a stand on these issues. As far as I can tell my view could go in either direction.

about ‘fun’. As with MacFarlane’s arguments, my goal is to show that once PCR is adopted, the kind of intuition pattern Lasersohn attempts to explain (that provides the central motivation for his version of relativism) is compatible with *any* view of the semantics for ‘fun’, as long as it is combined with PCR. I show this with respect to two versions of invariantism and one version of contextualism.

12.7.3 Two Invariantist Explanations of I2

According to I2 there’s a content that A asserts and B denies. The structure of the invariantist’s explanation could be the following. There’s a stable semantic content expressed by all utterances of ‘roller coasters are fun’. A asserts it and B denies it. So far, so good.²⁴ The challenge is to say something about what this stable semantic content is. I won’t try to settle that question here. I simply consider Lasersohn’s objection to two invariantist options and show how PCR undermines those objections.

- *Proposal #1*: ‘fun’ semantically denotes the property of triggering a certain kind of reaction in *some* people (where ‘certain reaction’ is spelled out in more detail).

No one disputes that this property exists. The question is whether it is the semantic value of ‘fun’. About that proposal, Lasersohn says:

... this analysis cannot be right, because the truth conditions are simply much too weak. Suppose Mary doesn’t like roller coasters at all. I think she can sincerely say *This is not fun* as she rides a roller coaster, even if she knows that many other people do enjoy them. But under this analysis, she could not say this, at least not sincerely. (p. 653)

If PCR is true, the analysis *can* be right. The assumption underlying the reply is that Mary (saliently) says the semantic content of the sentence she’s uttering. Lasersohn assumes that speakers utter their sentences with the intention of communicating and committing to the semantic content of the uttered sentences. If PCR is correct, that’s not what speakers typically do. According to PCR, speakers often assert propositions they are not intending to assert and are not committing themselves to. Someone who utters “Roller coasters are not fun” would be surprised to find that she had denied *that roller coasters trigger reaction R in some people*. She certainly didn’t intend to *commit* to this view. That, however, is nothing peculiar to sentences containing ‘fun’ and other predicates of taste. According to PCR this phenomenon is ubiquitous.

Since we are comparing truth- and content-relativism here, it is also important to note that the truth-relativist is not in a position to complain over this feature of PCR. According to Lasersohn’s version of truth-relativism, the speakers deny

²⁴ Note, however, that an invariantist need not insist that the shared content is the semantic content. She could appeal to some other part of the speech act content that’s shared. I think that’s what an invariantist should do in many cases, but I won’t pursue that option here since my goal is to show how Lasersohn’s arguments fail and he doesn’t consider this option for the invariantist.

and affirm a 'proposition' that is standard *neutral*. This neutral object is what they allegedly disagree over. But that's not substantive disagreement. To see this, note that on Lasersohn's view, there's no more disagreement between two speakers who utter "Roller coasters are fun" and "Roller coasters are not fun" than there is between two speakers who utter "Tim is wearing blue socks" at different times. On the analogous Kaplan semantics for tensed sentences, utterances, the two speakers have asserted temporally neutral propositions, and these are evaluated at (world, time) pairs. There is no way in which we would consider these two speakers as disagreeing. The situation for predicates of taste is exactly analogous.²⁵

Proposal #2: This is the view that 'fun' is a 1-place predicate, with no relativization to a standard or a group. Think of this as the objectivist view of *fun*, where the semantic value is the one and only standard for fun (whatever that might be—a proponent of this view would have to tell us).

About this kind of view, Lasersohn says:

In a loose sense, this approach could perhaps be regarded as the analog in the area of predicates of personal taste to Williamson's (1994) analysis of vague predicates. Williamson argues that the meanings of apparently vague predicates actually have sharp boundaries, so that, despite appearances, there is a definite fact of the matter as to whether someone is thin or not, for example; and we would likewise be claiming that there really is a definite fact of the matter as to whether roller coasters are fun or not. (p. 655)

Lasersohn's comparison to Williamson's account of vagueness is helpful here and does provide support for the invariantist. Lasersohn, however, finds the comparison problematic. He points out that Williamson's account of vague terms depends on the idea that we *cannot* come to know where the sharp boundaries lie. In the case of predicates of taste, Lasersohn argues, it's the other way around: we are epistemically *privileged*. He says:

But with predicates of personal taste, we actually operate from a position of epistemic privilege, rather than the opposite. If you ride the roller coaster, you are in a position to speak with authority as to whether it is fun or not; if you taste the chili, you can speak with authority as to whether it is tasty. (p. 655)

If we assume PCR, this alleged disanalogy disappears: you can be in an epistemically privileged relation to one part of the speech act content, but not to another. In this particular case you won't be in an epistemically privileged position with respect to the semantic content, but you will be with respect to some other part of the speech act content. If you sincerely utter 'Roller coasters are fun' you have asserted one proposition with respect to which you are not in an epistemically privileged position and one proposition with respect to which you are in an

²⁵ Again, I should emphasize that my point here is not to endorse this as the semantics for 'fun'. I just want to point out that when PCR is introduced, the kind of reply Lasersohn gives doesn't work.

epistemically privileged position. The former is the semantic content and the latter is, say, the proposition *that roller coasters are fun by your standards*.²⁶

These are but two illustrations of how an invariantist could use PCR to respond to Lasersohn's objections. What these considerations show is not that invariantism is the correct semantics for 'fun', but that Lasersohn's strategy for responding to invariantism fails. One way to make it perspicuous that these objections to Lasersohn do not favor invariantism is to note that the same kind of defense can be given of contextualist semantics for 'fun'. I turn to that next.

12.7.4 Two Contextualist Explanations of I2

Here is how a contextualist can account for much of the data about predicates of taste, in particular, the data about agreement and disagreement. The semantic content of A's utterance of "Roller coasters are fun" in C1 is *that roller coasters are fun by A's standards*, the semantic content of B's utterance of "Roller coasters are not fun" in C2 is *that roller coasters are not fun by B's standards*. So far this is a version of contextualism about 'fun'. Suppose someone judges, from a context of interpretation C3 that A and B disagree. How could that be? Here are two moves a contextualist who endorses PCR could make:

- (i) Let's say E is the person making the report in C3, i.e., she is the person uttering, "A and B disagree" in C3. Here is one possible candidate for P: the proposition *that roller coasters are fun for E*. This would be surprising, of course. How could A and B make assertions about E? They might not have been thinking about E, they might not even know who E is. So how can they make assertions about her? Well, PCR is a radical view, and this would be a particularly radical version. But it is not entirely implausible—in particular, it should not seem particularly implausible to someone who is prepared to endorse truth-relativism. The truth-relativist thinks that the truth-value of what A asserted and B denied depends on *E's standards of assessment* (i.e., according to the relativist, when E thinks about whether what A and B said is true, she thinks about whether roller coasters are fun for *her*). It is but a short step from this view to the view that P is what A asserted and B denied relative to E's context.²⁷
- (ii) It should be clear by now that this is not the only option for a PCR-contextualist. The object they disagree over could be one of the propositions appealed to by semantic invariantists. For the contextualist these propositions would not be the semantic contents, but they could be part of the speech

²⁶ This is on the assumption that we have privileged access to our own standards of taste, something there's reason to doubt. I leave those issues aside here.

²⁷ This view is easier to swallow if one endorses a strong version of content relativism, i.e. the view described as P2.2 in § 3 above.

act content relative to C3 and so, in some cases, explain agreement and disagreement judgments.²⁸

12.8 CONCLUSION: PCR, SEMANTIC BLINDNESS, AND THE FOG OF CONTENT

Here's a way to think of my view. If PCR is correct, there are, for each utterance, lots and lots of asserted propositions; what proposition(s) is (/are) salient will shift from one context of interpretation to another. So if you try to use your intuitions about what's (saliently) asserted as your guide to semantic content, you're bound to get confused—you'll find puzzling data patterns, inexplicable shifts across contexts of utterance and across contexts of interpretation. You're in a fog of content where it's hard (if not impossible) to distinguish the semantic content from all the other content that's floating around. It is impossible to *directly* access semantic content in this fog—none of our intuitions about content come with “True Semantic Content” stamped on them. The correct response, according to PCR, is to stop using these kinds of intuitions as your guide to semantics.

I'll end by mentioning two (of many) issues that should keep a proponent of PCR awake at night:

- In this paper I have not said anything about what I take to be the correct semantics for ‘know’ and ‘fun’. I have shown that both contextualism and invariantism can be defended against the kinds of objections MacFarlane raises, but that doesn't settle the question of which one of these is correct. In general, I favor what I call ‘minimalist’ and invariantist theories (see C&L 2004), but no argument in this paper depends on accepting that view.
- I have told you a lot about how *not* to do semantics if PCR is true, and I have told you very little about how to do it. Indeed, it might look like we have no way to do semantics if PCR is right. Appeal to intuitions about content is one of the semanticist's central tools. If those intuitions are completely unreliable, how do we proceed? Here's how I see it: there's overwhelming evidence in favor of PCR. Some, those I call Radical Contextualists (see C&L 2004, chapter 3), conclude from this that systematic semantics is impossible. That's not my conclusion. I think PCR presents a challenge to semanticists. The challenge is to figure out how to find stable semantic content in the midst of speech act content that varies, not only between contexts of utterance, but

²⁸ It is also useful to think about the following. Imagine B hearing A's utterance of “Roller coasters are fun”. Disliking roller coasters, she replies, “No!” What is she denying? Again, a proponent of PCR has a number of options. B is denying some proposition said by A relative to B's context of interpretation. This could be the proposition in which ‘fun’ is indexed to B's preferences (that would be the analogue of (i) in the text), or the unrelativized proposition *that roller coasters are fun* (the analogue of one version of (ii)). (Thanks to Elia Zardini for suggesting I mention this kind of case.)

also across contexts of interpretation. In C&L (2004), I say a little bit more about how that can be done. But I don't say much and I consider those proposals work in progress. So at this point, consider the claim that PCR does not mean the death of semantics a promissory note.²⁹

References

- Cappelen, H. (2005) 'Pluralistic Skepticism: Advertisement for Speech Act Pluralism', *Philosophical Perspectives* 19: 15–39.
- and Lepore, E. (1997) 'On an Alleged Connection Between Indirect Quotation and Semantic Theory', *Mind and Language*, vol. 12, 278–96.
- , — (1998) 'Reply to Richard and Reimer', *Mind and Language* 13(4): 617–21.
- , — (2002) 'Insensitive Quantifiers', in J. Campbell, M. O'Rourke and D. Shier (eds.), *Meaning and Truth: Investigations in Philosophical Semantics*, New York: Seven Bridges Press, 42–59.
- , — (2003) 'Context Shifting Arguments', *Philosophical Perspectives* 17: 25–50.
- , — (2004) *Insensitive Semantics*, Oxford: Blackwell.
- , — (2006) 'Shared Content' in E. Lepore and B. Smith (eds.), *Oxford Handbook of Philosophy of Language*, Oxford: Oxford University Press, 1020–55.
- Egan, A. (forthcoming): 'Epistemic modals, relativism, and assertion', *Philosophical Studies*.
- , Hawthorne, J., and Weatherson, B. (2005): 'Epistemic Modals in Context', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*, Oxford: Oxford University Press, 131–69.
- Kölbel, M. (2002) *Truth Without Objectivity* London: Routledge.

²⁹ MacFarlane and Egan, Hawthorne, and Weatherson (2005) discuss a view they call 'content relativism', or 'expressive relativism'. That view is not PCR because it doesn't include P1 and it doesn't take P2.1 into consideration. As a result most of the objections raised are off the mark with respect to PCR. I'll mention one objection and the kind of reply a proponent of PCR would give (this strategy generalizes to many of the other objections insofar as they can be seen as objections to PCR): MacFarlane, referring to an objection from Egan, Hawthorne, and Weatherson, says that if content relativism were true:

... we could no longer say, with Stalnaker 1978, that the effect of assertion is to add the proposition asserted to a 'common ground' of presupposed propositions, for there may be no common fact of the matter about which proposition was asserted. (2005, p. 220)

Quite frankly, I'm not sure what the essential function of assertion is. I'd like to have a view about that, but I don't. However, if you're attracted to Stalnaker's story, it's easy to get something like it if you endorse PCR. Here's one possibility:

The essential effect of assertion is to add the proposition(s) saliently asserted relative to the context of utterance to the common stock of presuppositions.

Of course, when the context of interpretation is not identical to the context of utterance, it is unclear how to apply the Stalnakerian framework. The theory is not meant to deal with inter-contextual cases. It's hard, therefore, to see it as a particular fault of PCR that it fails to give a natural answer to what goes on in those cases.

- Lasersohn, P. (2005) 'Context Dependence, Disagreement and Predicates of Personal Taste', *Linguistics and Philosophy* 28(6): 643–86.
- MacFarlane, J. (2005) 'Assessment Sensitivity of Knowledge Attributions', in Tamar Szabó Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology* 1, Oxford: Oxford University Press, 197–233.
- (2007) 'Semantic Minimalism and Nonindexical Contextualism', forthcoming in G. Preyer and G. Peter (eds.), *Content and Context: Essays on Semantics and Pragmatics*, Oxford: Oxford University Press.
- Reimer, M. (1998) 'What is Meant by What is Said? A Reply to Cappelen and Lepore', *Mind and Language* 13(4): 598–605.
- Richard, M. (1998) 'Semantic Theory and Indirect Speech', *Mind and Language*, 605–16.
- Soames, S. (2002) *Beyond Rigidity: The Unfinished Agenda of Naming and Necessity*, Oxford: Oxford University Press.
- (2005) 'Naming and Asserting', in Z. Szabó (ed.), *Semantics vs. Pragmatics*, Oxford: Oxford University Press.
- (forthcoming) 'The Gap Between Meaning and Assertion', in Martin Hackl and Robert Thornton (eds.) *Asserting, Meaning, and Implying*.
- Williamson, Timothy (1994) *Vagueness*, London: Routledge.

13

Faultless or Disagreement

Andrea Iacona

Among the various motivations that may lead to the idea that truth is relative in some non-conventional sense, one is that the idea helps explain how there can be “faultless disagreements”, that is, situations in which a person A judges that p , a person B judges that not- p , but neither A nor B is at fault. The line of argument goes as follows. It seems that there are faultless disagreements. For example, A and B may disagree on culinary matters without either A or B being at fault. But standard semantics has no room for such a case, so there is something wrong with standard semantics. The best way to amend it is to add a new parameter that is relevant to the truth or falsity of what A and B judge, presumably, something related to personal taste.¹ The aim of what follows is to show that this line of argument is flawed. It is true that standard semantics has no room for faultless disagreements. But there is nothing wrong with this, for it should not be assumed that such disagreements exist.

1. The claim that there are faultless disagreements is at odds with a schema that seems correct at first sight:

F If p and one judges that not- p , then one is at fault.

F seems correct because all it says is that if things are a certain way, it is wrong to judge that they are not that way. In other words, if it is true that p then it is not true that not- p , hence judging that not- p is judging something not true. The clash between F and the claim that there are faultless disagreements is shown by the following proof:

1	A judges that p and B judges that not- p	Assumption
2	Neither A nor B is at fault	Assumption
3	p	Assumption
4	B is at fault	1, 3 F
5	Not- p	2, 3, 4 RAA

¹ The expression ‘faultless disagreement’ is used by Max Kölbel. The line of argument appears in Kölbel (2003), Kölbel (2008) and Lasersohn (2005).

6	A is at fault	1, 5 F
7	Not-2	2, 2, 6 RAA

1 advances the hypothesis that A and B disagree. 2 adds to that hypothesis that the disagreement is faultless. 3 is an adjunctive assumption that leads from 1 to 4 with the help of F. 5 is derived from 2, 3 and 4 by *reductio ad absurdum*, so 3 is discharged and the following lines do not depend on it. 6 is warranted by F given the apparently innocuous condition that if one judges that *p* then one judges that not-not-*p*. For on that condition F entails that if not-*p* and one judges that *p*, then one is at fault. 7 is obtained again by *reductio ad absurdum* and depends solely on 1, given that 2 is discharged. From 1–7 it turns out that if A and B disagree then it is not the case that neither A nor B is at fault, which is classically equivalent to the conclusion that if A and B disagree then either A or B is at fault.²

The proof set out entails the negation of the claim that there are faultless disagreements, so either there is something wrong with it, or the claim is not true. An argumentative move in favour of the first alternative goes as follows. If one accepts the proof, one must have something to say about the disputes that qualify as good candidates for being faultless disagreements, that is, the disputes where two persons A and B seem to contradict each other—in that, say, A assertively utters a sentence and B assertively utters its negation—yet the topic is not one of those for which it is clearly inappropriate to suppose that neither A nor B is at fault. Two options are available, but neither of them is attractive. The first is to adopt a “realist” view according to which such disputes are really disagreements, but they are only apparently faultless. That is, either A or B is at fault, even if there may be no way of settling the issue and find out where the mistake is. Realism is unattractive in that it does not account for the fact that we are inclined to regard certain disputes as faultless, hence it is unable to explain the difference between such disputes and those we are not inclined to regard as faultless. The second option is to adopt a form of “indexical relativism” according to which the candidate disputes are really faultless, but they are only apparently disagreements. In this case the view is that the judgments conveyed by A and B are, or are equivalent to, those expressed by sentences involving reference to A and B themselves. For example, if A utters ‘*x* is tasty’ and B utters ‘*x* is not tasty’, the judgments they convey are, or are equivalent to, those expressed by ‘*x* is tasty for me’ and ‘*x* is not tasty for me’, so A and B do not really contradict each other. Indexical relativism is unattractive in that it doesn’t account for the fact that we are inclined to regard at least some candidate disputes as disagreements, hence it is unable to explain the difference between such disputes and the cases in which

² Similar derivations are considered in Wright (2001: 56), Kölbel (2002: 24–5), and Kölbel (2003: § 2).

two parties express mere preferences by using overtly indexical sentences.³ This reasoning seems to draw its appeal from two unquestionable facts. The first, invoked against realism, is that there are disputes about culinary matters or similar topics that we are inclined to regard as faultless. An example may illustrate. Two wine connoisseurs buy a bottle of Barolo and a bottle of Barbaresco for €100 each. After tasting both, one of them says “This is better”, and the other replies “No, that is better”. Acknowledging the divergence, each of them feels the need to reassert the favoured view in front of the other. But a couple of attempts unaccompanied by appreciable progress lead them quickly to drop the subject and start talking about something else. It is no accident that they do not embark on a thorough discussion on the virtues of Barolo and Barbaresco. Usually, when one argues for a view one thinks that there are reasons to accept the view, and hence that by stating those reasons one may be able to convince another person that initially does not accept it. There may be something about which that person is mistaken or ignorant, or there may be a type of experience of which that person is devoid. But in our case neither of the two disputants thinks that way.

The second fact, invoked against indexical relativism, is that there are disputes about culinary matters or similar topics that we are inclined to regard as disagreements. An example of such dispute may be the following. While the two connoisseurs are still talking, a friend of them shows up with some wine in a box that he bought for €0.99 at a discount shop. They offer him a glass of Barolo, and he offers his wine in return. After tasting the Barolo and learning that the bottle cost €100, he says “This is no better than my own”, and adds a disapproving look to comment on the price. Then one of the two connoisseurs retorts “Yes, it is!”, and all of them start discussing wine. Although the friend seems not to give in easily, the two connoisseurs keep arguing with him, because they think there are reasons to believe that Barolo is better. It is their deep conviction that anyone sufficiently exposed to the type of experience they had would agree with them.

Examples like the first make it plausible to claim that there are faultless disputes about culinary matters or similar topics, while examples like the second make it plausible to claim that there are disagreements about such topics. However, it is essential to recognize that acceptance of these two claims does *not* amount to acceptance of the claim that there are faultless disagreements. If it did, the reasoning would be blatantly circular. Clearly, one is bound to accept the two claims if one accepts the claim that there are faultless disagreements. If there is at least one faultless disagreement about culinary matters or similar topics, then there is at least one dispute about such topics that is faultless and at least one dispute about such topics that is a disagreement. But the converse does not hold. Assuming that there are disputes about culinary matters or similar topics that are faultless and disputes about such topics that are disagreements, it does not follow

³ This line of argument emerges from Wright (2001: § 1), and Kölbel (2003: §§ 3–4).

that there are faultless disagreements about such topics. For nothing guarantees that the first two sets of disputes have some member in common.

2. The upshot of the argument outlined is that one should think it over before accepting the proof. But perhaps one should think it over before accepting the argument. Undoubtedly, realism and indexical relativism as they have been depicted are quite unattractive options. However, one question is whether each of these two options is tenable, another question is whether the conclusion of the proof is true. The distinction is crucial, because the considerations that may be advanced against realism and indexical relativism do not by themselves undermine the proof.

It is clear how the two facts invoked give evidence against the two views. Realism is taken to entail that all the disputes that qualify as good candidates for being faultless disagreements are disagreements. Since we are inclined to say that some disputes of that kind are faultless, the realist is left with the problem that some disagreements seem faultless. In the same way, indexical relativism is taken to entail that all the disputes that qualify as good candidates for being faultless disagreements are faultless. Since we are inclined to say that some disputes of that kind are disagreements, the indexical relativist is left with the problem that some faultless disputes seem disagreements. That is, in one case the problem arises because the disputes the first inclination is about are assumed to be *disagreements*, while in the other it arises because the disputes the second inclination is about are assumed to be *faultless*.

What the argument tends to conceal, however, is that although each of these two assumptions is required by the respective view, neither of them is required by the conclusion of the proof. All that the proof entails is that, given any dispute, either it is not faultless or it is not a disagreement. Take a dispute where the first inclination is strong, such as that between the two connoisseurs. In this case it is plausible that the first disjunct does not hold, so the consequence we get is that the second disjunct must hold. But this doesn't seem to be a problem, unless it is assumed that any dispute of the kind considered is a disagreement. The same goes for a dispute where the second inclination is strong, such as that between the two connoisseurs and the friend. In this case it is plausible that the second disjunct does not hold, hence we get that the first must hold. Again, the consequence seems to pose no problem without the assumption that any dispute of the kind considered is faultless. Is there some other piece of evidence, besides the two unquestionable facts considered, that may be invoked to question the proof? It seems not. As far as evidence is concerned, the proof stands. Of course, one apt to doubt the proof will probably regard the claim that there are faultless disagreements as "intuitive" on its own. The point, however, is that it is far from obvious that the claim is more intuitive than F or than the logic employed to derive $\mathcal{I}$. Therefore, insisting on it certainly would not suffice to undermine the proof.

Actually, it is not even obvious that the claim is intuitive *at all*. It is hard to find a clear example of a dispute that appears to be faultless and at the same time appears to be a disagreement. For any case that we are able to describe with sufficient accuracy, our inclinations seem to push us one way or the other. Thus in the case of the two connoisseurs we have no strong inclination to think that there is genuine disagreement. A natural thing to say about this case is that the divergence boils down to the fact that one of them prefers Barolo while the other prefers Barbaresco. Similarly, in the case of the two connoisseurs and the friend we have no strong inclination to think that the dispute is faultless. It is plausible that the connoisseurs are right and the friend is wrong. After all—we are tempted to say—expensive Barolo *is* better than boxed cheap wine.

3. Unlike what some may have thought, the thesis that no disagreement is faultless leaves room for a simple and straightforward account of the two unquestionable facts considered. All that is needed to make sense of those facts is to acknowledge *another* fact, which is so evident that one might wonder how it can go unnoticed. The fact is that there is an underlying ambivalence in our ordinary way of talking. Generally, when we express evaluative judgments on culinary matters or similar topics we assertively utter sentences of the form '*x* is *P*'. But there are different things we may want to say with such sentences. Sometimes we utter '*x* is *P*' to express a personal inclination towards *x*, that is, we mean that *x* is able to produce certain effects on us that characterize it as *P* for us. Call it *subjective* use of '*x* is *P*'. For example, in the case of the two connoisseurs 'This is better' and 'That is better' are intended to say nothing more than each of them would drink the referred wine if he had to choose. Neither of them is likely to oppose a paraphrase of the sentence uttered such as 'I prefer it' or 'I like it better'. At other times, instead, we utter '*x* is *P*' meaning that *x* is able to produce certain effects independently of our personal inclinations toward *x*. That is, we take *x* to be *P* for the average person, or for a normal person, or for any person under ideal conditions. Call it *objective* use of '*x* is *P*'. For example, in the case of the two connoisseurs and the friend, one of the two connoisseurs utters 'Yes, it is!', meaning that Barolo is better independently of his inclinations towards it. Presumably, he is not willing to rephrase what he utters as 'I prefer it' or 'I like it better'.⁴

Given that '*x* is *P*' may convey different kinds of judgment, the same goes for '*x* is not *P*'. Hence there are different kinds of dispute that two persons *A* and *B* may have if *A* utters '*x* is *P*' and *B* utters '*x* is not *P*', depending on what they actually mean. One kind of dispute is that in which *A* and *B* use '*x* is *P*' and '*x* is not *P*' subjectively, that is, *A* means that *x* is *P* for *A*, and *B* means that *x* is not *P* for *B*. In this case the dispute is faultless, but it is not a disagreement.

⁴ As Isidora Stojanovic nicely puts it, there are cases in which the two parties are likely to move from Yes/No dialogue to Ok/Ok dialogue, and cases in which they aren't. See Stojanovic (2007), § 2.

Another kind of dispute is that in which A and B use ' x is P' and ' x is not P' objectively, that is, A means that x is P independently of A's inclinations towards x , and B means that x is not P in the same sense. In this case the dispute is a disagreement, but it is not faultless. Obviously, there are also disputes in which A and B differ in the kind of judgment conveyed. But any such dispute involves a misunderstanding, so it is likely to turn into a dispute of one of the other two kinds if A and B become clear about their respective intentions.

This shows how the two unquestionable facts can coexist. The first, namely, that there are disputes about culinary matters or similar topics that seem faultless, is easily explained by saying that those disputes *are* faultless, that is, they are of the first kind. The same goes for two related facts. One is that it is natural to distinguish between faultless disputes and disputes that are not faultless.⁵ This is explained by saying that there are both disputes of the first kind and disputes of the second kind. The other is that it is natural to distinguish between topics on which faultless disputes can arise, such as culinary matters, and topics on which they cannot arise, such as mathematics.⁶ This is explained by saying that there are topics on which ordinary discourse is ambivalent in the way described, and topics on which it is not. Normally, when we talk about mathematics we take any assertive utterance of a sentence of the form ' x is P' to convey that x has the property of being P independently of the inclinations of the speaker.

The explanation of the second fact is analogous. The reason why some disputes about culinary matters or similar topics seem to be disagreements is that they *are* disagreements, that is, they are of the second kind. A related fact is that in some cases we talk about such topics assuming that there is a difference between using a sentence of the form ' x is P' and using an overtly indexical sentence, such as 'I like x '.⁷ This is explained by saying that in some cases ' x is P' does not merely express a personal inclination towards x .

To say that any dispute belongs to one of the two categories indicated is not the same thing as to say that, for any case in which a person assertively utters ' x is P' and another person assertively utters ' x is not P', it is clear what kind of dispute they are having. In many cases it is *not* clear what kind of judgment one conveys by uttering such a sentence, because it is not clear what one has in mind. This may be due to the fact that one has in mind something that is not immediately manifest at the moment of the utterance, but that would come to the surface if the discussion were to progress. Or it may happen that one doesn't even have in mind something definite at a given stage, although one could end up by endorsing a clear-cut position at a later stage. A consequence of this unclarity is that in many cases we have no strong inclination to say that the dispute belongs to a given category between the two, hence we have no strong

⁵ Wright (2001: 48), Kölbel (2003: 53).

⁶ Kölbel (2003: 58).

⁷ Wright (2001: 51–2), Kölbel (2002: 38–40), Kölbel (2003: 64).

inclination to deny that it belongs to the other. This may foster the illusion that there are intuitive faultless disagreements.

4. The line of argument sketched at the beginning of the paper rests on the claim that there are faultless disagreements. For according to that line of argument, the problem with standard semantics is that it is unable to account for such disagreements. What has been said so far entails that there is no reason to accept the claim, hence that there is no reason to regard the indicted inability as a problem. However, this leaves open the question of how ordinary discourse on culinary matters or similar topics can be accounted for in compliance with the distinction between disputes of the first kind and disputes of the second kind. Therefore, it might still be suspected that standard semantics is in trouble. This last section is intended to dissipate that suspicion. The ambivalence in our use of sentences of the form ' x is P ' can plausibly be explained in accordance with standard semantics, so the same goes for the distinction between disputes of the first kind and disputes of the second kind.

There are basically two ways of explaining the ambivalence in our use of sentences of the form ' x is P '. One is to appeal to *pragmatic* considerations such as the following. Given the linguistic meaning of ' x is P ', what a speaker 'says' by uttering it is always that x has the property of being P . Semantically, the property of being P is just like any other property. There is no significant difference between ' x is P ' and, say, ' x is white'. However, not every utterance of ' x is P ' is intended to 'communicate' that x has the property of being P . There are utterances of ' x is P ' that convey something different, that is, something that could be said by uttering an indexical sentence such as 'I like x '. When that happens, the use of ' x is P ' is subjective. Otherwise, it is objective. According to this explanation, what matters to identify a dispute between A and B is what they communicate. Thus a dispute of the first kind is faultless, in that both A and B communicate something true. But it is not a disagreement, in that there is no contradiction in what they communicate. Similarly, a dispute of the second kind is a disagreement, in that A and B communicate what they say. But it is not faultless, in that they can't both say something true.

The other way is to couch the difference between subjective and objective use of ' x is P ' *semantically*, that is, as a difference in the interpretation of ' x is P '. According to this explanation, a dispute of the first kind is one in which the sentences uttered are interpreted "subjectively", hence in different ways. By contrast, a dispute of the second kind is one in which they are interpreted "objectively", hence in the same way. That is, in the first case A assertively utters ' x is P ', B assertively utters ' x is not P ', but their interpretations α and β of ' x is P ' differ in such a way that the sentence is true relative to α and false relative to β . In this case neither A nor B is at fault, in that ' x is P ' is true relative to α and ' x is not P ' is true relative to β . But there is no disagreement between A and B , because α and β differ. By contrast, in the second case A assertively utters ' x is P ',

B assertively utters ‘ x is not P ’, and there is an interpretation α of ‘ x is P ’ that A and B share. In this case A and B disagree, in that for A ‘ x is P ’ is true relative to α while for B it is false relative to α . But either A or B is at fault, given that ‘ x is P ’ can’t be both true and false relative to α .

Neither of these two ways of explaining the ambivalence in our use of sentences of the form ‘ x is P ’ seems to cause trouble to standard semantics. In the first case the explanation is trivially compatible with standard semantics, because it is compatible with any account of the linguistic meaning of the sentences in question that takes their surface grammar at face value. Semantics doesn’t really matter in this case, as the weight of the explanation is located outside its domain.

In the second case the explanation seems compatible with standard semantics, because the difference of interpretation can easily be described in terms of standard notions. One option is to adopt the notion of *index* to represent the context in which a sentence is uttered. A context is a concrete entity that includes innumerable features, such as a time, a place, a possible world, a conversational situation, and so on. An index is a collection of features of various sorts that may be “extracted” from the context. The notion of index can be combined with the assumption that the logical form of a sentence may not coincide with its surface grammar. On that assumption, a logical form has a number of open positions occupied by elements whose value can vary from context to context. These are either overt context-dependent expressions, such as indexicals, pronouns and demonstratives, or covert variables. Accordingly, an index can be defined as a number of values for those elements. A difference of interpretation can thus be described as a difference of index. That is, if the logical form of ‘ x is P ’ is something like ‘ x is $P(y)$ ’, where y is a covert variable that takes persons or groups of persons as values, a subjective interpretation of ‘ x is P ’ can be defined as one in which the value of y is the speaker, while an objective interpretation of ‘ x is P ’ can be defined as one in which the value of y is a privileged group of persons or an ideal person.⁸

Another option is to adopt the notion of *point of evaluation*. Like the index, the point of evaluation is a collection of contextual features. But these are not necessarily features that go together as part of the same context. For example, a point of evaluation may include a time that is not that of the utterance, or a possible world that is not the actual world. If the notion of point of evaluation is adopted to explain the ambivalence in our use of sentences of the form ‘ x is P ’, a difference of interpretation can be described as a difference in the point at which the sentences in question are evaluated. In this case the thought is that a point of evaluation provides something like an evaluative background on which an object x may “count” as being P .⁹

⁸ The notion of index is outlined in Kaplan (1989) and Lewis (1998).

⁹ Predelli (2005) suggests that this is the role to be assigned to the point of evaluation in order to account for the context-sensitivity of sentences such as ‘The leaves are green’.

One thing must be clear about this second option. To say that a difference of interpretation may lie in a difference of point of evaluation is to say that a difference of interpretation may not involve a difference of index. If an index gives us a “content” expressed by a sentence in a context, this means that different interpretations of the sentence are compatible with the same content. A faultless dispute can thus be described as one in which two parties do not really contradict each other by assertively uttering ‘*x* is *P*’ and ‘*x* is not *P*’, even though the two sentences express, as it were, contradictory contents. This is *not* what a relativist about truth should expect. Sometimes relativism is framed as a treatment of faultless disputes according to which the divergence between the two parties is located at the level of point of evaluation rather than at the level of index. The alleged advantage of such treatment over more conventional forms of indexicalism is that it preserves the apparent contradictoriness of contents.¹⁰ But the explanatory choice of the point of evaluation is not in itself relativist. To get relativism, it has to be assumed in addition that a difference of point of evaluation does not amount to a difference of interpretation. Or else, a faultless dispute turns out to be one in which the sentences uttered are absolutely true or false on different interpretations.¹¹

References

- García-Carpintero, M. (2008) ‘Relativism, Vagueness and What is Said’, this volume.
 Kaplan D. (1989) ‘Demonstratives’, in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes From Kaplan*, Oxford: Oxford University Press.
 Kölbel, M. (2002) *Truth without objectivity*, London: Routledge.
 ——— (2003) ‘Faultless Disagreement’, *Proceedings of the Aristotelian Society* 104, 53–73.
 ——— (2008) ‘Motivations for Relativism’, this volume.
 Lasersohn, P. (2005) ‘Context Dependence, Disagreement, and Predicates of Personal Taste’, *Linguistics and Philosophy* 28, 643–86.
 Lewis, D. (1998) ‘Index, Context, and Content’, in *Papers in Philosophical Logic*, Cambridge: Cambridge University Press.
 Predelli, S. (2005) ‘Painted Leaves, Context, and Semantic Analysis’, *Linguistics and Philosophy*, 28, 351–74.
 Stojanovic, I. (forthcoming) ‘Talking About Taste: Disagreement, Implicit Arguments and Relative Truth’, *Linguistics and Philosophy*.
 Wright, C. (2001) ‘On Being in a Quandary’, *Mind* 110, 45–98.

¹⁰ Kölbel (2003: § 7), Lasersohn (2005: §§ 2–4).

¹¹ García-Carpintero provides some elucidations about the notion of what is said that make this point vivid. Note that here the relativist has nothing to gain from the analogy with the case of ‘The leaves are green’. For in that case the point of evaluation is taken to account for what appears to be a difference of interpretation.

This page intentionally left blank

Presuppositions of Commonality: An Indexical Relativist Account of Disagreement

Dan López de Sa

In some philosophically interesting cases—including future contingents, predicates of personal taste, evaluative predicates in general, epistemic modals, and knowledge attributions—there seem to be possible contrasting variations in judgments about an issue in the domain that do not seem to involve fault on the part of any of the participants.

According to *relativism*, these appearances of faultless disagreement are to be endorsed. According to *moderate* relativism, this can be done within the general Kaplan–Lewis–Stalnaker two-dimensional framework, in which the basic semantic notion is that of a sentence s being true at a context c at the index i : it may in effect be the case that s is true at c (at its index i_c) but false at c^* (at i_{c^*}). According to *indexical* (moderate) relativism, this is so in virtue of the *content* of sentence s at c being different from that of s at c^* .

Indexical relativism thus seems to account straightforwardly for the *faultlessness* of the judgments that could be expressed by using s at c but not at c^* . What about the facts involving intuitions of *disagreement*, as revealed in ordinary disputes in the domain?

In recent debates on these issues, most presuppose that indexical relativism simply cannot account for these facts, and is thereby to be rejected. The arguments from disagreement that some provide, however, only target particularly simple versions of indexical relativism. My aim here is to show that they do not affect more complex versions thereof—more in particular, that they do not affect the view I have defended elsewhere,¹ which exploits presuppositions of commonality to the effect that the addressee is relevantly like the speaker. There are versions of indexical relativism, therefore, in a good position to account for the envisaged intuitions concerning disagreement.

In what follows, I unpack and argue for all these claims, in that order.

¹ See López de Sa (2003a, 2003b, 2007a), which elaborate on a suggestion by Lewis (1989).

14.1 FAULTLESS DISAGREEMENT AND RELATIVISM(S)

In some philosophically interesting cases—including future contingents, predicates of personal taste, evaluative predicates in general, epistemic modals, and knowledge attributions—there seem to be possible contrasting variations in judgments about an issue in a domain that do not seem to involve fault on part of any of the participants.

Take one of the cases where such an appearance is strongest. It does seem that Hannah and her wife Sarah can have contrasting judgments as to whether Homer Simpson is funny. Hannah may have a judgment that she could naturally express in a normal conversation by uttering ‘Homer is funny’ whereas Sarah may have a contrasting one she could express by uttering ‘Homer is not funny.’ As things may be, this divergence can be irremovable, not imputable to any lack of relevant information, attention, or the like—rather, to use Wright’s apt phrase, they may have just reached a point where, one would feel, they should just ‘agree to differ,’ as it were, without imputation of fault on either side. For, it seems, both may be judging what is true—and still disagree as to whether Homer is funny.

It is a *datum*, I take it, that in a lot of situations in domains like those mentioned there *seems* to be such cases of *faultless disagreement*.² This is not to say that such appearances need be endorsed in all such cases—only that rejecting them would require some ‘explaining away.’ The prospects of such explaining away might easily be dissimilar across the cases: what may be wildly implausible with respect to Homer being funny might turn out to be quite sensible as to whether Hannah knows after all that the bank will be open.³ Here I will not defend, with respect to any of the domains in question, that appearances of faultless disagreement should be endorsed. Rather, my aim is to dispel a worry that a certain kind of view is not capable of adequately endorsing them—*assuming* they are to be endorsed. For illustrative purposes, though, I will mainly focus in what I take to be the more favorable cases, such as those involving whether Homer is funny.

Following the lead of Crispin Wright (1992), one can conceive of *relativism* in general as this attempt to endorse appearances of faultless disagreement—for such need to involve, in one way or other, some relativity to contrasting features

² Some use ‘faultless disagreement’ in a more restricted sense, requiring that there be a single content or proposition which is contrastingly judged, see for instance Kölbel (2003). As we will see, according to this more restricted sense, it cannot just be taken as a *datum* for relativists and non-relativists alike that there are apparent faultless disagreements, nor do all versions of relativism endorse that there are in effect faultless disagreements in the relevant domains. These I take to favor my more liberal usage.

³ Or it may be plausible in just some but not all of the cases in a given domain. Notice that it is sufficient for relativism that *some* of the appearances of faultless disagreement, in the domain in question, are to be endorsed. Feldman (2001) gives examples concerning ‘wrong’ where some alternative explaining away might be more plausible. This being so does not contradict relativism, see Cohen (2001). See also Williamson (2005) for further discussion of similar examples.

of the subjects in question. And the different sources the relativity might be held to have are what gives rise to the different relativisms.

14.2 MODERATE RELATIVISMS: INDEXICAL VS NON-INDEXICAL

According to *moderate* relativism, endorsing appearances of faultless disagreement can be done within the general Kaplan–Lewis–Stalnaker two-dimensional framework, in which the basic semantic notion is that of a sentence *s* being true at a context *c* at the index *i*: it may in effect be the case that *s* is true at *c* (at its index *i_c*) but false at *c** (at *i_c**).

The jargon I adopt is from Lewis (1980). A *context* is a location—time, place, and possible world, or *centered world* for short—where a sentence is said. It has countless features, determined by the character of the location. An *index* is an *n*-tuple of features of context, but not necessarily features that go together in any possible context. Thus an index might consist of a speaker, a time before his birth, a world where he never lived at all, and so on. The coordinates of an index are features that can be shifted independently, unlike those of a context, and thus serve to represent the contribution of sentences embedded under sentence operators, such as ‘possibly’ or, more controversially, ‘somewhere,’ ‘strictly speaking,’ and so on. Given a context *c*, however, there is *the index of the context*, *i_c*: that index having coordinates that match the appropriate features of *c*. Given this uniqueness, the basic two-dimensional relation can be abbreviated in this special case: sentence *s* is true at context *c* iff *s* is true at context *c* at index *i_c*.⁴

Moderate relativism—or *contextualism*, as it has also sometimes been labelled in recent debates⁵—contrasts with what I label *radical* relativism, characterized by a departure from this classical two-dimensional framework. Radical relativism has it that it could be that a sentence *s* at context *c* at index *i* is true *from a certain perspective* but false from another—where perspectives are to be thought of as the same sort of thing as contexts, but representing a location from where a sentence, as said in a (possibly different) location, is viewed or assessed.⁶

⁴ Sentence *s* is here assumed to contain no ambiguous, vague expressions, and the like. Besides, unless otherwise specified I use ‘sentences’ for sentence-types, but ‘utterances’ for utterance-tokens.

⁵ Although in some of these debates, particularly that concerning knowledge attributions, the label is some other times reserved for a more specific view, which has it that they are features of the attributor in contrast with those of the addressee that are relevant for the truth-value of the attribution.

⁶ Radical relativists in this sense include MacFarlane (2003, 2005, *inter alia*) and Laserohn (2005). I propose to use ‘perspectives’ instead of MacFarlane’s ‘contexts of assessment,’ as I think it helps to avoid confusions with ‘context of use/utterance’ (‘context’ here) and, more importantly, with ‘circumstance/point of evaluation’ (‘index’ here).

But, moderate views have it, no such departure is required. Appearances of faultless disagreement can be endorsed, as I said, by its possibly being the case that in effect *s* is true at *c* (at its index i_c) but false at c^* (at i_{c^*}). According to *indexical* (moderate) relativism, this is so in virtue of the *content* of sentence *s* at *c* being different from that of *s* at c^* .

The *content* of sentence *s* at *c* is the function that assigns to each index *i* the value of *s* at *c* at *i*. These arguably model that with respect to which attitudes such as belief or desire are relations, and what sentences in context say, at least in one understanding of the locution. They can then be called *propositions*, although the term is also sometimes reserved for functions from *worlds* to truth-values. The two usages coincide if, but only if, worlds are the only coordinates of indices. This is indeed a possible view, but one on which I rather remain neutral—hence, following Lewis (1980) again, I opt for ‘contents.’

According to the alternative *non-indexical* (moderate) relativist view, the content of *s* at *c* is the same as at c^* , but it determines with respect to index i_c a different truth-value than with respect to the different index i_{c^*} . This requires that the relevant feature be a coordinate of indices, such as (uncontroversially enough) a possible world and (more controversially) a time, a standard of precision, and so on.

Quite frequently one finds a characterization of a view in terms of the contention that “propositional truth is relative,” in the sense that contents of sentences in context are true or false only relative to some non-standard feature (other than a world, time, standard of precision and so on).⁷ By itself, the contention is neutral with respect to non-indexical (moderate) relativism and radical relativism. The view that contents’ truth-values are relative to, say, senses of humor or moral codes or bodies of knowledge or what-have-you is by itself neutral as to whether it is the context that determine the appropriate features of indices—so that one preserves that a sentence at a context is true just in case it is true at the context at the index of that context—or rather they are determined by the perspective from which the statements are assessed.⁸ I suspect, however, in most cases, it is the former moderate view that would be intended, if distinguished from the latter radical one.⁹

For indexical relativism, however, the content of sentence *s* at *c* is different from that of *s* at c^* —hence that it may in effect be the case that *s* is true

⁷ See for instance Kölbel (2004), Stanley (2005: ch. 7), and Zimmerman (2007).

⁸ The present point is also made in MacFarlane (2005: 323–38). In more recent, yet unpublished work, MacFarlane has emphasized that not only the alleged characterization, but also considerations (including some by himself) aiming to favor radical relativist views in contrast with contextualist ones, are in fact neutral between radical relativism proper and (his label) “non-indexical contextualist” positions of the sort considered in the text.

⁹ The list of defenders of such non-indexical but moderate relativism might turn out to include also Egan, Hawthorne, and Weatherston (2005) and Egan (2007). Further discussion of this, as well as a defense of the general taxonomy of relativist positions employed here, are included in my still-in-progress López de Sa (2007b).

at c (at its index i_c) but false at c^* (at i_{c^*}). Indexical relativism thus seems to straightforwardly account for the *faultlessness* of the judgments that could be expressed by using s at c but not at c^* . What about the facts involving intuitions of *disagreement*, as revealed in ordinary disputes in the domain?

14.3 DISAGREEMENTS REVEALED IN ORDINARY DISPUTES

In recent debates on these issues, most presuppose that indexical relativism simply cannot account for these facts, and is thereby to be rejected. Here is Wright's particularly vivid voicing of the worry:

If [indexical relativism] were right, there would be an analogy between disputes of inclinations and the 'dispute' between one who says 'I am tired' and her companion who replies, 'Well, I am not' (when what is at issue is one more museum visit). There are the materials here, perhaps, for a (further) disagreement but no disagreement has yet been expressed. But ordinary understanding already hears a disagreement between one who asserts that hurt-free infidelity is acceptable and one who asserts that it is not. (Wright, 2001: 451)

Essentially the same point is made in Kölbel (2004) and Lasersohn (2005). I think that this is one of the most serious charges that indexical relativism faces.¹⁰ Actually, this seems to be a killer objection for particularly simple *individualistic* versions of indexical proposal. I illustrate this with a toy proposal. Take the view which is exhausted by:

- (0) For each context c , 'Homer Simpson is funny' has the content that is true (at index i) iff (with respect to i) the speaker of c is amused by Homer.¹¹

So if c is the actual context centered at Hannah's actual location, and c^* the one at her wife Sarah's, then 'Homer Simpson is funny' at c has the content that is true with respect to the indices where Hannah is amused by Homer, but at c^* has the content that is true with respect to the indices where Sarah is amused by Homer. Thus 'Homer Simpson is funny' is true at c and false at c^* . But, insofar as (0) is concerned, this is so in as much and exactly the same sense in

¹⁰ The other main one concerns facts involving intuitions about indirect attitude and speech reports, see for instance Hawthorne (2004). I will not be concerned with this—arguably related but—different objection here, see for further discussion Ludlow (2005) and López de Sa (2007b).

¹¹ Of course, independently of the present issue about disagreement, this is merely a caricaturesque toy proposal. To begin with, things might be clearly funny even in situations where one is not amused by them, as a result of lack of attention, preoccupation, and so on—so that, at the very least, any remotely plausible proposal along these lines would replace *being amused* by *being disposed to be amused*, under certain specified conditions. In order to stress the relevant differences with respect to the involved subjects among the different proposals, I would abstract from these issues here, and stick to toy proposals such as (0).

which ‘I am tired’ is (we may suppose) true at c but false at c^* . However, as Wright observes, ordinary understanding would typically hear a disagreement in the former case but not in the latter, as Hannah and Sarah disagree as to whether Homer is funny, but not as to whether one of them is tired. I agree. Thus a proposal exhausted by something along the lines of (0) fails to capture our intuitions about disagreement.

14.4 TWO COMMUNITARIAN INDEXICAL PROPOSALS

One response sometimes made on behalf of indexical relativism is that the considered problem arises due to the individualistic character of the simple versions of indexical relativism such as the one considered. Hence, the thought would go, one would avoid it by replacing in proposals such as (0) the individualistic mention of the speaker of the context by a reference to a community or group, salient in the conversation taking place at the context, to which the speaker of the context most plausibly belongs. This seems to be the proposal developed in von Fintel and Gillies (2005). To illustrate, consider:

- (1) For each context c , ‘Homer Simpson is funny’ has the content that is true (at index i) iff (with respect to i) the community conversationally salient in c is amused by Homer.

Unfortunately, this by itself does not seem to improve upon (0) regarding the account for intuitions concerning disagreement. For the community conversationally salient in the context of one participant in a conversation might be precisely relevantly different from the community conversationally salient in the context of another participant—and this might be the very topic of the conversation:

Argentinean Pablo: We greet colleagues by kissing them once.

British John: Oh, we don’t. We just shake hands.

Spanish Pedro: Well, we kiss them twice—unless the two of us are men, in which case we shake hands.

As in Wright’s previous case concerning ‘I am tired,’ there might be the materials here, perhaps, for a (further) disagreement, but ordinary understanding does not hear a disagreement has yet been expressed.

So, provided that the relevant communities salient at c and c^* are relevantly different, a proposal along the lines of (1) would indeed secure that ‘Homer Simpson is funny’ has at c and c^* appropriately different contents, and thus that ‘Homer Simpson is funny’ is true at c and false at c^* . But again, insofar as (1) is concerned, this is so in as much and exactly the same sense in which ‘We greet colleagues by kissing them once’ is (we may suppose) true at c but false at c^* . However, as observed, ordinary understanding would typically hear

a disagreement in the former case but not in the latter, as Hannah and Sarah disagree as to whether Homer is funny, but not as to whether they greet colleagues by kissing them once. Thus a proposal along the lines of (1) fails to capture our intuitions about disagreement.

The problem has been that the groups as well as the individuals salient in the different contexts involved in a given conversation might be relevantly *dissimilar*, which fits nicely with the indexical relativist account of the faultlessness of the judgments, but does not put indexical relativism in a position to account for intuitions of disagreement just by going communitarian instead of individualistic. This seems to be so, at least, unless one amends the proposal by *requiring* that the group salient in a context include all the participants of the conversation. This would give a proposal of the following kind:

- (2) For each context *c*, ‘Homer Simpson is funny’ has the content that is true (at index *i*) iff (with respect to *i*) the community conversationally salient in *c* includes the participants of the conversation and is amused by Homer.

This is in effect the kind of ‘Single Scoreboard Semantics’ view recently defended by DeRose (2004) in the domain of knowledge attributions.

I am ready to grant that this improves upon the rest of proposals considered so far regarding the account of intuitions of disagreement. But the main problem with it is that this is done at the cost of jeopardizing the indexical relativist account for intuitions of *faultlessness*, as presented above. For what if the participants in a given conversation are *not* relevantly alike? That is, what if there is no community conversationally salient in *c* and *c** which includes both Hannah and Sarah? It depends on the details of the proposal—and the semantics one favors for sentences containing singular definite descriptions with non-satisfied matrices. But in such a case either ‘Homer is funny’ is *false* in both *c* and *c**, or it *lacks a truth-value* in both *c* and *c**—the latter being the option favored by DeRose himself, hence his labeling it *the GAP view*. In either case, if Hannah were to utter ‘Homer is funny’ she would be thereby asserting something that is *not true*. As truth seems to be the weakest norm for assertion, it is very hard to see any sense in which Hannah’s is faultless.¹² Therefore this kind of proposal does not seem to account appropriately for faultlessness.¹³

¹² In more recent, yet unpublished, work DeRose emphasizes the intuitions we have about cases in which the relevant subjects are not participants of one conversation—nor can (easily) become such. Although I think this emphasis is appropriate (see the discussion in Section 6), I understand that it does not suffice for vindicating our intuitions of faultlessness—as they seem to be present, and with equal strength, with respect to (actual or conceivable) cases where a conversation does take place.

¹³ This seems not to be the only problem for such a proposal. It is (arguably) very implausible that the *truth-value* of my utterance of ‘Homer Simpson is funny,’ as opposed to (say) its felicity or appropriateness, is affected by a change in *your* sense of humor—even if we are talking together. For an elaboration of a criticism along these lines, see Feldman (2004).

14.5 THE PRESUPPOSITION OF COMMONALITY APPROACH

The worry concerning disagreement jeopardizes a particularly simple individualistic version of indexical relativism. I have just considered two proposals that try to overcome it by going communitarian, and found them faulty. I want to argue that a more complex—even if equally individualistic—version of indexical relativism is capable of neatly accounting for intuitions of disagreement as revealed by ordinary disputes on the matter. The main complexity comes from the fact that, unlike those considered, the proposal is not exhausted by a claim about the *content* of the sentences in question at the different contexts, but crucially also comprises another component, concerning the *presupposition* that some expression in it triggers. This *presupposition of commonality approach*, which I have also defended in other works,¹⁴ elaborates on a suggestion due to David Lewis in his discussion of relativism about values, where he considers a similar objection from disagreement:

Wouldn't you hear them saying 'value for me and my mates' or 'value for the likes of you'? Wouldn't you think they'd stop arguing after one speaker says X is a value and the other says it isn't?—Not necessarily. They might always presuppose, with more or less confidence (well-founded or otherwise), that whatever relativity there is won't matter in *this* conversation. (Lewis, 1989: 84)

As I just said, I propose to elaborate this suggestion *via* the contention that the relevant expression involved triggers a presupposition of commonality to the effect that the participants of the conversation are all alike in the relevant respects. To illustrate, consider:

- (3) (a) For each context *c*, 'Homer Simpson is funny' has the content that is true (at index *i*) iff (with respect to *i*) the speaker of *c* is amused by Homer.
- (b) 'is funny' triggers the presupposition that the participants in the conversation are similar with respect to humor.

I assume here Stalnaker's account of speaker presuppositions and the derived notion of expressions triggering presuppositions. Here is a recent statement of his views:

Acceptance . . . is a category of propositional attitudes and methodological stances towards a proposition, a category that includes belief, but also some attitudes (presumption, assumption, acceptance for the purposes of an argument or an inquiry) that contrasts

¹⁴ In López de Sa (2003a, 2003b, 2007a). See also Egan (2007) for what seems to be a similar move—although it is not completely clear to me on behalf of exactly which position this is made, see López de Sa (2007b).

with belief, and with each other. To accept a proposition is to treat it as true for some reason. One ignores, at least temporally, and perhaps in a limited context, the possibility that it is false. It is common ground that ϕ in a group if all members *accept* (for the purposes of the conversation) that ϕ , and all *believe* that all accept ϕ , and *believe* that all *believe* that all accept ϕ , etc. The speaker presuppositions [are] the speaker's beliefs about the common ground. A *nondefective* [conversation] is a [conversation] in which the participants' beliefs about the common ground are all correct. Equivalently, a nondefective [conversation] is one in which all of the parties to the conversation presuppose the same things. (Stalnaker, 2002: 716–17)¹⁵

A given *expression* triggers a certain presupposition if an utterance of it would be infelicitous when the presupposition is not part of the common ground of the conversation—unless participants accommodate it by coming to presuppose it on the basis of the fact that the utterance has been produced.

Now consider Hannah and Sarah once again. The presupposition of commonality approach of (3) does not differ from the simple individualistic proposal of (0) with respect to the component about the content of the sentence at the different contexts. Hence the explanation of faultlessness is as straightforward as it was there—it can indeed be the case that 'Homer Simpson is funny' is true at c and false at c^* , and therefore that both Hannah and Sarah are (faultlessly) judging what is true. But, unlike the case of (0), the approach of (3) is not exhausted by the component about content. And this is what accounts for our intuitions about disagreement concerning the case. According to the approach, 'is funny' triggers a presupposition of commonality to the effect that both Hannah and Sarah are similar with respect to humor. Thus, in any non-defective conversation where Hannah uttered 'Homer is funny' and Sarah replied 'No, it is not,' it would indeed be common ground that Hannah and Sarah are relevantly alike, and thus that they are contradicting each other. After all, and provided they are alike, either both Hannah and Sarah are amused by Homer or they are not. Hence why, unlike in the case of 'I am tired,' ordinary understanding would indeed hear a disagreement as to whether Homer is funny.

Thus, I submit, the version of indexical relativism constituted by this presupposition of commonality approach seems to neatly account for both intuitions of faultlessness and intuitions of disagreement. I want to end by considering some objections that this is indeed so, which would clarify the prospects of the proposal further.

¹⁵ In the last two sentences of the quotation, Stalnaker originally speaks of a conversation's 'context' (in earlier works, 'context set', as it can be seen (when non-defective) as the set of worlds in which the relevant presuppositions are indeed true). I have substituted 'context' by 'conversation' here, in order to avoid confusion with the Lewisian usage I am adopting, which reserves 'context' for worlds centered at spatiotemporal points representing locations where a sentence could be said. Thus one conversation (and its "context set") involves—and indeed is determined by features of—many different contexts, corresponding to the locations of the different participants at the different stages of the conversation.

14.6 OBJECTIONS

- (A) *The presuppositional approach of (3) accounts merely for the appearance of disagreement, as opposed to genuine disagreement proper, in that it only explains why Hannah and Sarah would take themselves to disagree.*

Reply: Well, the proposal *does* account for intuitions of (“genuine”) disagreement (“proper”) as revealed in ordinary disputes, to the extent to which these are captured by something along the lines of the following:

DISAGREEMENT In any conversation, it will be common ground that utterances of (say) ‘Homer is funny’ and ‘Homer is not funny’ would contradict each other.¹⁶

In any conversation where the presupposition of commonality is fulfilled, the utterances will in fact contradict each other. And in any non-defective conversation it is common ground that such a presupposition is in fact fulfilled. This seems to me to be much stronger than mere appearance of disagreement.

- (B) *Maybe so, but the intuitions of disagreement are anyway stronger than that. They are intuitions to the effect that in any conversation whatsoever any such pair would in fact contradict each other—regardless of what participants think about each other. And that is something that (3) does not account for.*

Reply: According to the objector here, an indexical relativist proposal such as (3) would violate the following:

CONTRADICTION In any conversation whatsoever, it will be the case that utterances of (say) ‘Homer is funny’ and ‘Homer is not funny’ would in fact contradict each other.

Of course, (3) does not respect *this*. But, unlike the intuitions of DISAGREEMENT, I think it is quite dubious that there exist the unrestricted *intuitions* of CONTRADICTION. Of course, the fact that CONTRADICTION cannot be motivated by appealing to the intuitions revealed by ordinary disputes on the matter does not generally guarantee that it cannot be motivated *in some other way*. But such an independent case is still to be provided, as the only argument from disagreement actually given against indexical relativism seems to be the one involving intuitions I have been considering in this paper. (Furthermore, as observed previously in connection with the norm for assertion, this case would then amount to a case against faultlessness. So if any such case could be substantiated—which I very much doubt, at least

¹⁶ I take worlds containing the relevant non-defective conversations to be closer. If this is rejected, the subjunctive conditional should be amended by adding the qualification that the utterances be felicitous.

concerning cases like that of ‘the funny’—then it is *relativism* in general, no matter whether radical or moderate, indexical or non-indexical, which would be—*would* be—in trouble.¹⁷)

- (C) *The problem with [a proposal along the lines of (3)—DLdS] is that it explains only intra-conversational disagreement, leaving inter-conversational disagreement unaccounted for. This is not a stable resting point. Once the importance of accounting for disagreement has been conceded, one cannot limit oneself to disagreement within conversations. . . . [W]hen I overhear a group of ten-year-olds chattering about how “funny” certain knock-knock jokes are, I may think that they are wrong. These jokes just aren’t that funny. But the kids certainly don’t think of themselves as involved in a conversation with me—they may not even know I’m there. Nor do I think of myself as conversing with them.* (MacFarlane, 2007)

Reply: I agree that ordinary understanding would hear a disagreement being present here (see below). I am also inclined to think that, in some such cases at least, one would also regard the relevant pair of unconditionalized utterances as contradicting each other, despite the fact that the subjects are not participating in any *actual* conversation. But, I submit, this can be accommodated within the framework of the presupposition of commonality approach, to the extent that these cases are those where there are *possible* (close, easily actualizable) conversations, which would involve the people in question. Thus, although not actually conversational partners, they could (easily) become such.¹⁸ What the cases illustrate, then, is that the intuitions of DISAGREEMENT should not be understood as restricted to actual conversations.

- (D) *Surely there is an intuitive sense in which people disagree in the envisaged scenarios. And this sense is available even if, for whatever reason, participants do not presuppose that the others are alike—but may indeed come to presuppose otherwise.*

Reply: Agreed. Hannah and Sarah might disagree as to whether Homer is funny, and their respective distinctive senses of humor be perfectly apparent to them. But again I do not think this goes against the approach of (3)—rather it can be seen as precisely illustrating further the sense in which our intuitions support the appropriately weaker DISAGREEMENT as opposed to CONTRADICTION. Intuitively, I submit, the disagreement in our case is

¹⁷ But isn’t precisely *radical* relativism different here, in that it can accommodate both CONTRADICTION and faultlessness? I don’t think so, but I will not try to defend my view here—see again López de Sa (2007b). For responding to the objection, however, it is enough to realize that intuitions of disagreement as revealed by ordinary disputes in the domain only support something like DISAGREEMENT, and not CONTRADICTION.

¹⁸ What about cases where a conversation is not only not actual, but simply cannot take place—or, at the very least, not easily? “When I was ten, I used to go around saying ‘fish sticks are delicious’ (and meaning it!). Now I say ‘fish sticks are not delicious.’ It seems to me that I disagree with my past self. But I am not involved in a *conversation* with my past self” (MacFarlane, 2007). *Imagined*, or easily imaginable, conversations might play in some such cases the envisaged role—but see the discussion below.

constituted by the contrastive features of Hannah's and Sarah's senses of humor (say). In non-defective conversations where they presuppose they are alike, this disagreement would be naturally expressible by the relevant pair of (unqualified) contrasting utterances, which would be regarded as contradicting each other—and would in fact contradict each other, provided the presupposition is fulfilled. But in equally non-defective conversations where they do not presuppose they are alike, but may indeed presuppose they are not, their disagreement exists all the same, but it need not be so expressible. Rather, the prediction is that they could naturally express it by utterances which conditionalize on them or their relevant feature, thus canceling the presupposition of commonality. They would not contradict each other—but felicitously express their disagreement as to whether Homer is funny all the same. And this is—or so it seems to me—quite plausible with respect to our case:

Hannah: Oh, they're showing The Simpsons! Come and watch, Sarah! Homer is so funny!

Sarah: Funny for you, darling. You should remember that he doesn't amuse me at all.

CONCLUSION

In recent debates on these issues, most presuppose that indexical relativism simply cannot account for these facts, and is thereby to be rejected. The arguments from disagreement that some provide, however, only target particularly simple versions of indexical relativism. My aim here has been to show that they do not affect more complex versions thereof—more in particular, that they do not affect the view I have defended elsewhere, which exploits presuppositions of commonality to the effect that the addressee is relevantly like the speaker.

This, of course, does not amount to a full defense of the presupposition of commonality approach, nor of indexical relativism, in any of the domains mentioned. Realist or invariantist attempts to explain away the appearances of faultless disagreement have not been considered, no remotely plausible specific proposal concerning the content component has been mentioned, nor has the case for the presence of a further presuppositional element been substantiated.

These points notwithstanding, I hope I have made the case that there are versions of indexical relativism that are in a good position to account for the envisaged intuitions concerning disagreement.¹⁹

¹⁹ I am greatly indebted to the participants of the Epistemology, the Relativism, and the Vagueness Arché Seminars and to the members of LOGOS for discussions on the topic of this paper—and to Eline Busck, Richard Dietz, Philip Ebert, Manuel García-Carpintero, Patrick Greenough, Carrie Jenkins, Max Kölbel, Paula Milne, Sebastiano Moruzzi, Julien Murzi, Pablo

References

- Cohen, S. (2001). 'Contextualism Defended: Comments on Richard Feldman's "Skeptical Problems, Contextualist Solutions"'. *Philosophical Studies*, 103: 87–98.
- DeRose, K. (2004). 'Single Scoreboard Semantics'. *Philosophical Studies*, 119: 1–21.
- Egan, A. (2007). 'Epistemic Modals, Relativism and Assertion'. *Philosophical Studies*, 133: 1–22.
- , Hawthorne, J., and Weatherson, B. (2005). 'Epistemic Modals in Context', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*. Oxford: Oxford University Press.
- Feldman, R. (2001). 'Skeptical Problems, Contextualist Solutions'. *Philosophical Studies*, 103: 61–85.
- (2004). 'Comment on DeRose's "Single Scoreboard Semantics"'. *Philosophical Studies*, 119: 23–33.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Kölbel, M. (2003). 'Faultless Disagreement'. *Proceeding of the Aristotelian Society*, 104: 53–73.
- (2004). 'Indexical Relativism versus Genuine Relativism'. *International Journal of Philosophical Studies*, 12: 297–313.
- Lasersohn, P. (2005). 'Context-Dependence, Disagreement, and Predicates of Personal Taste'. *Linguistics and Philosophy*, 28: 643–86.
- Lewis, D. (1980). 'Index, Context, and Content', in S. Kanger and S. Öhman (eds.), *Philosophy and Grammar*. Dordrecht: Reidel. (Reprinted in *Papers in Philosophical Logic*, Cambridge: Cambridge University Press, 1998 (q.v.).)
- (1989). 'Dispositional Theories of Value'. *Proceeding of the Aristotelian Society* (suppl. vol.), 63: 113–38. (Reprinted in *Papers in Ethics and Social Philosophy*, Cambridge: Cambridge University Press, 2000 (q.v.).)
- López de Sa, D. (2003a). '(Indexical) Relativism about Values: A Presuppositional Defense'. *MS*.
- (2003b). *Response-Dependencies: Colors and Values*. Barcelona: Universitat de Barcelona, Ph.D. Dissertation.
- (2007a). 'The Many Relativisms and the Question of Disagreement'. *International Journal of Philosophical Studies*, 15: 269–79.
- (2007b). 'The Many Relativisms: Index, Context, and Beyond'. *MS*.
- Ludlow, P. (2005). 'Contextualism and the New Linguistic Turn in Epistemology', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*. Oxford: Oxford University Press.
- MacFarlane, J. (2003). 'Future Contingents and Relative Truth'. *Philosophical Quarterly*, 53: 321.

- MacFarlane, J. (2005). 'Making Sense of Relative Truth'. *Proceeding of the Aristotelian Society*, 105: 321–39.
- (2007). 'Relativism and Disagreement'. *Philosophical Studies*, 132: 17–31.
- Stalnaker, R. (2002). 'Common Ground'. *Linguistics and Philosophy*, 25: 701–21.
- Stanley, J. (2005). *Knowledge and Practical Interest*. Oxford: Oxford University Press.
- von Fintel, K., and Gillies, A. S. (2005). '“Might” Made Right'. *MS*.
- Williamson, T. (2005). 'Knowledge, Context, and the Agent's Point of View', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*. Oxford: Oxford University Press.
- Wright, C. (1992). 'Truth and Objectivity'. Cambridge, Mass.: Harvard University Press.
- (2001). 'On Being in a Quandary'. *Mind*, 110: 45–98. (Reprinted in *Saving the Differences*, Harvard University Press, 2003 (q.v.).)
- Zimmerman, A. (2007). 'Against Relativism'. *Philosophical Studies*, 133: 313–48.

Index

- acceptance of a relativistic proposition, *see*
relativistic proposition
- actuality 85, 98–101, 119–23
- agreement 232–3
- Aristotle 45, 46, 58
- assertion 9–10; *see also* truth
correctness of 225, 228–9
norms of 226, 229–30
of a relativistic proposition, *see* relativistic
proposition
and truth 225, 228, 232
- assessment sensitivity 90–3
- Austin, J. 42, 61, 132, 135
- Ayer, A. J. 226
- Bach, K. 22 n. 37, 132, 137, 269 n. 10
- Barker, C. 152 n. 31
- Barwise, J. 47, 49, 54, 55, 57, 58
- belief in a relativistic proposition, *see* relativistic
proposition
- Belnap, N. 82, 83 n., 85, 86 n., 87, 99 n. 19
- Bennigson, T. 208 n. 5
- Bilgrami, A. 207 n. 2
- Blackburn, S. 236
- Boghossian, P. 168 n., 208 n. 5
- Bonnay, D. 32 n., 33, 104, 111 n., 113, 114
- Borg, E. 2 n. 2
- Borges, J. L. 81
- Brandom, R. 142, 148, 210
- Brännmark, J. 60
- Burnyeat, M. 207 n. 1
- Cappelen, H. 2, 11 n. 17, 26 n., 34, 131 n.,
175 n. 22, 188 n. 4, 208 n. 5, 216 n. 25,
245 n. 15, 260 n., 266, 267, 268, 269
n. 10, 271 n. 11, 276 n. 16
- Carston, R. 2 n. 2
- Chalmers, D. 3 n., 8 n. 12
- circumstance 42–3, 66, 68–72, 75–6
- cognitive limitation 182
- Cohen, S. 22 n. 38, 298 n. 3
- Condoravdi, C. 245 n. 17
- content 295
see also perspectival content
relativism 270–2
representational 168, 170, 173
of the sentence 47–9; *see also* *lekton*
truth-conditional 164
of the utterance 44–5, 48; *see also* situation
of evaluation
- context 65, 68–76, 294, 295, 299
of assessment 70, 90–1, 161, 164, 166–8,
172, 174, 180
sensitivity 2, 6, 13–14
shifting 110–15, 119, 122, 126
of use 83
- contextualism 123–7, 299
about epistemic modals 242–3
indexical 11–16
non-indexical 300
non-standard 164
in pragmatics 25–7
- contingency 136
- Davidson, D. 164 n. 10, 207 n. 2
- DeRose, K. 21 n., 22 n. 38, 197 n. 12, 242
n. 8, 255 n. 32, 303
- de se* 139–40
- Dietz, R. 178 n. 25, 203 n., 222 n., 223 n.,
245 n. 17
- disagreement 279
contextualist explanations of 283–4
correct 247
faultless 57–61, 129, 137, 193, 210, 226,
230–3, 287, 288, 289, 290, 293, 298,
304–5
impossibility of explanation 218–21
invariantist explanations of 281–3
relativistic 211
- disquotational schema 95
- Dokic, J. 104, 118, 123–4
- double indexing 6
- Dreier, J. 25 n. 43
- Dummett, M. 46–7, 182 n. 30
- Edgington, D. 119 n.
- Egan, A. 22 n. 36, 56, 57, 63 n., 130, 146–9,
178 n. 25, 187 n. 1, 193 n. 9, 197 n. 13,
208 n. 6, 216 n. 23, 242 n. 6, 244, 247
n. 22, 266 n. 3, 285 n., 300 n. 9, 304 n.
- Egré, P. 32 n., 33, 104, 107 n., 111 n., 113,
114, 118, 123, 124
- Einheuser, I. 31 n. 58, 33–4, 244
- epistemic modals 197–9; *see also*
contextualism; relativism
contextualism 241–3

- epistemic modals (*cont.*)
 egocentric uses 244, 255
 epistemic possibility 240–1
 epistemic situation 240–1
 evidentiality 246
 ignorant assessor scenarios 250–2
 indexical contextualism 259–60
 knowing assessor scenarios 247–50
 non-factuality 246
 non-indexical contextualism 259–60
 relativism 243–4
 speaker inclusion 242–3
- Evans, G. 20 n. 32, 32, 43–6, 47 n., 130, 139, 141–2, 14–8, 151–2, 209 n. 11, 231 n. 10, 232
- Evans' anti-relativist argument 142, 148
- expressivism 226, 236
- fact-relativity 208 n. 5
- Fara, D. G. 105, 107
- Feldman, R. 298 n. 3, 303 n. 13
- Field, H. 30 n. 52
- Fine, K. 31 n. 55, 150 n. 30
- Fintel, K.v. 22 n. 36, 59 n., 240 n., 241 n. 4, 242 n. 5, 243 n. 10, 245 n. 17, 246 n. 19, 257 n. 4
- Frege, G. 1, 2, 41, 43, 45, 47, 49, 225–6, 227, 228
- future
 contingent 81, 93–4
 open 81
- García-Carpintero, M. 20 n. 32, 33, 44 n., 132 n. 4, 137 n. 13, 150 n. 30, 295 n. 11
- Geach, P. 44, 226
- Gibbard, A. 219 n. 32
- Gillies, A. 22 n. 37, 59 n., 61, 241 n. 4, 242 n. 5, 243 n. 10, 245 n. 17, 246 n. 19, 257 n. 34, 302
- Glanzberg, M. 245 n. 15
- Gomez-Torrente, M. 105, 107
- gradable adjectives, context-dependence of 145, 149–50
- Graff, D. 126 n., 146, 149, 150
- Green, M. 86 n.
- Grice, H.P. 20 n. 31, 131–2
- Hacking, I. 255 n. 32
- Hales, S. 208 n. 5
- Halpern, J. 113
- Harman, G. 25
- Hawthorne, J. 22 n. 36, 63 n., 130, 144 n. 21, 146–9, 159–60, 178 n. 25, 187 n. 1, 247 n. 23, 266 n. 3, 285 n., 300 n. 9, 301 n. 10
- Higginbotham, J. 139
- Hintikka, J. 33 n., 44, 46, 110, 113, 125
- Iacona, A. 35, 218 n. 28
- illocutionary force 9–10
- index 83, 299
- indexicality 63–5, 67–71, 73–6, 108, 113, 114, 115, 118, 121, 136
- initial redundancy 98
- intension 5
- intensionality 63–4, 66–71, 75–6
- interpretation 70–1
 function 65–6
 semantic 64, 67, 69, 75
- introspection 107, 110–13, 118
- Jordan, J. 166 n. 12
- Kamp, H. 104, 113, 119–20, 125–6
- Kaplan, D. 1, 2, 6–8, 10, 14 n. 21, 32, 42–3, 48 n., 55 n., 64, 65 n. 2, 66–74, 76, 82–4, 95 n. 14, 98, 111 n., 129, 136, 138–40, 143–4, 148, 157, 201 n. 21, 245 n. 14, 279, 282, 294 n. 8, 297, 299
- Kennedy, C. 143, 144 n. 20, 149, 150 n. 28
- King, J. 3 n., 42 n., 18, 139–41, 267
- knowledge:
 factivity of 214–15
 higher-order 107, 110–13, 118, 123–5
 perceptual 108, 110–13, 116, 122, 123–5
- Kölbel, M. 58, 59, 68 n. 6, 105 n., 141, 150 n. 30, 169 n., 226–31, 233, 266 n. 3, 287 n., 298 n. 2
- Kompa, N. 24 n. 40
- Kratzer, A. 242 n. 5
- Kripke, S. 66 n., 106, 137
- Lasersohn, P. 19 n. 29, 58–9, 60, 208 n. 6, 212 n., 226, 230 n., 231 n. 8, 243 n. 11, 265 n. 1, 266, 279–83, 287 n., 299 n. 6, 301
- lekton* 47–8, 59–61
- Lepore, E. 2, 11 n. 17, 26 n., 131 n., 175 n. 22, 188 n. 4, 266–8, 271 n. 11
- Lewis, D. 1, 2 n. 2, 3 n., 6, 8, 22 n. 38, 30 n. 52, 65 n. 2, 70, 82, 83, 85, 87, 98, 99 n. 20, 101, 131, 133, 135 n. 8, 136–42, 144, 148, 157, 294 n. 8, 297, 299–300, 304
- linguistic evidence 179, 182
- logical consequence 84
- logical implication 84
- López de Sa, D. 15 n. 23, 35, 297 n., 300 n. 9, 301 n. 10, 304 n., 307 n. 17
- Ludlow, P. 301 n. 10
- luminosity 118

- McDowell, J. 8 n. 13
- MacFarlane, J. 4 n., 9 n. 14, 17 n. 25, 18, 19
 n. 29, 20, 22 n. 34, 24, 27 n., 29–30,
 31 n. 56, 32, 33 n., 34, 44–5, 48 n., 58
 n. 10, 63 n., 69–70, 82, 86 n. 7, 92 n.,
 93 n., 97 n., 129–30, 135, 141–2, 145
 n. 22, 148, 149 n. 27, 161 n. 7, 163 n.,
 171, 173 n., 174 n., 178 n., 179–80, 187
 n. 1, 197 n. 14, 200 n. 18, 208–11,
 212 n., 217 n. 27, 219, 221, 223, 227–8,
 231 n. 10, 241 n. 4, 242–4, 245 n. 15,
 247, 248 n., 249 n. 28, 252–3, 257 n. 34,
 259 n. 38, 260, 265 n. 1, 266, 273,
 274–8, 279 n. 21, 281, 284, 285 n., 299
 n. 6, 300 n. 8, 307
- McGee, V. 88 n. 9
- McGinn, C. 53–4
- Machina, K. F. 88 n. 8
- McLaughlin, B. 88 n. 9
- margin for error principle 103–27
- Mates, B. 44
- Meiland, J. W. 16 n., 208 n. 3
- modalities, historical 85
- moderate relativism 10, 18, 48, 130, 142, 299
 weak 49
 strong 49
- Moltmann, F. 219 n. 31
- Moruzzi, S. 34, 184 n., 216 n. 25
- natural language semantics 6–11
- necessity
 historical 84
 logical 84
- Newton-Smith, W. 207 n. 2, 208 n. 3
- non-cognitivism 226, 236
- non-indexicalism 244–5
- non-representationality conjecture 170–1
- operators 31–2
- parameter 287
- Percival, P. 209 n. 11
- Perry, J. 3 n., 42, 55 n., 137, 138 n., 175 n.
- perspectival content
 asserting 191–2
 evaluating 195, 199
 modeling 191
- perspectival facts 190–1, 198
- perspective 227–8, 235–6, 299
- pluralism 267–8
 and semantic content 269–70
- pluralistic content relativism 266
 and contextualism 277–8
 and faultless disagreement 280
 and scepticism 275–6
 and semantic blindness 284
 and stability data 273–5
 and variability data 273–5
- point of evaluation 83
- possibility, historical 85
- practical conflict 234
- Predelli, S. 2 n. 2, 6, 7 n. 11, 8 n. 12, 27 n. 46,
 30 n. 53, 32, 65 n. 2, 69, 76 n., 145
 n. 22, 294 n. 9
- presupposition 304–5
- Price, H. 30 n. 51, 221 n., 246 n. 20
- Prior, A. 2, 44, 46–7, 58, 82, 138, 140 n. 17,
 142
- proposition 1–2
 Austinian 47, 49–50
 de se 2–3
 relativistic, *see* relativistic proposition
 temporal 41–4
- Protagoras 157, 166 n. 12, 207
- Protagorean tradition 158, 165
- Rabinowicz, W. 113, 117, 122 n., 124
- realism 288, 289, 290
- Recanati, F. 2 n. 2, 42, 50 n. 7, 111 n., 175
 n. 22
- Reimer, M. 272
- relativism 270–2, 298
 about content, *see* content relativism
 definition of 4
 about epistemic modals 21–2, 243–4
 factual 188, 202
 and fiction 30
 folk 158
 and future contingents 27–30
 ground-level 158
 indexical 188, 288, 289, 290, 300
 about knowledge attributions 22–4
 about matters of taste 16–20
 moderate, *see* moderate relativism
 moral 233–5
 about moral value 24–5
 see also pluralistic content relativism
 about probability 30
 propositional 188, 200
 radical 18, 44–5, 58, 130, 142, 147, 299
 self-refutation of, *see* self-refutation of
 relativism
 temporal, *see* temporalism
 about Travis cases 25–7
 about truth 163–5
 about vagueness 30–1
 weather-report 174–5
- relativistic proposition
 acceptance of 209, 213
 assertion of 210, 217–22
 belief in 216–22

- representation 168–72, 200–2
 perspectival 169
 Richard, M. 31 n. 56, 32, 34, 42, 50–9,
 129–30, 140 n. 17, 145–9, 151–3, 208
 n. 6, 217 n. 26, 272
 Rosenkranz, S. 34, 218 n. 30
 Rysiew, P. 22 n. 38
- Salmon, N. 132 n. 3
 Schiffer, S. 135 n. 10
 Schlenker, P. 14 n. 21
 Segerberg, K. 113, 117, 119, 122 n., 124
 self-refutation of relativism
 classical argument 207
 operative argument 221–2
 semantic confusion 150–1
 Shapiro, S. 126 n., 227
 Siegel, H. 208 n. 3
 situation of evaluation 49, 59–61
 Soames, S. 126 n., 260, 266, 267 n., 268, 269
 n. 10, 270, 271 n. 12
 Socrates 41, 44, 166 n. 12
 Sorensen, R. 103
 sorites 108, 125–7
 speech reports 13–14
 Sperber, D. 2 n. 2
 Stalnaker, R. 1, 3 n., 6 n., 56, 135 n. 8, 137,
 138 n., 285 n., 297, 299, 304–5
 standards of the community 61
 standard of taste 17
 possession of 19
 Stanley, J. 2 n. 2, 50 n. 6, 135 n. 8, 144 n. 21,
 149 n. 27, 159–60, 175 n. 22, 214 n. 16,
 300 n. 7
 Stephenson, T. 178 n. 25, 243, 247 n. 22
 Stojanovic, I. 8 n. 12, 32, 60, 65 n. 3, 68 n. 7,
 252 n. 31, 291 n.
 supervaluations 86–9
 Swanson, E. 246 n. 21, 257 n. 35, 258 n.
 Swyer, C. 166 n. 12, 208 n. 3
 Szabo, Z. 50 n. 6
- Taschek, W. 227
 Teller, P. 241 n. 4
 temporalism 2, 41, 50, 138
 moderate 51
- Thomason, R. H. 82, 84 n. 4, 86–7
 traditional misgiving 66
 Travis, C. 2 n. 2, 25–6, 27 n. 46, 144–5
 truth 133
 absolute 228, 235
 and assertion 141, 225, 228, 232
 determinate 96–8
 disquotational 147
 logical 84
 of occurrences 8
 propositional 93–6, 227
 relative, *see* relativism about truth
 in semantics 5–7
 of utterances 1, 93–4
 truth-functionality 87–8
 two-dimensional semantics 113–15, 119–23
- utterance, time of 44–7, 49
- vagueness
 epistemicist theory of 126
 higher-order 112–13
 validity 63–6, 71–7
 Veltman, F. 246 n. 21
- Weatherston, B. 22 n. 36, 63 n., 130, 146–9,
 178 n. 25, 187 n. 1, 266 n. 3, 285 n., 300
 n. 9
 what is said 131
 Williamson, T. 33, 88, 103–12, 113 n.,
 114–16, 118, 122–3, 125–6, 132, 147
 n. 25, 164 n. 9, 172, 176–7, 282, 298
 n. 3
 Wilson, D. 2 n. 2
 worlds, branching 84–6
 Wright, C. 22 n. 35, 33, 158 n. 2, 165 n.,
 168 n., 170 n., 172, 178 n. 25, 180 n.,
 210 n. 13, 216 n. 25, 226–7, 243 n. 10,
 245 n. 16, 288 n., 289 n., 292 n. 5, 298,
 301–2
- Yablo, S. 246 n. 21
 Yalcin, S. 246 n. 21
- Ziff, P. 132 n. 3