

July 2008

A DEFENSE OF MODERATE INVARIANTISM

Leo W. Iacono

University of Nebraska at Lincoln, leo.iacono@gmail.com

Follow this and additional works at: <http://digitalcommons.unl.edu/philosophydiss>



Part of the [Philosophy Commons](#)

Iacono, Leo W, "A DEFENSE OF MODERATE INVARIANTISM" (2008). *Philosophy Dissertations, Theses, & Student Research*. 1.
<http://digitalcommons.unl.edu/philosophydiss/1>

This Article is brought to you for free and open access by the Philosophy, Department of at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in Philosophy Dissertations, Theses, & Student Research by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

A DEFENSE OF MODERATE INVARIANTISM

by

Leo Iacono

A DISSERTATION

Presented to the Faculty of

The Graduate College at the University of Nebraska

In Partial Fulfillment of Requirements

For the Degree of Doctor of Philosophy

Major: Philosophy

Under the Supervision of Professor Albert Casullo

Lincoln, Nebraska

August, 2008

DISSERTATION TITLE

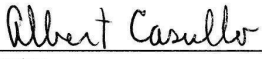
A Defense of
Moderate Invariantism

BY
Leo Iacono

SUPERVISORY COMMITTEE:

Approved

Date



Signature

July 1, 2008


Albert Casullo
Typed Name



Signature

July 1, 2008

John Gibbons
Typed Name



Signature

July 1, 2008


Joseph Mendola
Typed Name



Signature

July 1, 2008

Mark van Rooijen
Typed Name



Signature

July 1, 2008

Sarah Sawyer
Typed Name



Signature

July 1, 2008

Allan Donsig
Typed Name

A DEFENSE OF MODERATE INVARIANTISM

Leo Iacono, Ph.D.

University of Nebraska, 2008

Adviser: Albert Casullo

This dissertation is a defense of moderate invariantism, the traditional epistemological position combining the following three theses: invariantism, according to which the word ‘know’ expresses the same content in every context of use; intellectualism, according to which whether one knows a certain proposition does not depend on one’s practical interests; and anti-skepticism, according to which we really do know much of what we ordinarily take ourselves to know.

Moderate invariantism needs defending because of seemingly powerful arguments for contextualism, the view that, like ‘I’ and ‘now’, ‘know’ expresses different contents in different contexts. It has been argued that only contextualism can properly reply to skeptical arguments, and that only contextualism can account for our tendency to go from judging that a knowledge claim is true to judging that it is false in response to shifts in the context of use. Moderate invariantist replies to these arguments have largely failed. I propose new replies on behalf of moderate invariantism, while critiquing earlier, less successful attempts.

Chapter one is introductory. In chapters two and three, I examine contextualist replies to skeptical arguments arising from radical skeptical hypotheses (such as the possibility that one is a brain in a vat) and from considerations involving lotteries. I argue that if contextualism can adequately respond to these arguments, then there are equally effective replies available to the moderate invariantist. In the next two chapters, I examine Stewart Cohen’s “Airport Case,” which elicits intuitions about knowledge claims that supposedly only contextualism can accommodate. In chapter four, I argue that none of the invariantist replies to the case that depend on denying the intuitions succeed; in chapter five, I accept the intuitions, but argue that

contextualism does not follow from them. Finally, in chapter six, I evaluate two interesting invariantist critiques of contextualism; according to the first, contextualism collapses into the radical position that *every* English expression is context-sensitive; according to the second, 'know' fails a test for context-sensitivity involving indirect speech reports.

For Millie, who puts things into context

ACKNOWLEDGMENTS

Several current or former Nebraska philosophers have provided me with helpful comments and advice at various stages of writing this dissertation. My greatest debt, by far, is to my advisor Albert Casullo, who has been an exemplary teacher and mentor to me throughout my graduate education. Al provided prompt, careful, and incisive feedback on multiple drafts of each of the chapters of this dissertation, as well as related papers that I sent to journals and conferences, and several other proto-chapters that didn't get into the final version. In addition to his help with my dissertation, Al has been an unfailing source of support and good advice about the practice of philosophy and the philosophical profession. I couldn't have hoped for a better advisor. Thanks also to Tim Black and Sarah Sawyer for comments on a draft of chapter 2, and to Joseph Mendola for his comments on a late draft of the entire dissertation.

Much of the material contained here was initially presented at several meetings of the UNL philosophy department's graduate student colloquium. I thank all those who attended those talks, and in particular Brent Braga, David Chavez, Mark Decker, Cullen Gatten, Danielle Hampton, Howard Hewitt, Clayton Littlejohn, Tim Loughlin, Sruthi Rothenfluch, and Steve Swartzler.

Finally, love and thanks to my wife Amanda Gailey, whose patience, encouragement, good humor, and wit helped make writing this dissertation, and much else besides, a pleasure.

CONTENTS

ABSTRACT	iii
ACKNOWLEDGMENTS	vi
TABLE OF CONTENTS	vii
Chapter	
1. INTRODUCTION.....	1
1. Moderate invariantism and its competitors.....	1
2. Traditional epistemology and EC: business as usual?	3
3. Further preliminary comments.....	8
4. Outline of chapters 2–6.....	17
2. CONTEXTUALISM AND SKEPTICISM.....	20
1. Introductory remarks.....	20
2. Constraints on an adequate response to the skeptical paradox	21
3. Invariantist and contextualist responses to AI	23
4. Representing different knowledge relations	24
5. DeRose-style responses to AI.....	25
6. The moderate standards argument from ignorance.....	27
7. Critiques from semantic ignorance	31
8. Restricting the contextualist’s options	38
9. The paradox presented by MAI	49
3. CONTEXTUALISM AND THE LOTTERY	53
1. Introductory remarks.....	53
2. The lottery conviction	56

3. Lottery arguments	57
4. EC and lottery arguments.....	65
5. SSI and lotteries	69
6. A moderate invariantist reply to lottery arguments	79
7. Rejecting the lottery conviction.....	85
4. THE AIRPORT CASE (I): INTUITION–DENYING REPLIES	96
1. Introductory remarks.....	96
2. Warranted assertability maneuvers	100
3. Implicature based replies	114
4. Bach and belief removal	126
5. Williamson and psychological bias.....	128
6. Stanley and salient questions	131
5. THE AIRPORT CASE (II): SPEECH ACT PLURALISM	136
1. Introductory remarks.....	136
2. The response from speech act pluralism.....	137
3. A new argument for EC	143
4. An objection to the New Airport Argument	146
5. Conclusion	153
6. CAPPELEN AND LEPORE’S CRITIQUE OF CONTEXTUALISM.....	157
1. Introductory remarks.....	157
2. The collapse argument	157
3. A test for context-sensitivity.....	170
WORKS CITED.....	194

CHAPTER 1
INTRODUCTION

1. Moderate invariantism and its competitors

Epistemological inquiry has traditionally relied on two assumptions that have only recently come under heavy scrutiny. The first, *invariantism*, holds that when one uses the word ‘know’, the semantic content one is thereby expressing does not depend upon the conversational context in which one utters it. If invariantism is correct, then when I utter ‘*S* knows that *p*’, I attribute the very same epistemic property to *S* whether I am in the seminar room or the barroom, whether my life depends on the question of whether *p* is true or I am utterly indifferent to that question. The second assumption, *intellectualism*, holds that one’s practical situation is irrelevant to the question of whether one knows a given proposition. It might strike you as plausible that when the consequences to *S* of being wrong about *p* are especially bad, *S* needs more or better evidence for *p* in order to know that *p* than she would if *p* were of no practical importance to her. If intellectualism is correct, that impression is mistaken. A further traditional assumption, one that remains largely unchallenged, is *anti-skepticism*: the view that radical skepticism is either false or is at least a terrible cost that an epistemological theory ought to avoid if at all possible. Conjoin invariantism, intellectualism, and anti-skepticism, and you have *moderate invariantism*, the view to be defended in this dissertation.

Moderate invariantism was virtually ubiquitous in epistemology twenty years ago, and remains very popular. But it is far from clear that it is a defensible position. The challenges to moderate invariantism began with the advent of *epistemic contextualism* (EC). According to EC, ‘know’ is context-sensitive: the semantic content expressed by ‘know’ varies depending on the

context in which it is uttered. In ordinary contexts, the semantic content of ‘know’ is a relation that obtains between a subject S and a proposition p provided that S has satisfied some fairly modest epistemic requirements; in other contexts, the requirements become more difficult, or even impossible, to achieve. Since it holds that ‘know’ is context-sensitive, EC is inconsistent with moderate invariantism’s intellectualist component.

EC is supported by both linguistic and epistemological considerations. The primary linguistic consideration is that EC seems to best account for our shifting intuitions regarding the truth-values of knowledge attributions and denials (utterances of sentences of the form ‘ S knows that p ’ and ‘ S does not know that p ’) in response to shifts in the underlying contexts in which the attributions and denials take place. The primary epistemological consideration is that EC promises a satisfying response to skeptical arguments.

Some invariantists have responded that invariantism can accommodate much of the evidence for EC provided that it is granted that whether S knows that p depends in part on certain features that have not traditionally been thought epistemically relevant, most notably S ’s practical interests with respect to p . The version of invariantism that accords epistemological significance to such non-traditional factors is *subject sensitive invariantism* (SSI). Because it makes practical interests epistemically relevant, SSI is incompatible with the intellectualist component of moderate invariantism.

If invariantism’s only response to the arguments for EC requires falling back on SSI, then moderate invariantism is doomed. To defend their position, moderate invariantists must show that the arguments for EC are unsuccessful, without relying on any assumptions that would commit them to SSI. Several attempts in this direction have been made, but none are fully adequate; some are dismal failures, while others are on the right track but require adjustments or

further development. The primary goal of this dissertation is to defend moderate invariantism by providing fully adequate responses to the arguments for EC that do not require falling back on SSI. A secondary goal is to evaluate previous invariantist replies to EC, making use of and extending the most promising of these replies.

2. Traditional epistemology and EC: business as usual?

It is quite natural for epistemologists to suspect that the question of whether EC is correct is orthogonal to the epistemological problems that they are interested in, and to assume that they can continue to work on those problems without having to worry about whether EC is true. After all, EC is a linguistic thesis about the word ‘know’, but the problems that have traditionally occupied epistemologists concern non-linguistic concepts like knowledge and justification. Of course, since EC purports to provide a satisfying response to the problem of skepticism, it is at least potentially relevant to this part of epistemology. But when it comes to other fundamental epistemological issues about which EC has nothing to say—for example, the debates labeled ‘foundationalism versus coherentism’, ‘internalism versus externalism’, and ‘rationalism versus empiricism’—it might seem that the question of whether EC is true or false is simply irrelevant to how one should go about addressing these issues.

The view that, due to its focus on language, EC is irrelevant to traditional epistemological concerns is quite popular; we find it expressed in the following passages:

The main thesis of epistemic contextualism (EC) has considerable plausibility as a thesis in linguistics or in philosophy of language. In applying it to epistemology, however, it is possible to overreach. (Sosa 2000, 3)

The issues of greatest interest in epistemology thus seem independent of contextualism. (Sosa 2000, 6)

Some of the most important epistemological views in recent years—e.g., contextualism and subject-sensitive invariantism—have been grounded in large part in a close consideration of the semantics of knowledge ascriptions ... the proper focus in epistemology should begin to shift back to traditional considerations involving justification, evidence, warrant, and so on—which is where it should have been all along. (Lackey 2007, 619)

If the question of whether or not EC is true is irrelevant to most traditional epistemological problems, then responding to the arguments for EC might seem to be of limited epistemological interest. There are, however, good reasons for thinking that the question of whether or not EC is true *is* relevant to virtually any epistemological inquiry, and thus that any epistemologist should be *very* concerned about whether the arguments for EC can be resisted. This is not because EC really does have something to say about every important epistemological issue; rather, it is that if EC is correct, then epistemologists are prone to a kind of error that renders suspect the results of virtually any epistemological inquiry.

If EC is true, then it is a truth that nobody seems to have recognized until quite recently. Furthermore, EC remains highly controversial; as we shall see in subsequent chapters, many epistemologically and linguistically sophisticated philosophers continue to argue, vehemently, against EC. Now, if competent speakers could easily recognize when the content of ‘know’ shifts with corresponding contextual shifts, as they can with clearly context-sensitive words like ‘I’ or ‘now’, then EC would not be a controversial claim, just as the context-sensitivity of ‘I’ or ‘now’ is in no way controversial. So if EC is correct, then even sophisticated, competent speakers must be very bad at recognizing when ‘know’ expresses different contents in different contexts.

Here is how this fact about EC relates to epistemological inquiry. A central, and seemingly indispensable, feature of epistemological methodology is the reliance on counterexamples to refute epistemological theories. The way this goes, roughly, is that philosopher *A* offers a theory, (T), of some epistemic feature *E*,

(T) Necessarily, *E* obtains only if *C* obtains,
and philosopher *B* thinks up a seemingly possible situation *S* and claims:

(CX) In *S*, *E* obtains but *C* does not obtain.¹

If there is a robust and widely shared intuition that (CX) is true, then *S* becomes generally accepted as a counterexample to (T), and (T) is widely viewed as refuted. Epistemologists then try to provide alternatives or refinements to (T) in hopes of devising an account of *E* that is resistant to counterexamples. The paradigm case of this dialectical process is of course the long series of attempts to provide an analysis of knowledge in the wake of Gettier's (1963) counterexamples to the JTB account of knowledge, but the same general process can easily be discerned in the literature on every other important epistemological issue.

If EC is correct, then the process of finding and accepting counterexamples to epistemological theories is utterly unreliable and inadequate as it stands. In order for the truth of (CX) to refute (T), '*E*' must express the same content in (CX) and in (T). But (T) and (CX) are articulated in different contexts. In order to have reason to believe that *S* is a counterexample to (T), then, we must have reason to believe that '*E*' expresses the same content in the context in which (T) is articulated and the context in which (CX) is articulated. Suppose '*E*' is an expression containing 'know', for example, '*S* knows that *p*'. We have granted that if EC is true, then we are very bad at recognizing when 'know' expresses different contents in different contexts, and yet there are such situations. So if EC is correct, then for all we know 'know' expresses different contents in the context of (T) and in the context of (CX). If EC is correct, then, we are simply not entitled to rely on our intuitions that a given situation is (or is not) a counterexample to any theory involving knowledge in particular (as opposed to theories of justification, evidence, etc.) This would turn back the epistemological clock to 1963 at least,

since Gettier's counterexamples to the JTB account are themselves subject to this type of objection.

Epistemologists who focus on issues related to justification rather than knowledge might insist that EC remains irrelevant to their concerns, since the theories they articulate and the counterexamples they produce involve the word 'justified' rather than 'know'. There is, however, some reason to think that if the word 'know' is context-sensitive, then so is the word 'justified'. It is quite plausible that '*S* is justified in believing that *p*' is true only if *S* has the degree of justification for *p* required in order for 'knows that *p*' to apply to *S*; but on one very natural way of construing EC, the degree of justification for *p* that *S* must have in order for 'knows that *p*' to apply to *S* depends on the context in which 'knows that *p*' would be uttered. So it is quite plausible, given the truth of EC, that 'justified' is also context-sensitive. The preceding is not intended as a decisive argument that if EC is true, then 'justified' is context-sensitive; still, it should at least cause the epistemologist focused on justification to worry that 'justified' might be context-sensitive, and so to worry that the method of counterexamples is unavailable to her as well.

An epistemologist who recognizes that EC is true, and thus that she cannot blithely continue to rely on counterexamples in her own area of interest, might instead try to rehabilitate the method of counterexamples by investigating the details of EC; for example, if she could determine the set of pairs of contexts in which 'know' expresses different contents, and could then verify that the pair consisting of the context in which (T) is articulated and the context in which (CX) is articulated is not a member of this set, then she could legitimately present *S* as a counterexample to (T). My point is not that all epistemological inquiry is doomed if EC is true,

but rather that if EC is true, then further epistemological inquiry must be put on hold until a thorough investigation of EC is undertaken.

In light of the sorts of methodological disruptions that I have argued would be required if EC turns out to be true, epistemologists who are focused on traditional epistemological problems and are not otherwise interested in linguistic issues should still be very concerned with the question of whether EC is true, and consequently with the question of whether the arguments in favor of EC are successful. Furthermore, traditionally minded epistemologists should be reluctant to embrace the invariantist strategy of responding to the arguments for EC by falling back on SSI. Unlike EC, a linguistic thesis about the semantic properties of the word 'know', SSI is a substantive and radical account of the conditions of knowledge itself, with several surprising and counterintuitive implications (Stanley 2005, 124-125). For example, SSI appears to have the consequence that an expert in a field who has a lot to lose if the claims in her recent article turn out to be false does not know those claims, whereas a neophyte who is reading the expert's article but doesn't care about the field can know that the article's claims are true. This is a substantive, and bizarre, epistemological result that has clear implications for the theory of testimonial knowledge. So, whereas endorsing EC would result in serious methodological disruptions, endorsing SSI would result in serious theoretical disruptions. Only moderate invariantism requires no substantive methodological or theoretical disruptions. Traditionally minded epistemologists therefore ought to be very interested in the question of whether moderate invariantism is a defensible position.

Epistemologists focusing on traditional problems want to go about their business without worrying about whether EC is true. I have argued that they cannot do this unless moderate

invariantism is a defensible position. Thus, traditionally minded epistemologists ought to care about the project of defending moderate invariantism.

3. Further preliminary comments

The preceding discussion was meant to provide rough characterizations of some of the views to be discussed in this dissertation, and to motivate my project of defending moderate invariantism. There are three additional, largely unrelated, issues that it seems best to broach now rather than later. First, since semantic content and context-sensitivity are central notions in both moderate invariantism and EC, I want to articulate the conception of semantic content and context-sensitivity that I will be relying on throughout this work. Second, I want to justify my characterization of EC as the view that the *word* ‘know’ is context-sensitive, and compare it to a not obviously equivalent characterization of EC, according to which *sentences* containing ‘know’ are context-sensitive. Third, I want to state and explain my assumption that some form of an epistemic closure principle is correct.

Semantic content and context-sensitivity

Moderate invariantism and EC are claims about the semantic content of ‘know’. In order to fully understand these views, then, one must understand what semantic content is. Unfortunately, characterizing an expression’s semantic content and distinguishing it from other sorts of information or content that the expression might be used to convey—i.e., *pragmatic* content—is itself an important and controversial question in philosophy of language and linguistics. So it is not possible to provide an uncontroversial, informative definition of an expression’s semantic content. What I will try to do instead is sketch out how I conceive of an

expression's semantic content. I intend this picture to be very close to the mainstream or orthodox view in the philosophy of language, but I don't deny that it involves several assumptions that could reasonably be disputed.

Here's the picture. The semantic content of an expression is content that is determined by the conventional rules of the language (Soames 2002, 108-109). If the expression is complex, then it has constituents that have their own semantic contents and that bear certain syntactic relations to one another. There are conventional rules that determine the semantic content of a complex expression given the semantic contents of the expression's constituents and the syntactic relations among the constituents. The semantic content of a declarative sentence is a proposition, an abstract entity that can be either true or false. I will assume that propositions are structured: that they have individuals, properties, and relations as constituents, and that the semantic content of a declarative sentence is made up of the semantic contents of the sentence's constituents.²

So the semantic content of a complex expression is built out of the semantic contents of its constituents, whose semantic contents may themselves be determined by the semantic contents of *their* constituents, and so on. This cannot go on forever, though; eventually, we get to simple expressions that do not have further expressions as constituents. The semantic content of a simple expression relative to a context is given by a function that takes contexts as inputs and delivers semantic contents as outputs. This function from contexts to semantic contents is what Kaplan (1989) calls the expression's *character*. An expression's character is its conventional meaning: it is what a speaker must grasp in order to use the expression competently. Certain simple expressions may have stable characters: in such cases, the character function assigns the same semantic content to the expression relative to every context. Natural kind terms like 'water'

or 'gold' seem like good candidates for simple expressions with stable characters. Other simple expressions have unstable characters that assign different semantic contents to the expressions in different contexts. Expressions that have unstable characters are *context-sensitive*; paradigm examples of context-sensitive expressions are indexicals like 'I', 'now', 'here', 'this', and 'that'.

The preceding sketch should give you a rough sense of what a declarative sentence's semantic content is (a structured proposition composed of the semantic contents of the sentence's constituents), how the semantic content of a declarative sentence is related to the semantic contents of its constituents (the semantic contents of the constituents both compose and determine the semantic content of the sentence), and how the semantic content of a simple expression relative to a context is determined (it is given by the expression's character). What has not yet been discussed is the role semantic content plays in communication. Here I endorse a very plausible suggestion by Soames (2002, 56-57): that a declarative sentence's semantic content in a context is *asserted* whenever a competent speaker utters the sentence in that context and is speaking literally. Asserting a proposition p is stronger than merely communicating p : it involves *claiming* or *stating* that p is true, and a speaker who asserts that p thereby takes on a certain responsibility; she is in some sense vouching for the truth of p or providing assurances that p is true, and she is subject to criticism should her epistemic position with respect to p be unsuitable.³ In non-literal contexts, in which one is speaking ironically, metaphorically, loosely, or hyperbolically, one may succeed in uttering a declarative sentence without thereby asserting its semantic content. (No doubt one still *communicates* the sentence's semantic content in non-literal contexts, but one need not assert it.) When a competent speaker utters a declarative sentence and is speaking literally, however, she not only communicates, but also asserts the sentence's semantic content. So one role that a declarative sentence's semantic content plays in

communication is that it is the sort of content that gets asserted when a competent speaker who is speaking literally utters the sentence.

Another role that a declarative sentence's semantic content plays in communication is to allow hearers to grasp other, non-semantic, propositions conveyed by an utterance of the sentence. If I say, hyperbolically, 'I've been waiting in line forever', I am trying to convey, not that I have literally been waiting in line for all time, but rather that I have been waiting in line for a very long time. In order for a hearer to understand what I am trying to convey, however, she must first understand the sentence's semantic content, and then, recognizing that the semantic content could not possibly be true (and that I know this, and that I know that she knows this, etc.), draw an inference to the proposition that I am trying to convey. Similar points apply to other cases of non-literality, such as irony, loose use, and metaphor, as well as to propositions that are conversationally implicated by an utterance in the manner described by Grice (1989). In all of these cases, a hearer must grasp the semantic content of the uttered sentence in order to draw a reasonable conclusion about the non-semantic proposition that the speaker intends to convey.

Does EC really claim that 'know' is context-sensitive?

As I characterize it above, EC holds that 'know' (as well as its cognates, 'knows', 'knew', etc.) is a context-sensitive term. This way of characterizing EC can certainly be found articulated in print. For example:

The word 'I' is an indexical expression; it contributes potentially different semantic contents to the semantic content of sentences containing it relative to different contexts of use. Similarly, according to the contextualist, the word 'know' is an indexical expression. Relative to different contexts of use, it expresses different relations between persons and propositions. (Stanley 2004, 112)

Epistemic contextualism (EC)—the view that the content of the predicate ‘know’ can change with the context of utterance—has fallen into considerable disrepute lately. (Blome-Tillmann 2008, 29)

In the vast majority of cases, however, EC is characterized as a thesis about the context-sensitivity of *sentences* containing ‘know’. Consider, for example, the characterizations of EC one finds in general introductions to the topic:

EC holds that the truth conditions for tokens of sentences that include “knows” (and cognate expressions) vary with the attributor’s context. (Conee 2005, 51)

EC is the view that the proposition expressed by a given knowledge sentence (‘*S* knows that *p*’, ‘*S* doesn’t know that *p*’) depends upon the context in which it is uttered. (Rysiew 2008)

For the most part, the characterization of EC in terms of the context-sensitivity of ‘know’ and the characterization in terms of the context-sensitivity of knowledge attributing sentences are treated interchangeably and assumed to be equivalent. Some philosophers, however, have distinguished the view that ‘know’ is an indexical from other possible construals of EC. For example, Cohen writes:

How from the point of view of formal semantics should we think of this context-sensitivity of knowledge ascriptions? We could think of it as a kind of indexicality. On this way of construing the semantics, ascriptions of knowledge involve an indexical reference to standards. So the knowledge predicate will express different relations (corresponding to different standards) in different contexts.

But we could instead view the knowledge predicate as expressing the same relation in every context. On this model we view the context as determining a standard at which the proposition involving the knowledge relation gets evaluated. So we could think of knowledge as a three-place relation between a person, a proposition, and a standard. (Cohen 1999, 61)

According to the indexical construal of EC, ‘know’ expresses different relations in different contexts, where the relations are distinguished by the standards that a subject must meet in order to be in that relation with respect to a given proposition. According to the non-indexical

construal of EC, ‘know’ expresses the *same* relation in every context. However, this is a three-place relation relating a subject, proposition, and epistemic standard, and context provides the value for the epistemic standard.

Given these two possible construals of EC, a worry arises: is my initial characterization of EC (as the claim that ‘know’ is context-sensitive) inconsistent with the non-indexical construal of EC, according to which ‘know’ expresses the same relation in every context? If, as the non-indexical construal has it, ‘know’ expresses the same relation in every context, is it not correct to say that the semantic content of ‘know’ does *not* vary with the context? This worry seems to rest on the assumption that, according to the non-indexical construal, the semantic content of ‘know’ is limited to the three-place relation in question. But it seems more correct to say that on the non-indexical construal, the semantic content of ‘know’ consists of *both* a three-place relation *and* a value for one of the relation’s argument-places; namely, the argument-place that takes an epistemic standard as a value. (Consider that if one replaces ‘knows’ in ‘*S* knows that *p*’ with ‘wishes’, the truth-conditions of the resulting sentence ‘*S* wishes that *p*’ no longer involve epistemic standards at all; so ‘knows’ must be the word responsible for contributing an epistemic standard to the truth-conditions of ‘*S* knows that *p*’.) True, the epistemic standard is contextually determined, but this is no impediment to claiming that it is also (part of) the semantic content of ‘know’ on a particular occasion of use. (The individual to which ‘*I*’ refers on a given occasion of use is provided by the context, but this is no impediment to claiming that the semantic content of ‘*I*’ on that occasion of use is that individual.) If the semantic content of ‘know’ on a given occasion of use consists of both an (invariant) three-place relation and a (variable) epistemic standard, then the semantic content of ‘know’ clearly does vary from context to context, and ‘know’ is context-sensitive.⁴

It seems fair, then, to continue to assume that the formulation of EC in terms of the context-sensitivity of ‘know’ and the more common formulation in terms of the context-sensitivity of knowledge attributions are equivalent, and to move from one to the other as the situation dictates. That, at any rate, is what I shall be assuming in the chapters to come.

Epistemic closure

As mentioned in section 1, one of the most important considerations in support of EC is its alleged ability to provide an adequate response to various skeptical arguments. Consider, for example, the following skeptical argument:

- (P1) I don’t know that I am not a bodiless spirit being deceived by an evil genius;
- (P2) I know that I have a body only if I know that I am not a bodiless spirit being deceived by an evil genius;
- (C1) Therefore, I do not know that I have a body.

Arguments of this form will be discussed in chapters 2 and 3. There the moderate invariantist solution of denying (P1)—or the corresponding premise of some analogous skeptical argument—will be contrasted with the contextualist response, according to which the argument is sound whenever it is uttered but this does not threaten the truth of our ordinary knowledge attributions. There is, however, a third important way of responding to such arguments, namely denying (P2), or the corresponding premise in a skeptical argument of the same form. (P2) is known as the *closure* premise, because it is entailed by the following *epistemic closure principle*:

- (CL1) If *S* knows that *p* and *p* entails *q*, then *S* knows that *q*

Closure principles are so named because they claim that a set is *closed* under a certain relation: if *a* is a member of the set and *a* stands in the given relation to *b*, then *b* is also a

member of the set.⁵ (CL1) says that if a proposition p is in the set of propositions known by S , and p entails q , then q is itself in the set of propositions known by S . In other words, (CL1) states that the set of propositions known by S is closed under the entailment relation.

Denying (P2) might seem like a very attractive way of responding to the skeptical argument above, since (CL1) is very implausible. Suppose that p entails q . Somebody might both believe and know that p but not believe that q , because she is unaware that p entails q . If, as is usually thought, belief is a condition on knowledge, then in such a case the subject would know that p without knowing that q , even though p entails q .

Even though (CL1) is implausible, responding to the skeptical argument above by denying (P2), and justifying that denial by arguing that (CL1) is false, is not a good long term response to skeptical arguments. This is because even if (CL1) is false, there are good reasons for thinking that some other epistemic closure principle in the neighborhood is correct. But, no matter which epistemic closure principle one ends up defending, it seems that one can construct a skeptical argument that includes a premise entailed by the favored closure principle. For example, one plausible closure principle is (CL2):

(CL2) If S knows that p and S knows that p entails q , then S is in a position to know that q

A skeptical argument that depends on (CL2) instead of (CL1) is:

(P3) I am not in a position to know that I am not a bodiless spirit being deceived by an evil genius;

(P4) I know that having a body entails not being a bodiless spirit being deceived by an evil genius;

(P5) If I know that I have a body and I know that having a body entails not being a bodiless spirit being deceived by an evil genius, then I am in a position to know that I am not a bodiless spirit being deceived by an evil genius;

(C1) Therefore, I do not know that I have a body.

Other than the closure principle, the premises of this skeptical argument seem no weaker than the preceding argument based on (CL1): (P4) is clearly true, and (P3) seems no less plausible than (P1), since if I do not already know that I am not a bodiless spirit deceived by an evil genius, there seems to be no way in which I could come to know this; in particular, I could not come to know it via deduction from my belief that I have a body and my knowledge that having a body entails not being a bodiless spirit being deceived by a demon. Now, however, (P5) is an instance of the more plausible closure principle (CL2). Responding to this skeptical argument by denying the closure premise requires denying (CL2).

Since identifying weaknesses in less plausible epistemic closure principles like (CL1) does nothing to respond to skeptical arguments incorporating (instances of) far more plausible epistemic closure principles, like (CL2), there is little point in responding to a skeptical argument by denying the argument's closure premise unless one is willing to hold that even the most plausible formulations of epistemic closure principles are mistaken.

There are several problems with denying epistemic closure.⁶ For example, DeRose (1995, 200-201) has argued that denying a plausible epistemic closure principle commits us to the truth of "abominable conjunctions" like (AC):

(AC) While I don't know that I'm not a bodiless spirit being deceived by an evil genius, still, I know that I have a body.

If epistemic closure is false, and if this is supposed to help address skeptical arguments like the foregoing, then (AC) has to be true. But (AC) is virtually unintelligible, as even closure deniers like Dretske (2005, 17) concede.

There are several other compelling objections to denying closure, but discussing them here would take us to far afield.⁷ The cumulative force of these objections is to make the attempt to respond to skeptical arguments by denying closure highly unappealing. Furthermore, the three most important defenders of EC, Cohen, DeRose, and Lewis, all endorse epistemic closure—indeed, each of them argues for EC on the grounds that it can provide an adequate reply to skeptical arguments without denying closure. In what follows I will assume that some version of an epistemic closure principle is correct, and I will restrict myself to discussing versions of EC, like Cohen's, DeRose's, and Lewis's, that share this assumption.⁸

4. Outline of chapters 2–6

My goals in the following chapters are twofold. First, and most importantly, I will try to show that none of the important arguments in support of EC require abandoning moderate invariantism; for each of these arguments, there is an acceptable invariantist response that does not require falling back on a version of invariantism, like SSI, that is inconsistent with moderate invariantism.⁹ Second, I will evaluate several previous invariantist responses to EC.

Chapter 2 is the first of two dealing with EC's response to skeptical arguments. In this chapter I consider Keith DeRose's use of EC to respond to the Argument from Ignorance: I don't know that I am not a brain in a vat; if I don't know that I am not a brain in a vat, then I don't know that I have hands; therefore, I don't know that I have hands. I argue that if DeRose's

response to the Argument from Ignorance succeeds, then there is an equally good reply available to the invariantist.

Chapter 3 concerns skeptical arguments that arise from the *lottery conviction*: the widely held view that one cannot know that a certain lottery ticket will lose on the basis of the odds alone. EC and SSI both provide plausible responses to these *lottery arguments*; I offer a moderate invariantist reply. My reply involves denying the lottery conviction—I hold that one can know that a lottery ticket will lose on the basis of the odds alone, provided that it is sufficiently probable that it will lose. To defend this view, I offer a positive argument in its support, then respond to objections.

Chapter 4 is the first of two on so-called *context-shifting arguments*: arguments for EC that trade on cases in which a knowledge attribution and its denial both appear to be true, even though the epistemic position of the subject of the attribution and denial remains fixed. I focus on Cohen's (1999) Airport Case. I survey a range of invariantist replies to the Airport Case that rest on claiming that one of the intuitions elicited by the case is false, and argue that none of these intuition-denying replies succeed.

In Chapter 5, I defend an invariantist response to the Airport Case that accepts that the intuitions elicited by the case are correct, but denies that EC follows from this. My response relies on the thesis of *speech act pluralism*, according to which a single utterance of a declarative sentence is typically used to simultaneously assert several propositions.

Chapter 6 is part of my secondary project of assessing invariantist replies to EC. I consider two interesting anti-contextualist arguments presented by Herman Cappelen and Ernie LePore (C&L). According to the first, the same context-shifting arguments that are used to establish views like EC can be used to establish *radical contextualism*, the view that *every*

English expression is semantically context-sensitive. C&L then argue that radical contextualism is false; so, they claim, context-shifting arguments must be unsound. I argue that this objection fails, since C&L do not succeed in showing that context-shifting arguments inevitably lead to radical contextualism. C&L's second argument is that 'know' does not pass a test for context-sensitivity involving indirect speech reports. I argue that the proposed test for context-sensitivity is unsound.

¹ Instead, (T) might be of the form 'If *C* obtains, *E* obtains', and (CX) might have the form 'In *S*, *C* obtains but *E* does not obtain'. Or (T) might be a biconditional articulating necessary and sufficient conditions for *E*.

² An important alternative position on the nature of propositions holds that propositions are sets of possible worlds. This view is defended by Stalnaker (1984). For more on this debate, see King (2008).

³ The nature of assertion and the epistemic norms that govern it will be discussed further in chapter 4.

⁴ There is a related worry. In subsequent chapters, I will refer to the different relations expressed by 'know', and assume that according to EC, 'know' expresses different relations in different contexts. Is this not inconsistent with the non-indexical construal of EC? I do not think there is any inconsistency here, since whenever a subject *S*, proposition *p*, and epistemic standard *e* are in some three-place relation *R*, there is a corresponding two-place relation *R_e* such that *S* is related by *R_e* to *p*; (*S*, *p*) is in the extension of *R_e* if and only if (*S*, *p*, *e*) is in the extension of *R*. Given the non-indexical construal of EC, we can always state the truth-conditions of a knowledge attribution in a given context either in terms of a three-place relation where a contextually provided epistemic standard is one of the relata, or in terms of a two-place relation where the same contextually provided epistemic standard has been built into the two-place relation. These differences strike me as purely terminological. So, when I later refer to the relation expressed by 'know' in a given context, you can, if you wish, understand me to be referring to the two-place relation where the contextually provided epistemic standard has been built into the relation.

⁵ For a similar characterization of epistemic closure principles, see Kvanvig (2006).

⁶ Important attempts to respond to skepticism by denying closure include Dretske (1970) and Nozick (1981, 167-247).

⁷ Hawthorne (2005) presents an extremely powerful case against denying closure.

⁸ There are versions of EC that deny epistemic closure; for example, Heller (1999). I will ignore these in what follows, and treat EC as incorporating the assumption that epistemic closure is correct.

⁹ There are two other considerations for EC that I will not discuss here. The first, defended by Lewis (1996), is that EC can supposedly solve the Gettier problem; Cohen (1998) criticizes Lewis's solution. The second consideration is that EC can supposedly solve the problem of easy knowledge; see Cohen (2002).

CHAPTER 2

CONTEXTUALISM AND SKEPTICISM

1. Introductory remarks

Among EC's many alleged advantages, one of the greatest is supposed to be that it provides an adequate reply to the problem of skepticism. If, as is often thought, invariantism cannot provide an adequate reply to skepticism, then EC's supposedly satisfactory response provides us with one reason to be contextualists rather than invariantists.

It is my view that, whatever its other merits, EC is of no use in responding to skepticism. In this chapter, however, I will defend a somewhat narrower thesis. I will argue that DeRose's (1995) contextualist strategy for responding to the so-called *Argument from Ignorance* (AI) cannot succeed. Part of DeRose's solution involves attributing a certain kind of error to those who experience the paradoxical intuitions that AI elicits; it is claimed that they judge AI's conclusion to be false because they recognize that it would express a falsehood if uttered in an ordinary context. On the face of it, however, such a mistake does not seem easy to make. As it turns out, DeRose's options in explaining how such a mistake is possible are extremely limited. Most explanations do not square with some plausible condition on an adequate reply to skepticism. It will be shown that the only path that avoids violating one of these conditions must claim that a speaker who utters AI mistakenly intends to be expressing *another* skeptical argument. This other skeptical argument presents the contextualist with exactly the same philosophical problem that AI presents the invariantist; hence, a DeRose-style response to AI does not solve the skeptical problem.

DeRose's contextualist response to skepticism is the most fully developed and influential contextualist response yet articulated, and other contextualists have either tacitly or explicitly endorsed DeRose's broad approach.¹ Thus, showing that DeRose's response is unsuccessful goes a long way towards showing that EC is of no use in responding to skeptical arguments. Having said that, I do not deny that there may be other ways of employing EC to respond to skeptical arguments. Since these alternative approaches have (so far as I know) yet to be articulated, however, it seems fair to focus on DeRose's response in what follows.

2. Constraints on an adequate response to the skeptical paradox

The skeptical argument we shall be focusing on is:

Argument from Ignorance (AI)

- (P1) I don't know that I am not a BIV (a brain in a vat).
- (P2) If I know that I have hands then I know that I am not a BIV.
- (C1) So, I don't know that I have hands. (DeRose 1995, 183)

AI appears to be valid; its premises appear to be true; and yet, its conclusion appears to be false. In other words, AI presents us with a paradox: as Sainsbury (1995, 1) puts it, "an apparently unacceptable conclusion derived by apparently acceptable reasoning from apparently acceptable premises." The philosophical challenge posed by AI is to adequately resolve the paradox.

But what counts as an adequate solution to the paradox presented by AI? Clearly no solution can do justice to all of our intuitions, on pain of contradiction. At least one strongly held intuition will have to be denied. Furthermore, it is widely agreed that it is not sufficient to simply

identify the false intuition, or even to explain *why* it is false—one must also provide a plausible explanation for why we had the misleading intuition in the first place. DeRose puts it like this:

But we should hope for a better treatment of the argument than simply choosing which of the three individually plausible propositions—the two premises and the negation of the conclusion—seems least certain and rejecting it on the grounds that the other two are true. In seeking a solution to this puzzle, we should seek an explanation of how we fell into this skeptical trap in the first place, and not settle for making a simple choice among three distasteful ways out of the trap. We must explain how two premises that together yield a conclusion we find so incredible can themselves seem so plausible to us. (DeRose 1995, 184)

So a fully adequate reply to the paradox posed by AI must identify the false intuition, explain why it is false, and explain why we had the intuition in the first place.

In chapter 1, I articulated two assumptions that the contextualists with whom I engage—DeRose, Cohen, and Lewis—also share. First, the paradox cannot be adequately resolved by denying (P2)—the closure denying move. Second, the solution must not capitulate to a radical sort of skepticism; it must be consistent with the truth of a large proportion of our ordinary claims to know—claims to know everyday propositions made in non-philosophical contexts. In this chapter I will continue to assume that these are legitimate constraints on an adequate reply to skepticism.

Finally, I make another important assumption that will place further constraints on an adequate reply to AI. I will assume that the paradoxical intuitions elicited by AI are universal—that everybody who considers AI, to a greater or lesser degree, feels intuitively moved to accept AI's premises, to deny its conclusion, and to judge the argument to be valid. This is not to say, of course, that anybody *believes* every member of this inconsistent set of claims. But (I assume) anybody who resolves the inconsistency by denying one of the claims thereby denies an intuition that persists even after she has settled on her answer to the paradox. Just as I believe that the two lines in the Müller-Lyer illusion are of equal length even though, in a sense, it appears to me that

they are of different lengths, one who resolves the paradox by rejecting, say, (P1), believes that (P1) is false even though in a sense (P1) continues to appear true to her.²

This assumption has two important consequences. First, an adequate reply to AI cannot rest on a subject having some deficiency—say, not paying attention—unless one is willing to claim that everybody who considers AI has the same deficiency.³ Presumably, a universally shared deficiency has some systematic, underlying cause; an adequate reply to AI which postulates a universal deficiency should explain why everybody who considers AI suffers from this deficiency. The second consequence is that an adequate reply to skepticism must not stop at explaining what is going on in the case where one person, a skeptic, uses an argument like AI to convince another person that skepticism is true. Such two-person cases are often invoked by contextualists to illustrate and defend their solutions to skepticism.⁴ The fact is, however, that AI elicits the same paradoxical intuitions if a non-skeptic utters AI in the privacy of her office. An adequate response to skepticism must explain what is going on when the paradoxical intuitions are elicited in such a ‘single-person case’.⁵

In summary, I am assuming that an adequate reply to AI must: deny one of the contradictory intuitions that constitute the skeptical paradox; explain why the intuition is false; explain why we have the intuition; not deny (P2); be consistent with the truth of a large proportion of our ordinary claims to know; and be consistent with the assumption that everybody who considers AI experiences the paradoxical intuitions.

3. Invariantist and contextualist responses to AI

Given all these constraints, the invariantist has little choice in how to answer AI. If AI establishes the truth of (C1), then similar arguments can be deployed to show that almost all of

our ordinary claims to know are false. Given that an adequate response must avoid this radically skeptical result, the invariantist must either deny that AI is valid, or deny one of the premises. Since, according to invariantists, ‘know’ expresses the same relation throughout the argument, AI is clearly valid. The invariantist must therefore deny (P1) or (P2). But we are assuming (P2) is true; so the invariantist must deny (P1), and explain why it is false. Further, since we are intuitively moved to judge that (P1) is true, the invariantist must explain why we have this misleading intuition.

Contextualists have two further options. First, they might concede that AI is sound, but deny that the truth of (C1) threatens the truth of the majority of our ordinary knowledge attributions. This is DeRose’s approach, and the focus of this chapter. The second contextualist option, which has not been widely discussed or clearly distinguished from the first, is to deny that AI is valid. Roughly put, the idea is that AI equivocates because ‘know’ does not express the same knowledge relation throughout the argument. This approach respects our strongly held intuitions with respect to the truth-values of (P1), (P2), and (C1): it holds that (P1) and (P2) are true, and that (C1) is false. However, those who hold that AI equivocates must explain away the strongly held intuition that AI is valid. The prospects for this sort of contextualist response to skepticism will not be considered here.⁶

4. Representing different knowledge relations.

EC holds that ‘know’ may express different knowledge relations in different contexts of utterance, and that these knowledge relations differ with respect to the stringency of the epistemic standards one must meet in order to be in the knowledge relation with respect to a proposition. In what follows I will be useful to have some terminology that clearly refers to these

different knowledge relations. I will use ‘*h*-know’ to refer to the knowledge relation that according to EC is associated with very demanding requirements and is expressed by ‘know’ in special, skeptical contexts (‘*h*’ stands for ‘high-standards’). I will use ‘*m*-know’ to refer to the knowledge relation that according to EC is associated with less demanding requirements and is expressed by ‘know’ in ordinary contexts (‘*m*’ stands for ‘moderate-standards’). I will stipulate that these new words are not context-sensitive, and therefore express the same knowledge relation in every context. I will also use the words ‘*m*-knowledge’ and ‘*h*-knowledge’ to refer, respectively, to what one has when one *m*-knows or *h*-knows a proposition.⁷

5. DeRose-style responses to AI

According to DeRose, uttering AI’s (P1) tends to put in place impossibly high epistemic standards that must be met in order for an ascription of knowledge of a proposition about the external world to be true; (P1) is therefore true. Furthermore, these high-standards tend to stay in place, so that they obtain as one utters (P2) and (C1). For this reason, one’s utterance of AI’s conclusion is true (since high-standards for knowledge remain in place, almost all knowledge denials—sentences of the form ‘*S* does not know that *p*’—will express truths when uttered in this context).

Using our new terminology, we can put DeRose’s view as follows: when one utters AI, (P1), (P2), and (C1) express the same propositions as the following argument:

High-Standards AI (HAI)

(HP1) I don’t *h*-know that I am not a BIV.

(HP2) If I *h*-know that I have hands, then I *h*-know that I am not a BIV.

(HC1) So, I don’t *h*-know that I have hands.

Since *h*-knowledge involves impossibly demanding epistemic requirements, (HP1) is clearly true. Given that denying closure is off the table, (HP2) is true. And (HC1) clearly follows from (HP1) and (HP2). So HAI is a sound argument. Let two arguments be *equivalent* if their premises and conclusion express the same propositions. Since HAI is equivalent to AI, AI is sound.

Even though AI is sound, this does not have far-reaching skeptical implications. For example, it does not follow that our *ordinary* claims to know are systematically false. In an ordinary context, (P1) does not express the true proposition expressed by (HP1); instead, it expresses the proposition expressed by (MP1):

(MP1) I don't *m*-know that I have hands.

And, since *m*-knowledge is quite easy to come by, (MP1) expresses a falsehood.

So DeRose's solution seems to perform a minor miracle—giving the skeptic his due, by conceding that AI is sound, while protecting the truth of most of our ordinary claims to know. The worry, however, is that DeRose may be giving the skeptic *more* than he is due. Intuitively, AI is not sound, since its conclusion appears to be false. DeRose must explain why (C1) seems so implausible when uttered as the conclusion of AI, given his view that it expresses a truth in that context.

DeRose's answer is that our intuition that (C1) is false when uttered as the conclusion of AI is due to our recognition that in a more ordinary context, (C1) would express a falsehood:

The very strong pull that not-(P3) continues to exert on (at least most of) us even when the standards are high is explained [as follows]: Even while we're in a context governed by high standards at which we don't count as knowing an ordinary proposition, we sense that as soon as we find ourselves in more ordinary conversational contexts, it will not only be true for us to claim to know these very ordinary propositions that the skeptic now denies we know, but it will be wrong for us to deny that we know these things. It's easy, then, to think that the skeptic's present denial must be equally false and that it would be equally true for us now,

in the skeptic's presence, to claim to know these ordinary propositions. (DeRose 1995, 209)⁸

By a 'DeRose-style response to AI', I mean a response that takes the following three stances: first, it claims that AI is a sound argument whenever it is uttered; second, it denies that the truth of AI's conclusion threatens the truth of our ordinary claims to know, because such ordinary claims are made in a different context in which 'know' expresses *m*-knowledge; finally, it explains one's intuition that AI's conclusion is false by claiming that one realizes that it would be false if uttered in an ordinary context, and concludes that it is false in one's present context.

A DeRose-style response to AI need *not* be committed to DeRose's particular account of the truth-conditions of knowledge attributions or the mechanisms responsible for triggering and preserving context shifts. So the thesis of this chapter, that a DeRose-style response to AI cannot succeed, applies to any version of contextualism that adopts DeRose's broad approach to AI, regardless of the theory's details.

6. The moderate standards argument from ignorance

We have seen that a DeRose-style response to AI claims that AI is equivalent to HAI, and that the truth of HAI's conclusion poses no threat to the truth of our ordinary claims to know, which express *m*-knowledge, rather than *h*-knowledge. As Sosa (2004, 46-48) and Pritchard (2005, 6) have noted, however, AI's first premise can be reformulated so that it explicitly refers to *m*-knowledge rather than *h*-knowledge. This reformulated premise can then be used in a skeptical argument which is untouched by a DeRose-style response to AI and which threatens the truth of our ordinary claims to know. The argument is:

Moderate-Standards Argument from Ignorance (MAI)

(MP1) I don't *m*-know that I am not a BIV.

(MP2) If I *m*-know that I have hands, then I *m*-know that I am not a BIV.

(MC1) So, I don't *m*-know that I have hands.

Since MAI's conclusion does threaten the truth of our ordinary claims to know, it poses at least a *prima facie* challenge to the contextualist. Note that EC, a claim about the context-sensitivity of 'know', will be of no help in responding to MAI, since MAI does not contain the word 'know'. Given the contextualists' commitment to closure and to anti-skepticism, they can't accept (MC1) or deny (MP2). So, since MAI is obviously valid, they must deny (MP1). It might seem, then, that the contextualists are in the same boat with respect to MAI that invariantists are with respect to AI—each must plausibly claim that we *do* know, *by ordinary standards*, that we are not BIVs, and explain how this is possible.

There is, however, a crucial difference between the contextualist's position with respect to MAI and the invariantist's predicament with respect to AI. It is accepted by all parties that (P1) is an intuitively compelling premise, and hence that invariantists must explain why we have what is according to them the misleading intuition that (P1) is true. Contextualists, however, claim that (P1) denies *h*-knowledge; for them, intuitions concerning (P1) are irrelevant to (MP1), which denies *m*-knowledge. Contextualists are in the same boat as invariantists only if (MP1) is itself a highly intuitive premise. Now, Sosa claims that a premise like (MP1) *is* highly intuitive:

Moreover, it seems to me that the skeptic's main premise exerts about as powerful an intuitive pull even if we think of it *explicitly* as the claim that *by ordinary standards* we do not know ourselves to be free of radical illusion. I do not myself agree with that premise, but I do find it very difficult to say why it is not right, given how plausible it is that we would have the very same experiential basis if we were at this moment undergoing such illusion. (Sosa 2004, 60, n. 19)

But it seems that the contextualist could reply that Sosa goes wrong when he attempts to consult his intuitions about whether he knows, by ordinary standards, that he is free from radical illusion, and that he accidentally registers his intuition that he doesn't know *by high standards* that he is

free from such illusion. The fact that ‘knows by ordinary standards’ and ‘*m*-know’ are pieces of newly introduced, philosophical terminology should make us worry that our intuitions concerning the truth values of sentences containing these terms will be influenced by our judgments concerning the truth values of similar sentences that contain more familiar terminology, such as ‘know’ all by itself.

For all that has been said so far, it seems fair for the contextualist to simply deny that he is under any obligation to explain the alleged intuitive plausibility of (MP1). If that is right, then in order to respond to MAI the contextualist must deny (MP1) and explain why it is false, but need not account for any intuitions concerning (MP1). This means that the contextualist will have a fairly easy time providing truth-conditions for (MP1); if it is objected that the truth conditions seem intuitively wrong, the contextualist can again reply that these intuitions are really about *h*-knowledge, not *m*-knowledge.

Let me try to make this point more concrete. The intuition that (P1) is true seems largely driven by the thought that knowledge requires evidence, and that one does not appear to have evidence that one is not a BIV. Exactly why it seems that one does not have evidence that one is not a BIV is an interesting and difficult question, but surely it has something to do with the fact that a BIV is stipulated to be having just the experiences that one is actually having.

Suppose an invariantist denies (P1) by claiming that we can, after all, know that we are not BIVs without evidence. This response will clash with our intuition that knowledge requires evidence. The invariantist will be forced to explain why our intuition that knowledge requires evidence is false, and why we have the misleading intuition.

Now, Cohen, DeRose, and Lewis each offer accounts of *m*-knowledge according to which one *m*-knows that one is not a BIV without evidence, thus making (MP1) come out false.

For Cohen (1988, 112), this *m*-knowledge stems from the ‘intrinsic rationality’ of believing that one is not a BIV, where an intrinsically rational belief is explicitly defined as one that it is reasonable to have without evidence. For Lewis (1996, 233), the *m*-knowledge derives from the fact that the alternative that one is a BIV is not relevant, and hence need not be ruled out in order to *m*-know that one is not a BIV. But in the case of ‘fanciful’ skeptical alternatives, evidence against the alternative is not required for it to be non-relevant, according to Lewis. So according to Lewis one can *m*-know that one is not a BIV without evidence. Finally, for DeRose (1995, 208), the *m*-knowledge stems from the fact that the set of possible worlds that is relevant to assessing the truth of knowledge ascriptions in ordinary contexts does not include a world in which one is a BIV. Again, evidence is not required in order to *m*-know that one is not a BIV.

Since these contextualists’ accounts of *m*-knowledge entail that one can *m*-know that one is not a BIV without evidence, it might be objected that this conflicts with the powerful intuition that knowledge requires evidence. But now the contextualist has a reply that is unavailable to the invariantist: he can claim that the intuition that knowledge requires evidence is an intuition about *h*-knowledge, not *m*-knowledge. By focusing on the skeptical hypothesis that one is a BIV, one has placed oneself in a high-standards context, and when one consults one’s intuitions about what knowledge requires in a high-standards context, one gets intuitions about *h*-knowledge. One might, like Sosa, object that one is consciously trying to think about *m*-knowledge rather than *h*-knowledge. The contextualist, however, can simply deny that one can reliably access one’s intuitions about *m*-knowledge when one is in a high-standards context, even when one consciously tries to do so.

It appears, then, that the contextualists’ position with respect to MAI is dialectically stronger than the invariantists’ with respect to AI. In section 9, I will argue that adopting a

DeRose-style reply to AI deprives contextualists of this advantage. Contextualists who adopt a DeRose-style response to AI really are in the same dialectical boat as invariantists, because a DeRose-style response to AI commits them to claiming that we judge AI's (P1) to be true *because* we mistakenly take it to express the same proposition as (MP1), and we judge that proposition to be true. So (I will argue) contextualists who adopt a DeRose-style response must concede that (MP1) is as intuitively compelling as (P1).

7. Critiques from semantic ignorance

DeRose's explanation for the implausibility of (C1) when uttered as AI's conclusion has been criticized as requiring an unacceptable kind of semantic ignorance. This criticism, however, has been marred by errors. In this section, I will discuss two critiques from semantic ignorance, by Schiffer (1996) and Rysiew (2001), and explain what I view as the mistakes involved in those critiques.

Schiffer's critique

Schiffer's critique of contextualist solutions to skepticism is chiefly directed at DeRose's explanation for the implausibility of AI's conclusion, according to which we incorrectly judge (C1) to be false when uttered as AI's conclusion because we recognize that it would be false if uttered in an ordinary context. Schiffer finds a tension between this "error theory" and EC's semantic account: the error theory holds that we do not know what we are saying when we utter (C1), but "speakers would know what they were saying if knowledge sentences were indexical in the way the Contextualist requires"; thus, "the semantics is refuted by the error theory" (Schiffer 1996, 325-328).

Schiffer likens the contextually provided epistemic standards that according to EC are part of the semantic content of a knowledge attribution to the contextually provided location that is part of the semantic content of an utterance of ‘It’s raining’. Just as ‘It’s raining’ uttered in London expresses (roughly) the proposition that it is raining in London, ‘*S* knows that *p*’ expresses the proposition that *S* knows that *p* relative to such-and-such epistemic standard.⁹ Schiffer finds it implausible that one could utter ‘I know that I have hands’ in a skeptical context, thus expressing the false proposition that one knows by high-standards that one has hands, yet believe that one is instead expressing the true proposition that one knows by moderate-standards that one has hands. He claims that this would be as if “a fluent, sane, and alert speaker, who knows where she is, were actually to assert the proposition that it’s raining in London when she mistakenly thinks she’s asserting the proposition that it’s raining in Oxford” (Schiffer 1996, 326). It seems impossible that a speaker with the specified qualities could make such a mistake, and Schiffer’s point is presumably that the mistake DeRose imputes to us when we utter AI is likewise impossible.

While highly suggestive, Schiffer’s argument against DeRose is unfortunately not entirely clear. In attempting to interpret the argument, Rysiew (2001, 483-484) tentatively suggests that Schiffer is simply denying that one can fail to know the proposition literally expressed by one’s utterance. If this is Schiffer’s argument then it rests on a premise that is clearly false: one *can* fail to know the proposition literally expressed by one’s utterance, if for no other reason than that one can be mistaken about what the words in one’s utterance literally mean. Rysiew’s interpretive suggestion, however, assigns no significance to Schiffer’s comparison between our situation when we utter AI and the situation of the speaker in the London case, who is stipulated to be alert, sane, fluent, and to know where she is. It is more

charitable to take this comparison into account and interpret Schiffer's argument as an argument by analogy, along the following lines:

(P4) Suppose subject *S* utters AI and judges (C1) to be false. If *S* makes the mistake that DeRose imputes to her in order to explain this judgment, that is like a sane, fluent, and alert speaker, who knows that she is in London when she utters the sentence 'It is raining', falsely believing that her utterance expresses the proposition that it is raining in Oxford.

(P5) But the London case is impossible: a sane, fluent, and alert speaker who knows that she is in London could not make such a mistake.

(C2) Therefore, the mistake that DeRose imputes to *S* is impossible.

Let's concede (P5) and grant that given the stipulated qualities of the speaker, it is impossible for her to make such a mistake. The real trouble with Schiffer's argument (interpreted as an argument by analogy) is (P4), which he leaves entirely undefended. More specifically, Schiffer fails to argue for the claim that when we are intuitively moved to judge that AI's conclusion is false, we have the same (or relevantly similar) qualities as the speaker in the London case. But this is a crucial issue, since it seems that it is the possession of these qualities that makes the speaker's mistake in the London case impossible. Suppose, for example, that we deprive the speaker of knowledge that she is in London. If she falsely believes she is in Oxford, then she will falsely believe that 'It is raining' expresses the proposition that it is raining in Oxford. So for the analogy to hold, when we consider AI we must have a quality that is relevantly similar to the speaker in the London case's quality of knowing that she is in London. Presumably, that relevantly similar quality is knowing that one is in a 'high-standards context', in which 'know' expresses *h*-knowledge rather than *m*-knowledge. But EC *per se* is not

committed to the claim that one must know that one is in a high-standards context when one utters AI. The question of whether one can fail to know that one is in a high-standards context turns in part on which features of a context determine that high-standards are in place; EC, however, is a general thesis that is compatible with a whole spectrum of more specific contextualist theories that differ with respect to how they answer this question. So EC itself does not provide any support for (P4). Furthermore, in presenting his objection, Schiffer does not discuss the contextual features that DeRose, his main target, says make for a high-standards context. So Schiffer's argument is incomplete, because it fails to show that the speaker in the London case and a person who experiences the paradoxical intuitions when considering AI have relevantly similar qualities.

Rysiew's critique

Although Rysiew thinks it is not only possible, but common, for speakers to fail to know the propositions literally expressed by their utterances, he does not think it is possible for a speaker to be ignorant of her own *semantic intentions*—she can't be mistaken about the proposition she *intends* to express by using the sentence. Rysiew uses this claim to argue against contextualist replies to skepticism:

The contextualist wants to tie what the sentences we utter mean *very* closely to what we mean in uttering them: remember, the contextualist holds that changes in the content of knowledge-attributing sentences track changes in 'context'; and both Cohen and DeRose conceive of context in terms of the purposes, intentions, interests, beliefs, and so forth, of the *speaker*. But this just means that the independently plausible idea that we can be wrong about what *the sentences we utter* mean becomes, on the contextualist view, the manifestly *implausible* claim that we're actually systematically mistaken about what *we* (knowledge attributors) mean in uttering knowledge-attributing sentences. (Rysiew 2001, 485)

Rysiew's argument appears to be this: Suppose a speaker utters AI in a single-person case and experiences the paradoxical intuitions, including the intuition that her utterance of (C1) is false. According to contextualists, the content of the speaker's utterance of (C1) is very closely tied to her semantic intentions in uttering (C1); contextualist replies to skepticism, however, are committed to her ignorance of the proposition expressed by her utterance of (C1); hence, according to contextualists the speaker is ignorant of the proposition she *intended* to express by uttering (C1); but this is "manifestly implausible." Rysiew concludes as follows:

It seems to me that the contextualist thus tries to save our pedestrian knowledge of such matters as whether we have hands against the sceptic's attacks by giving up on the idea that we know what our communicative intentions are. But that is simply to trade one form of skepticism for another: epistemically speaking, we gain the world but lose our minds. (Rysiew 2001, 285)

There are two problems with Rysiew's argument, as I've construed it. First, it is not obvious that a DeRose-style response to skepticism really is committed to the speaker's ignorance of the proposition expressed when she utters AI's conclusion. According to DeRose, the speaker finds AI's conclusion implausible because she recognizes that (C1) would be false if uttered in an ordinary context. This is obviously a mistake, but it doesn't clearly follow that the speaker doesn't know the proposition literally expressed by AI's conclusion. After all, if she were asked what she said when she uttered AI's conclusion, she could sincerely report, 'I said that I don't know that I have hands'. If this report is itself made in a high-standards context, then it expresses the proposition that she said that she doesn't *h*-know that she has hands. It is plausible that a person sincerely reports that *p* only if she believes *p*. So, there is a case to be made that the speaker truly believes that she said that she doesn't *h*-know that she has hands—and that she therefore has a true belief concerning the proposition literally expressed by her

utterance of AI's conclusion. If she has this true belief, it is not clear why it should *not* amount to knowledge.

Second, there is a gap in Rysiew's argument. In order for the argument to be valid an additional premise is required:

If the speaker's semantic intentions in uttering a sentence are very closely tied to the utterance's content and the speaker does not know the utterance's content, then the speaker does not know her own semantic intentions.

This premise appears to be false unless 'very closely tied' just means *identical*. If one supposes that the content of a knowledge ascription just *is* the proposition that the speaker intends to express, then it is plausible that if the speaker knows what she means to say, then she knows the utterance's content. On the other hand, suppose that the speaker's semantic intentions are one factor among many that are involved in determining, in some fairly complex way, the content of the utterance; then it will be possible for the content of the utterance to differ from the speaker's semantic intentions; furthermore, if the speaker is unaware of some aspect of how the utterance content is determined, she may fail to know the utterance content and yet know her own semantic intentions.

In order for Rysiew's argument to work, then, we must assume that the proposition the speaker intends to express with (C1) is simply identical to the semantic content of (C1) in the context of AI. No version of contextualism that I know of, however, simply identifies the semantic content of a knowledge ascription with the proposition the speaker intends to express with the ascription. Even worse, far from identifying the content of a knowledge ascription with the semantic intentions of the speaker, DeRose's version of contextualism appears to allow *no* role for the speaker's semantic intentions in determining content. *Pace* Rysiew, on DeRose's

account the mechanisms that are involved in contextual shifts and content determination are entirely independent of psychological facts about the speaker. To see this we need to look more closely at the details of DeRose's account.

According to DeRose, knowledge ascriptions are subject to the *rule of sensitivity*: when one utters a sentence of the form '*S* knows that *p*' or '*S* does not know that *p*', the sphere of possible worlds relevant to determining the truth-value of the utterance expands, if need be, to include the nearest world *w* in which *p* is false (DeRose 1995, 205-207). An utterance of '*S* knows that *p*' is true only if *S* does not believe that *p* in *w*; correspondingly, an utterance of '*S* does not know that *p*' is true provided that *S* believes that *p* at *w*. This rule can result in a context shift: an expansion of the sphere of relevant worlds corresponds to a shift to more demanding standards for knowledge, because in order for a knowledge ascription to be true, *S*'s beliefs must now 'track the truth' through more worlds than was required in order for earlier knowledge attributions to be true. In addition, DeRose holds that once more demanding standards for knowledge have been imposed, these standards tend to prevail for some time, so that subsequent knowledge ascriptions are also true only if *S*'s beliefs track the truth through the expanded sphere of worlds.

The rule of sensitivity and the 'stubbornness' of high standards are essential to DeRose's response to AI: when the speaker utters (P1), the relevant sphere of worlds expands to include the nearest world *w* in which the speaker is a BIV. In *w*, the speaker believes that she is not a BIV; hence her knowledge denial is true. The elevated standards for knowledge persist as she continues to utter AI; therefore, when she utters (C1), this utterance is true provided that she believes that she has a hand in some world in the expanded sphere in which she doesn't have a hand. The speaker, however, believes that she has a hand in *w*—a world in the expanded sphere

in which she is a BIV and is therefore handless; hence, (C1) is true when the speaker utters it as the conclusion of AI—the result DeRose is after, given his position that AI is sound.

Notice that none of the context-shifting mechanisms discussed above rely on the speaker's intentions or any other fact about her psychology. The rule of sensitivity alone can engender a context shift, and the rule triggers a context shift whenever a knowledge ascription or denial with a certain embedded content is uttered. A speaker might therefore utter 'I know that I am not a BIV', intending to say truly that she *m*-knows that she is not a BIV; nonetheless, the rule of sensitivity would trigger a shift to higher standards for knowledge, and the utterance would turn out to falsely express that the speaker *h*-knows that she is not a BIV. Furthermore, for all DeRose says, the principle that high standards are relatively hard to dislodge does not depend on the speaker's intentions. A speaker might intend to say that she *m*-knows that she has hands, but fail to do so because the high standards in place have not yet been dislodged; in such circumstances the speaker would unintentionally and falsely say that she *h*-knows that she has hands.

Rysiew's argument, then, is ineffective against DeRose's version of contextualism, because the mechanisms that initiate and preserve a context shift operate independently of the speaker's semantic intentions. Furthermore, it is doubtful that Rysiew's argument is effective against *any* version of contextualism, since no version of contextualism that I know of simply identifies a speaker's semantic intentions with the content of her knowledge ascription.

8. Restricting the contextualist's options

I have argued that neither Schiffer nor Rysiew succeed in presenting a convincing objection to contextualist replies to skepticism in general, or DeRose's response in particular.

Here I take a different approach, which is nevertheless informed by Schiffer's and Rysiew's efforts. I will argue that in order to satisfy the criteria for an adequate reply to AI outlined in section 2, a DeRose-style response to AI must attribute a particular set of semantic intentions to a speaker who utters AI and experiences the characteristic paradoxical intuitions; specifically, the reply must claim that the speaker intends to express MAI and falsely believes that AI expresses MAI.

According to DeRose, we incorrectly judge that AI's conclusion is false because we recognize that (C1) would be false if uttered in an ordinary context. Such a mistake, however, does not seem at all easy to make when it comes to words that are widely acknowledged to be context-sensitive, such as 'here' and 'tall'. Consider, for example, the following two cases:

Case 1.

It is winter, and you are in your living-room sitting by a roaring fire. Feeling warm, you say 'It is warm here'. As you utter these words, however, you recognize that you will soon need to go outside to get more wood for the fire, and in that context, the words you are speaking would express a falsehood. You therefore judge that even now, in the warmth of your living room, your utterance expresses a falsehood.

Case 2.

Your 6-foot tall friend Hal is a superb amateur basketball player. Wondering whether Hal could make it in the NBA, you begin to make a list of pros and cons. Realizing that Hal is really quite short for an NBA player, you say, 'Hal is not tall'. However, you realize that in a different context—say one in which you want Hal to get a box from the top shelf—it would be false to say 'Hal is not tall'. You

therefore judge that even now, wondering about Hal's prospects in the NBA, your utterance expresses a falsehood.

It is obvious that an ordinary English speaker like you could not make the mistakes involved in these cases. In both cases, you would know that the context-sensitive word in question *is* context-sensitive, and you would know what it expresses in your present context, and what it would express in a different context. To suppose that what the word would express in a different context is what it expresses in your present context is therefore inconsistent with something that you know. So you could not make this sort of mistake unless you were unusually distracted or irrational. But you are—I trust—neither of these things; so you couldn't make this sort of mistake.

For DeRose's explanation of the implausibility of (C1) in the context of AI to work, then, he must hold either that we don't know that (C1) expresses different propositions in the context of AI and in an ordinary context, or that we are inattentive or irrational. Otherwise, these cases really would be analogous to the case of AI, and we could argue, *à la* Schiffer, that DeRose's imputes to us a mistake that is impossible for us to make, given what we are like.

We can quickly set aside the possibility that the intuition that (C1) is false stems from a failure to pay attention or from irrationality. As discussed in section 2, any adequate reply to AI must be consistent with the assumption that the paradoxical intuitions elicited by AI are universal; therefore, if DeRose's explanation for the implausibility of (C1) rests on a failure to pay attention, it must rest on a universal failure to pay attention. But it seems absurd that an argument should have the power to distract all those who consider it. By the same criterion, if DeRose's explanation for the implausibility of (C1) rests on some sort of irrationality, it would have to rest on universal irrationality. But universal irrationality seems implausible on its face;

besides, an adequate reply to AI would also have to explain the nature and causes of this universal irrationality. I doubt contextualists would find these directions very tempting.

Consider it settled that when you utter AI, you experience the characteristic paradoxical intuitions in spite of being attentive and rational. To avoid a Schiffer-style argument by analogy with Case 1 or Case 2, DeRose must hold that when you utter AI you don't know that (C1) expresses different propositions in your present context and in an ordinary context.

This claim, however, runs up against an argument to this effect: if you know the linguistic meaning of 'know', and you know that you are not in an ordinary context, then you *do* know that (C1) expresses different propositions in your present context and in an ordinary context.

The argument is best stated in terms of Kaplan's (1989) notion of *character*, presented in chapter 1 as part of the conception of semantic content that I am working with. Recall that an expression's character is a function from contexts of utterance to semantic contents. The character of a non-context-sensitive term is a constant function: every context of utterance gets mapped to the same content. The character of a context-sensitive term is not a constant function: for example, if I say 'I', the content of 'I' in that context is me; if you say 'I', the content of 'I' in that context is you; thus, the character of 'I' maps the context in which I say 'I' to me and the context in which you say 'I' to you.

For Kaplan, an expression's character, not its semantic content on a particular occasion of use, is to be identified with its linguistic meaning (Kaplan 1989, 520-521). This makes sense: it seems that a person can know what 'here' means even if she is locked in the trunk of a moving car and doesn't know what 'here' refers to when she utters it; on the other hand, if she doesn't

know that ‘here’ typically refers to the location of the speaker, we would judge that she doesn’t know the meaning of the word.

Back to the argument. If an expression’s character is its linguistic meaning, then knowing the linguistic meaning of ‘know’ consists in knowing the character of ‘know’. The character of ‘know’ is such that it maps high-standards contexts to *h*-knowledge, and ordinary contexts to *m*-knowledge; therefore, if you know the linguistic meaning of ‘know’, then you know that ‘know’ maps high-standards contexts to *h*-knowledge, and ordinary contexts to *m*-knowledge. Since you are attentive and rational, you know that (C1) thereby expresses different propositions in high-standards contexts and ordinary contexts. Suppose you know that when you utter AI, you are in a high-standards context. Then you *do* know (always assuming that you are attentive and rational) that (C1) expresses different propositions in an ordinary context and when you utter AI.

In order to salvage DeRose’s explanation for the implausibility of (C1), a contextualist must respond to this argument. It seems that there are two ways to respond. First, the contextualist could deny that we know the character of ‘know’ when we utter AI. Second, the contextualist could deny that we know that we are in a high-standards context when we utter AI.

Consider first the response that we don’t know the character of ‘know’ when we utter AI. This response fails, because it is inconsistent with the assumption that the paradoxical intuitions elicited by AI are universal. There is no problem with an attentive, rational person simply having a mistaken belief about the meaning of a word in her lexicon, and as a result forming mistaken beliefs about the propositions literally expressed by utterances containing that word; so there is no problem with an attentive, rational person failing to know the character of ‘know’. If the intuitions elicited by AI are universal, however, then *nobody* knows the character of ‘know’. Nobody who considers AI knows the character of ‘know’, since by assumption they all

experience paradoxical intuitions which are explained, according to this response, by ignorance of the character of ‘know’; furthermore, there is no reason to think that those who haven’t considered AI possess some quality that gives *them* knowledge of the character of ‘know’; hence, this response is committed to *nobody* knowing the character of ‘know’.¹⁰ But words can’t have characters that nobody knows about. The linguistic meaning of a word is not handed down from on high, in a way that is independent from what we think and how we act; speaking in the most general possible way, the linguistic meaning of a word is determined by the beliefs and practices of actual speakers; so, if ‘know’ expresses *h*-knowledge in high-standards contexts and *m*-knowledge in ordinary contexts, this fact must be determined by our beliefs and practices. It seems incredible, though, that these beliefs and practices could determine the linguistic meaning of ‘know’ without those whose activities determine the linguistic meaning coming to know it.

It might be replied that a person can fail to have fully explicit, propositional knowledge of the character of ‘know’, even if she engages in the activities that determine its character. That seems right—but she can’t fail to have tacit, implicit knowledge of the character of ‘know’ and engage in such activities; in fact, it is arguable that engaging in such activities *constitutes* such tacit knowledge. Tacit knowledge of the character of ‘know’ would be manifested in a tendency to assent to (C3) in high-standards contexts, and to deny it in ordinary contexts; so, if one had tacit knowledge of the character of ‘know’ when considering AI, this would be manifested by a tendency *not* to have the paradoxical intuitions universally elicited by AI. The contextualist response under consideration must therefore deny that anybody even has tacit knowledge of the character of ‘know’; but this is inconsistent with the fact that the character of ‘know’ must be grounded in activities of actual speakers that would result in, and arguably constitute, tacit knowledge of the character of ‘know’.

Note that the preceding argument relies on the claim that it is impossible for every member of a language community to fail to have even tacit knowledge of an expression's character, and not on the broader assumption that it is impossible for every member of a language community to fail to have even tacit knowledge of an expression's *application conditions*. This latter assumption is certainly false, since it is possible for every member of a language community to be ignorant of, or even radically mistaken about, features of the thing to which their expression refers, and so for every member of the language community to systematically misapply the term.¹¹ Consider the following case, for example:

Warm Earth

Warm Earth is a planet much like our own. In particular, the same natural kinds found on Earth can be found on Warm Earth; like Earth, the oceans, lakes, and rivers of Warm Earth are filled with water. Furthermore, people on Warm Earth speak the same languages as we do on Earth: English, French, Mandarin, etc. There are two important differences between Earth and Warm Earth: first, the inhabitants are technologically backwards (so no freezers); second, the temperature on Warm Earth is always significantly warmer than freezing. Because of this, the inhabitants of Warm Earth have never seen ice. It seems possible in this case for every English speaker on Warm Earth to agree that 'water' never applies to anything solid; they might even agree that water is essentially liquid or gaseous, and that it follows from the meaning of 'water' that it cannot be solid. Now suppose that Warm Earth undergoes global climate change, and gets much colder. The English speakers on Warm Earth begin to notice that the water in their oceans, lakes, and rivers is transformed into a hard, translucent substance when it

gets sufficiently cold. They conclude that this substance is water in its solid form, and that all their previous claims that water cannot be solid were mistaken.

Before Warm Earth started cooling, there was nothing in the beliefs and practices of the English speakers on Warm Earth to indicate even tacit knowledge that ‘water’ might apply to solid things in certain cases. Nevertheless, it was true even before the cooling began that ‘water’ applied to certain solid things, since after the cooling the English speakers admit that they were *mistaken* when they made claims like ‘water is necessarily non-solid’. Here, then, is a case in which every member of a language community fails to have even tacit knowledge of the application conditions of an expression in their language.

Ignorance of an expression’s character can lead to ignorance of the expression’s application conditions.¹² Furthermore, the Warm Earth case shows that universal ignorance of an expression’s application conditions is possible. But none of this shows that universal ignorance of an expression’s application conditions can result from universal ignorance of the expression’s character. I want to continue to insist that this is impossible, because universal ignorance of an expression’s character is impossible.

In the Warm Earth case, the English speakers on Warm Earth were all mistaken about the application conditions of ‘water’. What explains this mistake? It appears that in this case, the application conditions of ‘water’ are not entirely determined by the linguistic beliefs, intentions, and practices of the English speaker on Warm Earth; the nature of water itself is involved in determining the application conditions of ‘water’. In other words, the application conditions of ‘water’ do not supervene on use. It is the fact that the application conditions do not supervene on use that allows for universal ignorance of the application conditions. Cases like Warm Earth, then, should not motivate us to suspect that there can be universal ignorance of an expression’s

character unless there is reason to think that an expression's character can fail to supervene on use. I cannot imagine, however, what factor external to the linguistic beliefs, practices, and intentions of English speakers, could make it the case that an expression expresses such-and-such content in one context, but some other content in a different context.

The idea that an expression's character fails to supervene on use is especially ludicrous when we apply it to standard indexicals like 'I'. Imagine that, although our collective use determines that in almost all contexts 'I' refers to the speaker, some factor external to our use makes it the case that whenever anybody utters 'I' in the presence of Balthazar, 'I' refers to Balthazar. So, when I say 'I am hungry' in Balthazar's presence, the semantic content of my utterance is that Balthazar is hungry. Of course I do not know that this is so, because I am wholly ignorant of this external factor and its power to determine the content of 'I' in certain contexts. So I, and everybody else, proceed to use 'I' as if it always refers to the speaker of the context. Nevertheless, 'I' occasionally refers to Balthazar rather than the speaker. This scenario is ludicrous, not just because we would like to know what is so special about Balthazar, but also because, first, if we all proceed to use 'I' as if it always refers to the speaker of the context, then that fact appears sufficient to make it the case that 'I' always refers to the speaker of the context, and, second, because it is utterly mysterious what sort of factor or condition other than use could make 'I' acquire this odd character.

It is deeply implausible that an expression's character could fail to supervene on use. If that is right, then cases like Warm Earth, which show that universal ignorance of an expression's application conditions is possible, give us no reason to suspect that universal ignorance of an expression's character is possible.

Having satisfactorily established that the paradoxical intuitions elicited by AI cannot be explained by postulating ignorance of the character of ‘know’, let us consider instead the explanation that the paradoxical intuitions result from ignorance of the fact that one is in a high-standards context when one utters AI. As discussed in section 6, contextualist accounts will differ on the question of what features of a context make it a high-standards context in which ‘know’ expresses *h*-knowledge. Without a specific contextualist account of the features that make for a high-standards context, there is no reason in principle to deny that a rational, attentive speaker who knows the character of ‘know’ might nevertheless falsely believe that these features are not present when she utters AI. Furthermore, it has already been shown that the particulars of DeRose’s version of EC allow for such a mistake. Instead of futilely arguing that such a mistake is impossible, I suggest we focus instead on the question of how to characterize the semantic intentions of a rational, attentive speaker who knows—though perhaps only tacitly—the character of ‘know’ and believes—falsely, according to the contextualist—that she is in an ordinary context as she utters (C1) as the conclusion of AI.

When one is speaking literally, one intends to express a certain proposition *p*, by uttering a declarative sentence whose semantic content (in the context *c* that one believes oneself to be in) is *p*. If one then mistakenly utters a sentence whose semantic content in *c* is not *p*, then that mistake entails that one did not know (tacitly or otherwise) the sentence’s semantic content in *c*. In other words, if one *does* know (tacitly or otherwise) that the semantic content of *s* in *c* is *p*, and if one is speaking literally and believes oneself to be in *c*, then one will not utter *s* unless one intends to express *p*.

We can apply this plausible result to the case of the speaker who knows (though perhaps only tacitly) the character of ‘know’, and who utters (C1) as the conclusion of AI while believing

herself to be in an ordinary context. Since she has at least tacit knowledge of the character of ‘know’, she knows, at least tacitly, that the semantic content of (C1) in an ordinary context is the proposition that she does not *m*-know that she has hands. Furthermore, we have every reason to think that she is speaking literally as she utters (C1). It therefore follows from the plausible result above that by uttering (C1) the speaker intends to be expressing the proposition that she does not *m*-know that she has hands.

So far, all this accords well with the speaker’s intuition that her utterance of (C1) is false—it *would* be false, if it expressed the proposition that the speaker does not *m*-know that she has hands—and DeRose’s claim that the speaker judges (C1) to be false because she recognizes that it would be false if uttered in an ordinary context—since she believes that she is in an ordinary context, *of course* her judgment about the truth-value of her utterance of (C1) will match her judgment about the truth-value of (C1) when uttered in an ordinary context.

The problems for a DeRose-style response begin when we try to characterize the speaker’s semantic intentions in uttering the *premises* of AI. Suppose the speaker knows that she is in a high-standards context when she utters (P1) and (P2); since she is speaking literally her semantic intentions in uttering AI would be expressed by the following argument:

Invalid Argument from Ignorance (IAI)

(HP1) I don’t *h*-know that I am not a BIV.

(HP2) If I *h*-know that I have hands, then I *h*-know that I am not a BIV.

(MC1) I don’t *m*-know that I have hands.

But this argument is clearly invalid. Since IAI is supposed to express the speaker’s semantic intentions, and hence her thought process as she utters AI, attributing these semantic intentions to her contradicts the assumption that she is rational and attentive.

Only one plausible option remains: perhaps the speaker falsely believes that she is in an ordinary context as she utters (P1) and (P2), as well as (C1). Then the semantic intentions underlying her utterance of AI can be represented as follows:

(MP1) I don't *m*-know that I am not a BIV.

(MP2) If I *m*-know that I have hands, then I *m*-know that I am not a BIV.

(MC1) I don't *m*-know that I have hands.

But this is just MAI.

The upshot is that a DeRose-style explanation for the implausibility of AI's conclusion either fails to adequately respond to AI or must claim that the speaker's semantic intentions in uttering AI are captured by MAI.

9. The paradox presented by MAI

I've argued that a DeRose-style response to AI can succeed only if it holds that what the speaker intends to express by uttering AI is an argument equivalent to MAI. In the envisaged situation, the speaker takes herself to be speaking literally, and simply has false beliefs about what propositions AI literally expresses. Now, the following general principle seems obviously true:

If an attentive, rational speaker *S*

- (i) utters *u*,
- (ii) is attempting to speak literally in uttering *u*,
- (iii) believes that *u* literally expresses a proposition *p*, and
- (iv) judges that *u* is true/false,

then *S* judges that *p* is true/false.

In uttering (P1), the speaker is attempting to speak literally and takes her utterance of (P1) to express the proposition that she does not *m*-know that she is not a BIV. Since she judges (P1) to be true, it follows from our principle that she judges that she does not *m*-know that she is not a BIV.

Recall that in section 6 MAI was introduced as a *prima facie* challenge to the contextualist: a skeptical argument left untouched by a DeRose-style response to AI whose conclusion, if true, threatens the truth of our ordinary claims to know. It was noted that MAI may not be as challenging to contextualists as AI is to invariantists, because contextualists can plausibly deny that (MP1) is intuitively compelling. But it has now been shown that a DeRose-style response to AI must concede that the proposition expressed by (MP1), that I don't *m*-know that I am not a BIV, *is* intuitively compelling. Further, it is intuitively compelling in just the same way that (P1) is intuitively compelling, since our judgment that (P1) is true *stems* from our judgment that we don't *m*-know that we are not BIV's, and our mistaken belief that (P1) expresses this proposition.

Therefore, after a DeRose-style response to AI has done its work, the contextualist is left with MAI: a valid argument whose conclusion threatens our ordinary claims to know and whose premises are intuitively compelling. So a DeRose-style response to AI does not succeed, since after it has done its work a challenging skeptical paradox remains.

One might wonder whether the thesis of contextualism has the resources to respond to the paradox posed by MAI. It appears not. First, MAI does not contain the word 'know', so it is hard to see how a semantic thesis about 'know' could help. Second, it has been shown that a DeRose-style response is committed to claiming that the *proposition* expressed by (MP1) is intuitively compelling. It is plausible that a semantic thesis about a word's context-sensitivity could explain

why one has a misleading intuition about the truth value of a sentence uttered in a context (by making a mistake about the content expressed by the word in question); it is highly implausible, however, that a thesis about the context-sensitivity of a word could do any work in explaining why we have an intuition about the truth of a proposition that itself has nothing to do with words.

I have argued that a DeRose-style response to AI cannot succeed because after it has done its work, the contextualist is left with a skeptical paradox as challenging to him as AI is to the invariantist. Furthermore, the thesis of contextualism is of no help in addressing this new paradox. Suppose, then, that the contextualist finds non-contextualist resources to answer MAI; he presents truth-conditions for (MP1) according to which (MP1) turns out to be false, and he plausibly explains the misleading intuition that one doesn't *m*-know that one is not a BIV. It appears that the invariantist can simply appropriate this part of the contextualist's response and use it as an equally acceptable response to AI. After all, from the invariantist's perspective, the difference between AI and MAI is purely terminological; so, if the contextualist can adequately respond to the paradoxical intuitions elicited by MAI without using the resources of contextualism, he will have done just what the invariantist needed to do in order to adequately reply to AI. So, if the contextualist were to offer an adequate reply to AI (which would require much more than just a DeRose-style response to AI), he would have furnished the invariantist with a response that is at least as good. To sum up: unless contextualists can devise plausible alternatives to DeRose's approach, the problem of skepticism gives us no reason to choose contextualism over invariantism.

¹ Cohen (1999, 77-79) tacitly endorses DeRose's approach by defending it from the attacks of Schiffer (1996), and Neta (2003a, 26) also appears to endorse DeRose's broad approach.

² I think the assumption that the paradoxical intuitions elicited by AI are universal is widely-shared; personally, I think it is probably true. It is, however, the sort of claim that really stands in need of empirical confirmation.

³ Otherwise, the response hasn't dealt with those who experience the paradoxical intuitions but don't have the deficiency in question. One could, however, offer two or more distinct explanations for the paradoxical intuitions which collectively explain everybody's intuitions.

⁴ A case in point is DeRose (2004).

⁵ In what follows I will focus exclusively on single-person cases, which seem to me to present a greater challenge to EC than two-person cases.

⁶ The strategy of denying AI's validity, while not explicitly stated or distinguished from DeRose's approach, is at least suggested by Cohen (1988, 110-111) and Neta (2003b, 401-402).

⁷ I will adopt the simplifying assumption that according to EC 'know' only expresses two relations: *h*-knowledge and *m*-knowledge.

⁸ For clarity, I have replaced DeRose's 'C' and 'O' with '(C1)' and 'ordinary proposition', respectively. The unaltered quote is: "The very strong pull that (not-C) continues to exert on (at least most of) us even when the standards are high is explained in the manner outlined in section 2: Even while we're in a context governed by high standards at which we don't count as knowing that O, we at the same time sense that as soon as we find ourselves in more ordinary conversational contexts, it will not only be true for us to claim to know these very Os that the skeptic now denies we know, but it will be wrong for us to deny that we know these things. It's easy, then, to think that the skeptic's present denial must be equally false and that it would be equally true for us now, in the skeptic's presence, to claim to know that O."

⁹ In addition to this "hidden indexical" account, Schiffer considers other ways of fleshing out EC's bare claim that 'know' is context-sensitive, but this discussion need not concern us here.

¹⁰ One might demur by claiming that it is the act of considering AI that causes one to *lose* knowledge of the character of 'know', which most people otherwise have. This looks unpromising, however, since one would need to explain how an argument could deprive someone of knowledge of a word's linguistic meaning.

¹¹ See Burge (1986, 714-718).

¹² For example, a small child might fail to recognize that the context-sensitivity of 'tall' is such that it can, in the right context, apply to dandelions. If an adult said, of a dandelion three inches taller than any of the surrounding dandelions, 'that is a really tall dandelion', the child might object, on the grounds that the dandelion is much shorter than she is, and that she is not tall.

CHAPTER 3

CONTEXTUALISM AND THE LOTTERY

1. Introductory remarks

In the last chapter, I considered the sort of response EC affords to AI, a skeptical argument involving the radical skeptical hypothesis that one is a BIV. Similar contextualist replies could be provided to the Descartes-inspired skeptical argument exploiting the radical skeptical hypothesis that one is deceived by an evil demon into falsely believing that there are material objects. What makes the skeptical hypotheses involved in these arguments *radical* is that they involve a bizarre scenario in which the world is drastically different from how we take it to be, and yet we cannot even conceive of evidence that would allow us to rule out the possibility that this bizarre scenario is unfolding in the actual world (Cohen 1988, 111). Although it is not easy, to say the least, to explain how one can know that these radical hypotheses do not obtain, it is also quite difficult to take them seriously—it would be irrational, it seems to me, to seriously wonder whether one is being deceived by an omnipotent demon. As DeRose (1999, 3-4) observes, the philosophically uninitiated tend to respond to skeptical arguments involving radical hypotheses with something along the lines of “give me a break” or “aw, come on!” Though I don’t think that this sort of response is philosophically adequate in itself, it may well be motivated by considerations that could, at the end of the day, show why skeptical arguments involving radical hypotheses fail. Why, after all, should a possibility that it is irrational to take seriously nevertheless be such as to prevent us from knowing anything whatsoever about the external world?

What I am trying to suggest is that if skeptical arguments involving radical hypotheses do ultimately fail, then that failure may well be due to the fact that the arguments employ radical

hypotheses. Skeptical arguments that do not involve radical hypotheses might not succumb to the same critique. Thus, skeptical arguments that do not employ radical hypotheses may ultimately be more effective than skeptical arguments involving radical hypotheses, both because they are more apt to be taken seriously—they are less likely to provoke the “give me a break” response—and because they are immune to criticisms based on the *radical* nature of radical hypotheses.

In this chapter I consider a class of skeptical arguments that are structurally similar to AI, but where the radical hypothesis has been replaced by a possibility that we would all recognize as genuine, if extremely improbable: either the possibility that a certain lottery ticket is the winning ticket, or possibilities that don't involve actual lotteries, but with respect to which we are in essentially the same epistemic position as we are with respect to the possibility that a certain lottery ticket will win. It is widely thought that we cannot know that such possibilities do not obtain on purely statistical grounds, and it is this conviction (the *lottery conviction*) that motivates the skeptical arguments to be discussed in this chapter.

These skeptical arguments involving lotteries or lottery-like situations are relevant to the topic of this dissertation because, once again, EC is alleged to provide a satisfying response to lottery arguments that respects our intuitions and preserves the truth of our ordinary knowledge attributions. If invariantism cannot provide a satisfying response of its own, then that is one good reason to prefer EC to invariantism.

This simple picture is complicated by the fact that one brand of invariantism, SSI, also claims to offer an adequate response to skeptical arguments involving lotteries. If that is right, then the contextualist response to these skeptical arguments does not constitute a reason to prefer EC to invariantism. Since, however, my ultimate goal is to defend *moderate* invariantism, the fact (if it is a fact) that EC and SSI both provide their supporters with adequate replies to the

arguments is cold comfort for me—if moderate invariantism cannot adequately respond to the arguments but EC and SSI can, then that is a good reason to reject moderate invariantism in favor of one of the alternatives.

My plan in this chapter is to defend a moderate invariantist response to skeptical arguments involving lotteries. This response requires a highly unorthodox move: I must deny the lottery conviction. Many, I think, would view this move as simply a non-starter. For one thing, the lottery conviction seems to enjoy strong intuitive support. Furthermore, the lottery conviction does not rest on intuition alone—if one attempts to deny it, one immediately faces two potent objections. First, if one could know that a lottery ticket was a loser on the basis of the odds alone, then that knowledge would seem to rationalize actions that are manifestly *not* rational—for example, it would seem rational to throw the ticket away before the drawing, or to sell it for a penny, regardless of the expected value of the bet. Second, there is a version of Kyburg's (1961) lottery paradox that applies specifically to knowledge.¹ One of the premises that generate the knowledge version of the lottery paradox is precisely the claim that one can know that a certain lottery ticket is a loser on the basis of the odds alone. So there appears to be an easy reply to the knowledge version of the lottery paradox, provided that one is willing to deny this premise. If we uphold the premise, however, then it is unclear how to avoid the lottery paradox's seemingly unacceptable consequences.

These are legitimate worries, but I believe there are good replies to all of them. At the end of the day, I believe that one can defensibly reject the lottery conviction and provide an adequate moderate invariantist response to the skeptical arguments. In sections 2–5, I discuss the lottery conviction, the skeptical arguments, and the EC and SSI based replies to the arguments in

greater detail; then, in sections 6–7, I present and defend a moderate invariantist reply to the skeptical arguments based on denying the lottery conviction.

2. The lottery conviction

One of the most widely held and seemingly unshakable convictions in epistemology is that one cannot know that a certain lottery ticket will lose if one's only reason to believe that it will lose is that the probability that it will win is exceedingly small; one finds it very frequently stated that this is true no matter how small (short of 0) one makes the probability of winning.

This claim about lottery tickets generalizes in important ways. There are many other propositions that, while not specifically about lotteries, are importantly similar to the proposition that a certain lottery ticket will lose. For example, Vogel (1990, 16) claims that if you park your car in a high crime area and then go to the theater, then you have in effect entered your car in a lottery, the “winner” of which gets stolen.

According to Vogel, the proposition that your car has not been stolen resembles the proposition that a certain lottery ticket will lose in the following respects:

- (i) The possibility that your car has been stolen is not “abnormal, in some intuitive sense.”
- (ii) There is no special reason to think that your car in particular will not be stolen; it would be arbitrary for you to believe that your car will not be stolen unless you were also willing to believe, of each of the other cars parked in that neighborhood at that time, that it will not be stolen.
- (iii) There is “some statistical evidence” in favor of the proposition that your car has been stolen. (Vogel 1990, 16-17)

Vogel calls all propositions that share these features *lottery propositions*, and holds that one cannot know that a lottery proposition is true.

Notice that it follows from clauses (ii) and (iii) that whether a proposition is a lottery proposition depends not just on its content, but also on a subject's epistemic position with respect to it. Therefore, the same proposition (individuated by content) could be a lottery proposition for one subject, but not for another who is in a different epistemic position with respect to the proposition. For example, you can know that your car has been stolen if you walk back from the theater and see that your car is gone, and you can know that your lottery ticket is a loser if you see the drawing on television; when you are in those situations, however, these propositions are not lottery propositions for you.

Although there may be room for quibbles about Vogel's analysis of lottery propositions (for one thing, certain expressions in the analysis, such as 'abnormal', are rather unclear), it is unquestionably part of epistemological orthodoxy that there are propositions that are not specifically about lottery tickets but that are importantly similar to the proposition that a certain lottery ticket will lose, and that in virtue of this similarity one cannot know that any such lottery proposition is true. The *lottery conviction* is the widely held view that one cannot know that a lottery proposition is true.

3. Lottery Arguments

The lottery conviction can be used to generate skeptical arguments that are structurally similar to AI. The idea is that for a great many of the propositions that one would ordinarily claim to know, the putatively known proposition, *o*, entails a lottery proposition, *l*; since, according to the lottery conviction, one cannot know *l*, it seems to follow from epistemic closure

that one does not know o . So the lottery conviction motivates skeptical arguments with the following structure:

I do not know l (where l is a lottery proposition);

If I do not know l , then I do not know o (closure);

So, I do not know o (where o is a proposition we would ordinarily claim to know).

I will call skeptical arguments of this type *lottery arguments*. Let's consider some examples from the literature. The first, due to Hawthorne, involves an actual lottery:

Safari

I do not know that Sam will not win the lottery.

If I do not know that Sam will not win the lottery, then I do not know that Sam will not have enough money to go on an African safari this year.

So, I do not know that Sam will not have enough money to go on an African safari this year. (Hawthorne 2004, 2-3)

Suppose Sam is quite poor and that he has no immediate prospects (other than his lottery ticket) for making money. Then it seems that we could correctly claim to *know* that Sam will not have enough money to go on an African safari this year. The preceding argument appears to show otherwise.

The second example, due to Vogel, does not involve an actual lottery:

Car Theft

I do not know that my car has not been stolen and driven away from Avenue A, where I parked it.

If I do not know that my car has not been stolen and driven away from Avenue A, then I do not know that my car is parked on Avenue A.

So, I do not that my car is parked on Avenue A. (Vogel 1990, 20)

Again, though it would ordinarily seem unobjectionable to claim to know where one's car is parked even when one is not currently looking at it, the lottery conviction allied with epistemic closure seems to show that this claim would be false.

Although it would not normally be out of bounds to claim to know that Sam will not be able to afford an African safari, or to claim to know where one's car is parked, it must also be said that the conclusions of these arguments, unlike AI's conclusion that one does not know that one has hands, do not strike us as ridiculous or plainly false. In fact, it is fairly easy to get oneself into a frame of mind where the arguments' conclusions seem exactly right, and the arguments seem to be doing an especially effective job of illuminating *why* we fail to know certain propositions. Isn't it just obviously wrong that one could know that Sam will not be able to afford a safari, if he's got a lottery ticket in his pocket that just might win? If I've been away from my car for a few hours and I don't know that a thief hasn't driven it away, how on earth could I know that it is still where I parked it? In sharp contrast to how we tend to respond to AI, there is a strong temptation (at least on my part) to simply accept these arguments' conclusions.

If, however, one accepts the conclusions of the preceding arguments, then one seems forced to accept the conclusions of other lottery arguments with less appealing conclusions, and eventually to accept a quite far-reaching form of skepticism. Vogel (1990, 20-21) was the first to point out that lottery arguments seem to lead to *semi-skepticism*, the view that we do not know much of what we take ourselves to know about those parts of the world that we are not currently perceiving. Vogel claims that for many such propositions that we take ourselves to know, there is a corresponding lottery proposition that is entailed by the putatively known proposition;

consequently, a lottery argument can be constructed concluding that the putatively known proposition is in fact not known. Here are two more of Vogel's examples:

President

I do not know that George Bush has not had a fatal heart attack in the last five minutes.

If I do not know that George Bush has not had a fatal heart attack in the last five minutes, then I do not know that George Bush is the current President of the United States.

Therefore, I do not know that George Bush is the current President of the United States. (Vogel 1990, 20)

Meteorite

I do not know that the Golden Gate Bridge was not just demolished by a falling asteroid.

If I do not know that the Golden Gate Bridge was not just demolished by a falling asteroid, then I do not know that it stands at the mouth of the San Francisco Bay.

Therefore, I do not know that the Golden Gate Bridge stands at the mouth of the San Francisco Bay. (Vogel 1990, 21)

Vogel is not too clear about how much of our putative knowledge about the world that we are not immediately perceiving is threatened by lottery arguments—he just says that much of it is. However, it would be a mistake to say that all, or almost all, of our supposed knowledge of those parts of the world that we are not immediately perceiving is threatened by lottery arguments. This is because there are also a great many propositions about those parts of the

world that we are not immediately perceiving that appear *not* to be susceptible to lottery arguments.

For example, at the moment I am typing in the den of my house, and cannot see what is going on in my living room. I think I know, however, that my living room's walls are painted blue. It does not appear that there is a lottery argument that applies to this bit of putative knowledge, at least if we use the preceding examples as our model. All of those examples involve some event with the following characteristics: first, while perhaps very rare, events of that type do undeniably occur; further, if the event had occurred, the putatively known proposition would be false; finally, if the event had occurred, it would occur without one's immediate knowledge (one might eventually come to know that it had occurred, but one would not know it at the time at which one is supposed to have the knowledge). Is there an event like that with respect to the proposition that my living room's walls are blue? If an asteroid, airplane, or other flying object had hit that wing of my house and destroyed my walls, I would have noticed. Having one's walls silently disappear or having them repainted by preternaturally fast burglar-housepainters are not the sorts of event that even occasionally happen. The same goes for skeptical hypotheses that cast doubt on the veracity of memory or perception (e.g., the hypothesis that I am a brain in a vat or that I was slipped some memory altering pills). Going by the examples of lottery arguments we have considered so far, then, there does not appear to be a lottery argument that applies to my putative knowledge that my living room's walls are blue.

In the preceding example, the sort of event that could make the putatively known proposition false is of the type whose occurrence one would have noticed. In addition to this kind of example, propositions supported by our knowledge of what certain types of things cannot do, or what they typically do, seem to be immune from lottery arguments. For example, though I am

not presently observing any dogs, I still know both that no dogs are playing violin concertos today, and that some dogs are barking. I also know various propositions supported by my knowledge of human nature—I know, for example, that the president will not deliver his next televised press conference in the nude. Finally, I seem to know a great deal about the state of the world beyond my immediate surroundings in the recent past. Let's say it would take less than an hour for news of the destruction of the Golden Gate Bridge to reach me. Then, since I have not received such news, I know that the Golden Gate Bridge was standing at the mouth of the San Francisco Bay an hour ago.

So, while lottery arguments of the sort we have discussed so far seem to show that we know far less about the world beyond our immediately perceived environments than we ordinarily would claim to know, there are also very many propositions about the world beyond our immediately perceived surroundings that do not appear to be threatened by these sorts of arguments. Furthermore, each lottery argument concluding that we don't know something comes with its own rationale for why we don't know it. If there is some good reason to reject the rationale, then we can reject the argument's conclusion. If there is no good reason to reject the rationale, then it seems indefensible to continue to insist that one knows the proposition in question.

The standard approach to responding to lottery arguments has paralleled the standard to approach to responding to AI: to view the semi-skepticism that follows from lottery arguments as a completely unacceptable consequence, and thus to seek a "solution" to the lottery arguments. However, I think we should seriously consider the possibility that, unlike AI, whose conclusion is blatantly implausible, there is nothing about lottery arguments that requires solving—that

lottery arguments simply uncover the fact that we know quite a bit less about the world beyond our immediately perceivable environments than we would ordinarily be willing to claim.

I have been suggesting that we simply accept the lottery arguments we have considered so far, along with their skeptical conclusions. Not only is the semi-skepticism that follows from arguments of this sort fairly restricted—it is consistent with us knowing a great deal, even about the world beyond our immediate surroundings—also, the lottery arguments we have considered appear to be *good* arguments that shed new light on the limits of our knowledge.

Although the kind of semi-skepticism that follows from the lottery arguments we have considered thus far may be something we can live with, Hawthorne has proposed another kind of lottery argument that, if sound, entails a far more radical form of skepticism:

Suppose there is a desk in front of me. Quantum mechanics tells me that there is a wave function that describes the space of nomically possible developments of the system that is that desk. On those interpretations of quantum mechanics according to which the wave function gives probability of location, there is some non-zero probability that, within a short while, the particles belonging to the surface of the desk remain more or less unmoved but the material inside the desk unfolds in a bizarre enough way that the system no longer counts as a desk. Owing to its intact surface, the system would be reckoned a desk by normal observers. Call such a system a desk façade. I will be ordinarily inclined to think that I know by perception that there is a desk in front of me. But once the question arises, I will be far less inclined to think that I know by perception whether or not this is one of those unusual cases in which the desk recently developed into a desk façade. And, obviously, the example generalizes. (Hawthorne 2004, 4-5)

Hawthorne is offering the following argument:

- (P1) I don't know that my desk has not recently developed into a desk façade.
- (P2) If I don't know that my desk has not recently developed into a desk façade, then I don't know that there is a desk in front of me.
- (C1) So, I don't know that there is a desk in front of me.

A natural response to this argument would be to try to deny (P1). However, if quantum mechanics (as interpreted by Hawthorne) is correct, then there is a non-zero probability that, moments ago, my desk turned into a desk façade. Given that this statistical possibility is a consequence of a well-established scientific theory, it seems wrong to describe it as absurd. Provided that I don't cheat and look behind the desk to make sure it is not a façade, it seems that the only reason I have for believing that my desk has not recently turned into a façade is that the probability of this happening is so incredibly small; thus, my only reason for thinking that my desk has not developed into a façade is also a reason for thinking, of any other object, that *it* has not developed into a façade. It seems, then, that the proposition that my desk has not recently developed into a façade satisfies Vogel's three criteria for being a lottery proposition. If that is right, and Vogel's criteria are correct, then it follows from the lottery conviction that (P1) is true. Of course, one might argue that the proposition that my desk has recently developed into a desk façade is not a lottery proposition, either because it does not really satisfy Vogel's criteria, or because Vogel's criteria are off the mark. I, however, will assume both that Vogel's characterization of lottery propositions is correct, and that the proposition satisfies Vogel's criteria. Given these assumptions, denying (P1) requires denying the lottery conviction.

As Hawthorne points out, the example generalizes. So the lottery conviction seems to entail the view that you can never tell just by looking at something what sort of thing it is. That is a very radical form of skepticism, and about as deeply implausible as AI's conclusion that you don't know you have hands. Unlike the other lottery arguments we have considered, Hawthorne's lottery argument must be responded to somehow if we are to avoid a radical, and profoundly implausible, form of skepticism.

4. EC and lottery arguments

Having presented the skeptical problem posed by lottery arguments, let us now consider the sort of response that EC can provide to these arguments. Two important contextualists, Lewis (1996) and Cohen (1988), have offered responses to lottery arguments based on EC.² I will focus on Lewis's response here—at the appropriate level of generality, Cohen's treatment of lottery arguments is not importantly different from Lewis's. I will explain enough of Lewis's version of EC and his response to lottery arguments to make clear that Lewis's response requires denying the lottery conviction. More precisely—since EC complicates matters by introducing multiple knowledge relations—Lewis concedes that one can know, in the sense of 'know' that figures in our ordinary knowledge attributions, that a lottery proposition is true. In section 7, I will use this fact to defend my own invariantist response to lottery arguments, which also denies the lottery conviction.

Lewis's version of contextualism is a relevant alternatives theory that denies that knowledge requires either belief or justification. His account is built around the following central claim:

S knows that *P* iff *S*'s evidence eliminates every possibility in which not-*P*—
Psst!—except for those possibilities that we are properly ignoring. (Lewis 1996,
225)

As Lewis acknowledges, this is somewhat loosely stated and should really be expressed metalinguistically (Lewis 1996, 237-238). The following claim more accurately captures Lewis's view:

(LK) An utterance of '*S* knows that *p*' is true in a context *c* iff *S*'s evidence can eliminate every possibility in which not-*p*, except those possibilities that the speaker-hearers of *c* are properly ignoring.

Let's translate this into more standard relevant alternatives talk. The not- p possibilities that are *not* being properly ignored are the relevant alternatives which S must rule out in order to know that p ; the not- p possibilities that *are* being properly ignored are the non-relevant alternatives, which S need not rule out in order to know that p . Ruling out is understood as follows: S rules out a possibility w just in case S 's total evidence in w is different from S 's total evidence in the actual world α (Lewis 1996, 224).

As Vogel (1999, 163-168) argues, the greatest challenge for any relevant alternatives theory is providing a plausible and informative account of what makes an alternative relevant in a context. To his great credit, the bulk of Lewis's (1996) paper is devoted to articulating and defending several rules that collectively determine when an alternative is relevant in a context.

Lewis makes use of only two of these rules—the Rule of Actuality and the Rule of Resemblance—in his treatment of the lottery. The Rule of Actuality states that α , the actual world, is always relevant. The Rule of Resemblance states that if w is relevant, and w' saliently resembles w , then w' is itself relevant.

Lewis argues that the Rule of Actuality and the Rule of Resemblance explain why one cannot know that a lottery ticket is a loser on the basis of the odds alone. Lewis is envisioning a lottery in which there is a winning ticket, k ; it follows from the Rule of Actuality that the possibility that k is a winner is relevant. Consider some other ticket that will not in fact win, k' . The possibility that k' is a winner saliently resembles the possibility that k is a winner; so it follows from the Rule of Resemblance that the possibility that k' is a winner is relevant, and is thus a relevant alternative to the proposition that k' is a loser. But this alternative cannot be ruled out, since in the world in which k' wins, one has just the same evidence (the odds) that one has in α . From (LK), it follows that it would be false to say, 'I know that k' is a loser'.³

Now to Lewis's response to lottery arguments. Lewis (1996, 237) gives the example of "poor Bill," an inveterate gambler who "squanders all his spare cash on the pokies, the races, and the lottery." If we know about these bad habits, it seems that we can truly say, 'We know that Bill will never be rich'. But this is at odds with the conclusion of the following lottery argument:

(P3) We don't know that Bill will not win the lottery.

(P4) If we don't know that Bill will not win the lottery, then we don't know that Bill will never be rich.

(C2) We don't know that Bill will never be rich.

As we have seen, Lewis's Rule of Actuality and Rule of Resemblance entail the truth of (P3); further, Lewis is committed to epistemic closure, and so the truth of (P4). Since Lewis is committed to the truth of the argument's premises, and the argument is valid, he must be committed to the truth of (C2). But this seems at odds with Lewis's claim that we can truly say, 'We know that Bill will never be rich'.

Lewis's answer to this problem is to insist that (C2) is true in the context of the preceding lottery argument; nevertheless, we can truly say 'We know that Bill will never be rich' in other, more pedestrian contexts. This might seem puzzling, since given Lewis's commitment to epistemic closure it follows that 'We know that Bill will not win the lottery' is also true in those contexts—but Lewis's Rule of Actuality and Rule of Resemblance were supposed to preclude that. Lewis's reply is that the falsehood of 'We know that Bill will not win the lottery' does *not* follow from the Rule of Actuality and the Rule of Resemblance in pedestrian contexts in which it is true to say 'We know that Bill will never be rich', because in those pedestrian contexts the possibility that Bill will win the lottery is not salient, and so does not *saliently* resemble α .

In pedestrian contexts, then, Lewis holds that ‘We know that Bill will never be rich’ and ‘We know that Bill will not win the lottery’ are both true. Note that Lewis does not hold that it would be true to *say* ‘We know that Bill will not win the lottery’ in a pedestrian context—just the opposite:

It was true at first that we knew that Bill would never be rich. And at that point it was also true that we knew he would never lose—but that was only true so long as it remained unsaid! (And maybe unthought as well.) Later, after the change in context, it was no longer true that we knew he would lose. At that point, it was no longer true that we knew he would never be rich. (Lewis 1996, 237)

Again, it is important not to be misled by Lewis’s reluctance to put things metalinguistically. For complete accuracy, we should understand Lewis as saying that the *sentences* ‘We know that Bill will never be rich’ and ‘We know that Bill will lose the lottery’ are both true in the context in which the possibility that Bill will win the lottery has not been made salient, and are both false in the context in which the possibility that Bill will win the lottery has been made salient. Although ‘We know that Bill will lose the lottery’ is true in pedestrian contexts, it cannot be truly uttered. This is because by uttering the sentence, we make the possibility that Bill will win the lottery salient. The possibility that Bill will win the lottery then *saliently* resembles α , and it follows from the Rule of Actuality and the Rule of Resemblance that the sentence is false. Nevertheless, though the sentence cannot be truly uttered, it remains true in pedestrian contexts where the possibility that Bill will win the lottery has not been made salient.⁴

Lewis’s response to lottery arguments has many virtues: it preserves the truth of epistemic closure, it explains the intuitive appeal of lottery arguments (we find them appealing precisely because we recognize that they are sound), and it preserves the truth of our ordinary knowledge attributions, even though the propositions knowledge of which is attributed entail propositions that we intuitively do not know. All these benefits come at a cost, however: Lewis

explicitly states that we *do* know that Bill will lose the lottery on the basis of the odds alone, when the possibility that Bill will win has not been made salient. So Lewis's response apparently requires denying the lottery conviction.

It is, however, slightly inaccurate or misleading to simply state that Lewis's response denies the lottery conviction and leave it at that. According to Lewis, 'know' can express (at least) two epistemic relations. Consequently, there are really two versions of the lottery conviction, each involving one of those epistemic relations. Let's use '*m*-knowledge' to refer to the relation expressed by 'know' in contexts where the possibility that Bill will win the lottery has not been made salient, and use '*l*-knowledge' to refer to the relation expressed by 'know' in contexts where the possibility that Bill will win the lottery *has* been made salient. The first version of the lottery conviction holds that we can never *m*-know a lottery proposition. Lewis denies this version. The second version of the lottery conviction holds that we can never *l*-know a lottery proposition. Lewis affirms this version of the lottery conviction. What is most important for our purposes is that Lewis concedes that one can *m*-know a lottery proposition; that is, in the sense of 'know' that figures in ordinary knowledge attributions, he holds that one can know that a lottery proposition is true.

5. SSI and lotteries

Hawthorne (2004) proposes a treatment of various epistemological problems involving lotteries—including the skeptical threat posed by lottery arguments—that makes use of SSI.⁵ SSI holds that 'know' expresses the same relation in every context of utterance, but that whether this relation obtains between a subject *S* and a proposition *p* (and thus whether '*S* knows that *p*' is true) depends in part on features of *S*'s context that have not traditionally been thought

epistemically relevant, such as whether a certain possibility is salient to S , or features of S 's practical situation, such as the consequences to S should she turn out to be wrong about p . After presenting the details of Hawthorne's proposal, I will argue that Hawthorne's is committed to denying the lottery conviction. In section 7, I will use this fact to defend my moderate invariantist response to lottery arguments, which also requires denying the lottery conviction.

Hawthorne presents the problem posed by lotteries in a somewhat different light than I have so far. For Hawthorne, the central problem posed by lotteries is preserving three plausible but not obviously consistent claims: first, that we know a great deal, including many propositions that entail lottery propositions; second, that some form of epistemic closure is correct; finally, that one cannot use one's knowledge of a proposition p that entails a lottery proposition q to come to know that q . At first blush, these three claims appear to be incompatible, since if S knows that p and p entails a lottery proposition q , then epistemic closure would seem to tell us that S knows that q . The lottery arguments discussed earlier can be viewed as presenting the incompatibility in a different way: if S does *not* know any lottery propositions, then epistemic closure would seem to tell us that S does not know any proposition that entails a lottery proposition.

Hawthorne would reply that the preceding arguments rely on a rudimentary and quite obviously false version of epistemic closure:

(CL1) If S knows that p , and p entails q , then S knows that q .

This version of the closure principle is subject to several objections. First, if S does not know that p entails q , then it seems quite possible for him to know that p without ever coming to believe, or even think about, q . In such a case it is implausible to think that S would know that q . Even if S does know that p entails q , it seems possible for S not to perform the deduction from p

and p entails q to q , and so not believe that q . Again, it is implausible that S would know q in such a case. These objections suggest the following refinement of the closure principle:

(CL2) If S knows that p , competently deduces q , and thereby comes to believe q , then S knows that q . (Hawthorne 2004, 32-33)

However, Hawthorne argues that even (CL2) is not an entirely accurate statement of epistemic closure. Performing a deduction is a cognitive act that takes a certain amount of time, and it is possible for the epistemic status of a premise to change from t_1 , when the deduction begins, to t_2 , when it ends. In particular, it is possible to cease knowing p by the time one has competently deduced q . In that case, it would be implausible to claim that S knows q just because she has competently deduced q from p and knew p when she began the deduction. S must retain knowledge of p throughout the deduction. This consideration leads to Hawthorne's preferred formulation of epistemic closure:

(CL3) If S knows p , competently deduces q , and thereby comes to believe q , while retaining knowledge of p throughout, then S knows q . (Hawthorne 2004, 34)

Suppose that (CL3) is correct, and that at t_1 a subject S knows a proposition p that entails a lottery proposition q (for example, let p be the proposition that S will not be able to afford an African safari this year, and q be the proposition that S 's lottery ticket will lose). Given these assumptions, must we accept that S knows q ? Hawthorne would say 'no'. Suppose that at t_1 , S is not considering q at all, and *a fortiori* is not deducing q from p ; then it obviously does not follow from (CL3) that S knows that q . Now suppose S performs the deduction from p to q , and—though it is psychologically implausible to suppose that this would happen—thereby comes to believe that q at t_2 . It still does not follow from (CL3) that S knows that q , since S might lose knowledge of p between t_1 and t_2 . So, if performing the deduction from p to q could cause S to lose knowledge of p , then it is plausible to hold that each member of the not obviously

compatible trio is in fact true. Hawthorne's task is to provide a plausible account of how performing this deduction could cause S to lose knowledge of p .

Hawthorne begins by considering whether appealing to salience might deliver the desired account. Before getting to this, a quick word about the concept of salience: though 'salient' is sometimes used to mean no more than 'being consciously entertained', for Hawthorne it takes more than just consciously thinking about a possibility in order to make it salient. Hawthorne (2004, 62) characterizes salient possibilities as those that "one is worried about" or "takes seriously"; he views these two locutions as roughly synonymous.⁶

Here is how Hawthorne proposes to use salience to explain how deducing q from p can deprive S of knowledge that p . According to Hawthorne's version of SSI, if a certain not- p possibility is salient to S , then S does not know that p . By performing the deduction from p to q , the possibility that not- q (that S 's ticket is a winner) becomes, "as a matter of psychological fact," salient to S . Since p entails q , not- q is a not- p possibility; therefore, once S performs the deduction there is a not- p possibility that is salient to S , and so S ceases to know that p (Hawthorne 2004, 161).

It is certainly possible to put pressure on the preceding argument by questioning the version of SSI upon which Hawthorne is relying; for example, one might demand reasons for thinking that the mere salience of a not- p possibility can preclude knowledge that p . Hawthorne (2004, 169-172) addresses these issues and suggests several accounts that explain how salience could preclude knowledge; this discussion, however, is rather inconclusive. Though the issue is far from settled, I will not pursue this line of criticism here.

Hawthorne (2004, 173) acknowledges that the appeal to salience cannot fully explain why deducing q from p destroys S 's knowledge that p , even if his version of SSI is exactly right.

This is because the appeal to salience relied on the claim that when S deduces p from q , not- q becomes salient to S as a matter of psychological fact. But what if S is psychologically abnormal and the possibility that not- q does *not* become salient for her? Suppose that when S deduces q from p , she simply comes to believe that q and doesn't worry about the possibility that not- q at all. Since not- q is not salient to S , she appears to preserve her knowledge that p throughout the deduction; it follows from (CL3) that S comes to know that q .

It might be suggested that Hawthorne's version of SSI could be modified to deal with this sort of problem, by claiming that possibilities that *ought* to be salient to S , and not just those that happen to be salient to S , can preclude knowledge. Since, in the preceding case, the possibility that not- q ought to be salient to S , S does not know that p . The problem with this response is that there is no obvious reason why the possibility that not- q shouldn't be salient to S even before she performs the deduction. If S believes that she will not be able to afford an African safari this year, *shouldn't* the possibility that she will win the lottery be salient to her? If it should, then this modification of Hawthorne's SSI leads to the conclusion that S *never* knows that p . Of course, somebody might deny that the possibility that not- q ought to be salient to S before she performs the deduction, but if a plausible account of when a possibility ought to be salient is not provided, this move will seem hopelessly *ad hoc*.

Recognizing that the appeal to salience does not provide a fully adequate response to the problem posed by lotteries, Hawthorne considers the other "non-traditional" factor that SSI holds can be relevant to whether somebody knows, namely, the subject's practical situation. Consider the following piece of practical reasoning:

(P5) I will be in Morocco all summer.

(P6) If I will be in Morocco all summer, I should get somebody to water my plants while I am away.

(C3) So, I should get somebody to water my plants while I am away.

Since (P5) and (P6) are premises in an argument that draws a practical conclusion, they are *used as premises in practical reasoning*. Hawthorne's idea is that whether one knows that p depends on whether one is entitled to use p as a premise in practical reasoning. Since whether one is entitled to use p as a premise in practical reasoning is highly sensitive to features of the subject's practical situation, knowledge turns out to be highly sensitive to the subject's practical situation.

For example, consider the following, intuitively appalling, piece of practical reasoning:

(P5) I will be in Morocco all summer.

(P7) If I will be in Morocco all summer then I needn't buy flight insurance (because my plane to Morocco won't crash).

(C4) Therefore, I needn't buy flight insurance.

In this case I am not allowed to use the proposition that I will be in Morocco as a premise in practical reasoning; so, on Hawthorne's view, I do not know that I will be in Morocco all summer. However, in the previous piece of reasoning I was entitled to use the same proposition as a premise in practical reasoning; Hawthorne would say that in that case I did know that I would be in Morocco all summer. What makes the difference between the two cases?

Presumably, the difference is that in the first case I am faced with the practical problem of what to do with my plants, whereas in the second case I am faced with a different practical problem, that of deciding whether to buy flight insurance.

The following passage illustrates just how sensitive knowledge is to the subject's practical situation, given Hawthorne's proposal:

Allowing such a mechanism will make knowledge come and go with ease. One is offered a lottery ticket. At that point one doesn't know that one will be unable to afford a trip to Mauritius. One buys a ticket, forgets about the lottery, and goes to the bookstore. One chooses the 'local destination guide' over the much more expensive 'worldwide guide', reasoning from the premise 'I won't be able to afford to go to an exotic destination'. At that point you do know that you will be unable to afford a trip to Mauritius. (Hawthorne 2004, 176-177)

Interesting as it is, it is far from clear that Hawthorne's appeal to practical interests can solve the problem that he started out with. Let's review the problem. Suppose that at t_1 , S knows that p (the proposition that she will not be able to afford an African safari this year), and grant that (CL3) is the correct formulation of the closure principle. Hawthorne wants to ensure that S cannot come to know that q (that she will lose the lottery) by deducing it from p . If S performs the deduction and thereby comes to believe q , the only escape route provided by (CL3) is the possibility that some time after t_1 , S ceases to know that p . We have seen that Hawthorne's appeal to salience cannot explain why S would cease to know that p in those cases where not- q is simply not salient to S . We now have a new possible explanation for how S could cease to know that p : namely, that some time after t_1 , it becomes unacceptable for S to use p as a premise in practical reasoning. However, this new explanation will solve Hawthorne's problem only if there is *no* possible case in which S performs the deduction, not- q never becomes salient to her, and it remains acceptable for her to use p as a premise in practical reasoning throughout the deduction.

Deciding this issue requires answering the question: given that S may use p as a premise in practical reasoning at t_1 , what could make it unacceptable for S to use p as a premise in practical reasoning after t_1 ? The most natural answer to this question, and the one suggested by the passage cited above, is that an issue of some practical import for S arises after t_1 , and that S 's

deduction is a part of her deliberations over this practical issue. To modify one of Hawthorne's examples, suppose that somebody offers to buy *S*'s lottery ticket (expected value: 50 cents) for a penny. *S* could not reason as follows:

- (P8) I will not be able to afford an African safari this year.
- (P9) If I will not be able to afford an African safari this year, then my lottery ticket will lose.
- (C5) So, my lottery ticket will lose.
- (P10) If my lottery ticket will lose, I ought to sell it for a penny.
- (C6) So, I ought to sell my lottery ticket for a penny.

In this example it is quite plausible that after t_1 it becomes unacceptable for *S* to use p as a premise in practical reasoning, since something about *S*'s practical situation changes after t_1 —somebody offers to buy *S*'s lottery ticket for a penny.

If the only explanation for why *S* may not longer use p as a premise in practical reasoning after t_1 is that some new practical issue comes onto the scene, then Hawthorne's appeal to practical interests will ultimately not solve his problem, since we can easily construct a case in which no new practical issue arises for *S* after t_1 . Suppose that for an extended period of time after t_1 , nobody approaches *S* with an offer to buy her ticket, and no other practical issue having to do with either her lottery ticket or a possible African safari arises. Furthermore, the possibility that *S* will win the lottery never becomes salient to her throughout this period—she just never take seriously the possibility that she will win. Let's say this period lasts until t_2 . Since no relevant practical issue arises for *S* during the period from t_1 – t_2 , and since *S* knows that p and may use p as a premise in practical reasoning at t_1 , it is reasonable to expect that *S* may continue to use p as a premise in practical reasoning throughout the t_1 – t_2 interval. But now suppose that *S*

reasons from (P8) and (P9) to (C5) during the t_1 – t_2 interval. (Importantly, S does not draw any practical conclusions from this reasoning; specifically, she does not go on to reason as above to (C6).) Then it seems to follow from (CL3) that S can come to know that q via deduction, since none of the factors that Hawthorne suggests could cause S to cease knowing that p (the salience of not- q or the impermissibility of using p as a premise in practical reasoning) are present. If such a case is possible, then Hawthorne's appeal to practical interests does not solve his problem.

I have been considering the most natural explanation for why it becomes unacceptable for S to use p in practical reasoning after t_1 , which is that some practical issue to which p is relevant arises for S . I have argued that if this is the only explanation, then the appeal to practical interests does not ultimately solve Hawthorne's problem. It might be suggested instead that even if no practical issue arises for S from t_1 to t_2 , the mere act of performing the deduction from p to q makes it unacceptable for S to use p as a premise in practical reasoning. The problem with this explanation is making sense of how performing a deduction could in itself have any effect on one's practical situation.

The problem is especially glaring when the deduction in question involves premises and a conclusion that are entirely removed from anything S cares about or has a practical stake in. For example, consider Hawthorne's quantum physics example. Suppose I know that there is a desk in front of me, and to pass the time I deduce that the desk has not recently become a mere desk façade. Suppose further that the possibility that the desk *has* recently become a desk façade is not salient to me, because I just don't take it seriously.⁷ In order to avoid the conclusion that I can come to know via deduction that the desk has not recently become a desk façade, Hawthorne has to say that my practical situation changes as a result of performing the deduction in such a way that I can no longer use the premise that the object in front of me is a desk in practical reasoning.

But there is no reason to think that performing this deduction should have any effect on my practical situation—I am idly reasoning, with no practical goal in mind, and nothing of practical consequence hangs on the proposition that the desk in front of me has not recently become a desk façade (or, for that matter, on whether the object in front of me is a desk). So it is not at all plausible that performing the deduction alters one’s practical situation in such a way that one can no longer use the premise of the deduction in practical reasoning.

Hawthorne’s version of SSI succeeds in explaining many cases in which one apparently has knowledge of a proposition that entails a lottery proposition even though one cannot deductively acquire knowledge of the lottery proposition. Nevertheless, Hawthorne’s apparatus cannot explain all such cases. His appeal to salience cannot explain cases where the possibility that is supposed to destroy knowledge of the deduction’s premise is simply not salient to the subject; on the other hand, his appeal to practical interests does not help in cases where the subject’s practical situation does not change. Therefore, in many cases Hawthorne’s system appears to allow that we can come to know lottery propositions via deduction. It might be suggested that Hawthorne should just deny that we have knowledge of the premises of the deductions in those cases. However, since one such premise is the proposition that there is a desk in front of me, this would be capitulating to a radical form of skepticism. Hawthorne should instead concede that in many cases we can know the lottery propositions that are entailed by other propositions that we know.⁸ In short, Hawthorne’s SSI-based treatment of the lottery commits him to denying the lottery conviction.

6. A moderate invariantist reply to lottery arguments

In this section I will present a moderate invariantist reply to lottery arguments. My response will require denying the lottery conviction—I claim that, although there may be many lottery propositions that we cannot know, there are also many that we can know. Denying the lottery conviction faces several objections, to which I will respond in section 7. For now, I want to sketch the sort of response to lottery arguments I have in mind, on the assumption that the lottery conviction is false.

In section 3 I discussed several lottery arguments and noted that many of these arguments' supposedly unwelcome consequences were not clearly unacceptable. For example, it is just not clear to me that the conclusion that I don't know where my car is parked (having parked it in an urban area and not having seen it for several hours) is an unacceptable one that an adequate theory of knowledge must ensure comes out false. Similarly, it seems quite plausible to me that most claims about what I or other people will be doing in the future are unknowable, simply because (as lottery arguments make clear) people's plans can change, they can unexpectedly die, and so on. When I consider many of these lottery arguments, I find myself convinced by the arguments, I accept the conclusions, and I am inclined to view those knowledge claims I might have made that conflict with the conclusions as not strictly speaking true (though perhaps communicatively useful, only harmlessly false, and so on). Since I find so many of these arguments convincing, part of my response is to simply accept many of these arguments, and the limited form of skepticism that comes along with doing so.

On the other hand, there are other lottery arguments that I do not find entirely convincing and whose conclusions I am either ambivalent about or I wholeheartedly deny. For example, I find it quite plausible that I don't know what I will be doing six months from now, in spite of my

firm plans, since I don't know that I won't have a heart attack and die during the next six months. On the other hand, I am less inclined to say that I don't know what I will be doing in five minutes; but it does seem that I don't know that I won't have a heart attack and die in the next five minutes. This is a case where I am ambivalent about the argument's conclusion, even though I am fairly strongly inclined to accept the premises and the validity of the argument. Then we have Hawthorne's quantum mechanics argument. In this case I am very strongly inclined to deny the conclusion that I don't know that there is a desk in front of me, as well as the premise that I don't know that my desk has not recently turned into a desk façade.

So, I accept some lottery arguments, I am ambivalent about some others, and I firmly reject still others. What explains these disparate reactions? I want to suggest that my response to a lottery argument hinges on my assessment of the probability of the lottery proposition involved in the argument. When the probability that the lottery proposition is false appears ludicrously tiny, I judge that I know that the lottery proposition is true, and so reject the argument and its conclusion. For example, I find the proposition that my desk has recently turned into a desk façade so utterly improbable that I think I know that my desk has not recently turned into a desk façade. When the probability that the lottery proposition is false appears quite high, I judge that I do not know that the lottery proposition is true. Since the probability that I will die within the next six months seems relatively high to me, I judge that I do not know that I will not die in the next six months, and I accept the conclusion that I do not know what I will be doing in six months. In cases where I am ambivalent, I either do not have an adequate sense of the probability of the lottery proposition, or I do not have an adequate sense of how the probability of the lottery proposition bears on my knowledge of the lottery proposition. With respect to the argument concluding that I do not know what I will be doing in five minutes, either I do not have an

adequate sense of the probability that death or some plan-altering accident will not befall me in the next five minutes, or if I do have some measure of this probability, I do not have an adequate sense of whether the probability is great enough for me to know that it is true.

If my differing reactions to different lottery arguments are explained by my differing assessments of the probabilities of the lottery propositions involved in the cases, then perhaps this is because it really is possible to know a lottery proposition provided that it is sufficiently probable. One natural way of developing this thought would be to hold that there is some probability n such that one can know a lottery proposition q just in case the probability of q is at least n . According to this proposal, lottery arguments involving lottery propositions that are less probable than n are sound; lottery arguments involving lottery propositions whose probability is n or greater are (typically) unsound.⁹ Since Hawthorne's quantum mechanics examples involves a lottery proposition (that my desk has not recently turned into a desk façade) that is exceedingly probable, its probability is presumably greater than n and the argument is unsound. The suggested account can therefore avoid the radical sort of skepticism threatened by Hawthorne's argument. Vogel's car theft argument, on the other hand, involves a lottery proposition (that my car has not been stolen while I was away) whose probability, while still high, falls well short of one, and may fall short of n . If it does fall short of n , then the car theft argument is sound, and I truly do not know where my car is parked.

Note that the claim that one can know a lottery proposition just in case its probability is at least n does not entail that anybody knows exactly, or even roughly, what n is. It also leaves open the possibility that we are not particularly good at arriving at even a rough assessment of the actual probability of a lottery proposition. So the proposal we are considering is consistent

with a great deal of error and ambivalence in our judgments about whether we know lottery propositions, and, consequently, which lottery arguments are sound and which ones are unsound.

One problem for the sort of response to lottery arguments sketched above is that it involves a sharp cutoff n , such that one can know lottery propositions that are at least as probable as n , but one cannot know lottery propositions that are less probable than n . There are at least two reasons why one might find such a sharp cutoff deeply implausible.

First, consider two lottery propositions, q_1 and q_2 . Suppose that the probability of q_1 is exactly n , and suppose that the probability of q_2 is extremely close to, but just shy of, n . The sharp cutoff view presented above says that one could come to know q_1 , but one could not come to know q_2 . But it is implausible that this tiny difference in probability could make the difference between knowing and not knowing. Suppose, for example, that $n = 1 - 10^{-9}$. Then the sharp cutoff view says that one could come to know that one's ticket is a loser on the basis of the odds alone in a lottery with a billion tickets, but that one could not come to know that one's ticket is a loser on the basis of the odds alone in a lottery with a billion minus one tickets. This is hard to swallow—there seems to be no good explanation for why the presence of a single extra ticket in a lottery could make it the case that one can now come to know that one's ticket will not win.

A second, related objection is that whatever n happens to be, there seems to be no non-arbitrary reason to set n at *that* level, rather than slightly higher or lower. Nothing in our intuitive concept of knowledge seems to motivate a certain choice of a precise n . If it is nevertheless insisted that there is such an n , then this will leave it utterly mysterious how n came to have that value, rather than some other.

In light of these objections, the sharp cutoff view seems untenable. This, however, may not spell disaster for the view that one can know a lottery proposition provided that it is

sufficiently probable. Before I explain, I would like to point out that very similar issues arise for any number of very plausible accounts according to which having a certain property F requires being sufficiently G (where G is a property that comes in degrees). In epistemology, there is the view that in order to know that p , one must have sufficient justification for believing p . If this claim is understood to be claiming that there is a certain degree of justification, d , such that one may know that p just in case one's degree of justification for p is at least as great as d , then the very same problems will afflict this view (BonJour 2002, 43). The choice of d seems arbitrary, there is no good account for how d was determined, and there seems to be no good reason for allowing S_1 , who is justified in believing p to degree d , to know that p , but for denying that S_2 knows that p , when S_2 's justification for p falls just short of d . Outside of epistemology, we have the case of the heap. Clearly, in order for a collection of grains to count as a heap, it must contain sufficiently many grains. But if we say that the collection counts as a heap only if it contains at least m grains, we will face the same problems: the choice of m , rather than $m - 1$ or $m + 1$, seems arbitrary; there is no good account of how m was determined; and if a collection with m grains is a heap, there seems to be no good reason for saying that this other collection containing $m - 1$ grains is not a heap. In spite of these problems, we should not deny that in order for something to count as a heap, it must contain sufficiently many grains.

In the case of heap, the natural response is to deny that there is a sharp cutoff m such that x is a heap if and only if x contains m grains, yet insist that in order to count as a heap, x must contain sufficiently many grains. Now this sort of position faces several problems, most notably that it allows the Sorites paradox:

If x contains a million grains of sand, then it is a heap.

If x is a heap, then a collection with one less grain than x is also a heap.

By induction: a single grain is a heap.

Two very plausible premises have led to a patently absurd conclusion. Trying to resolve the Sorites paradox and other related problems that arise when one rejects sharp cutoffs has spawned a huge and sophisticated literature, which falls under the heading of vagueness. This is clearly not the place to take on these issues. My point is just this: the orthodox position in the vagueness literature is that vague terms do not admit sharp cutoffs, and that, while for a vague term v there are clear cases of v 's and clear cases of non- v 's, there are also borderline cases, which are neither clearly v 's nor clearly non- v 's. Against this backdrop of assumptions, most philosophers interested in vagueness then go on to address problems like the Sorites paradox.¹⁰

It is in line with the orthodox position in the vagueness literature, then, for me to deny that there is a sharp cutoff n such that a lottery proposition must be at least as probable as n if one is to know it, while at the same time insisting that if a lottery proposition is sufficiently probable, one can come to know it; further, I can hold that while there are clear cases in which one does not know a lottery proposition because it is not sufficiently probable, and clear cases in which one does know a lottery proposition because it *is* sufficiently probable, there are also borderline cases in which one neither clearly knows a lottery proposition nor clearly fails to know a lottery proposition; in such a borderline case, the probability of the lottery proposition will fall somewhere between the probability of the lottery proposition in the clear case of ignorance and the probability of the lottery proposition in the clear case of knowledge.¹¹ Of course there are problems for such a view—among other things, it will face a version of the Sorites paradox. These problems, however, arise from the phenomenon of vagueness generally, and not from any peculiarity in my particular response to lottery arguments. Similar problems will afflict orthodox treatments of 'heap'; however, the unorthodox alternatives—claiming that a

precise number of grains is required to make a heap, or denying that there are any heaps—seem even worse. So it seems reasonable to adopt a view that is in keeping with the orthodox position on vagueness and leave it to researchers on vagueness to hash out the details.

Summing up, I am proposing a response to lottery arguments according to which one may know a lottery proposition provided that it is sufficiently probable; I deny, however, that there is a specific probability n such that one may know a lottery proposition just in case the lottery proposition is at least as probable as n . I am assuming that any problems resulting from this rejection of a sharp cutoff are problems due to vagueness, and that it is reasonable to proceed without having to address these vagueness-related issues.¹²

7. Rejecting the lottery conviction

The moderate invariantist response to lottery arguments considered in the previous section requires denying the lottery conviction. The lottery conviction, however, is very widely shared and is often viewed as a kind of fixed point that any adequate theory of knowledge must accommodate. In order for my response to lottery arguments to be plausible, then, I have to show that the lottery conviction can be defensibly rejected. To that end, I will consider what I take to be the three most important considerations that support the lottery conviction, and argue that none of them decisively show that the lottery conviction is true.

First, it might be argued that the lottery conviction is just so intuitively compelling that it cannot be denied. It does seem correct that particular *instances* of the lottery conviction are very intuitively compelling; for example, the intuition that one cannot know that a ticket in an ordinary lottery with a million tickets will lose is very powerful. It seems to me, however, that

intuition actually supports the *denial* of the lottery conviction, understood as the *general* claim that one can never know any lottery proposition, no matter how probable.

Suppose that Brown, a young man, likes to buy one ticket a week in the state lottery. Being a creature of habit, Brown will continue buying a ticket every week until he dies at a ripe old age. Can we know that Brown will not win the lottery's grand prize every single week for the rest of his life? It certainly seems to me that I can know this, even though there is a non-zero probability that Brown will win the grand prize every week from now on, and even though my only reason for thinking he will not is that it would be so incredibly improbable.¹³ Or, suppose that an immortal flips a fair coin every three seconds for the next one hundred trillion years—or the end of time, whichever comes first—can't I know that his coin will come up heads at least once? To put it mildly, intuition does not unequivocally support the lottery conviction.

The second consideration in support of the lottery conviction exploits connections between knowledge and practical reason. Consider again Hawthorne's example in which somebody offers to buy your lottery ticket (expected value: 50 cents) for a penny. Suppose you reason as follows:

The ticket is a loser.

So if I keep the ticket I will get nothing.

But if I sell the ticket I will get a penny.

So I'd better sell the ticket. (Hawthorne 2004, 29)

This reasoning is manifestly *unacceptable* if your only reason for believing that the ticket is a loser is that the probability of winning is so low. However, if you *know* that your ticket is a loser, then the reasoning should be acceptable. (After all, if you read the lottery results in a newspaper

and thereby came to know that your ticket is a loser, the reasoning would be acceptable.) So, the argument goes, you cannot know that your ticket is a loser on the basis of the odds alone.

The preceding argument clearly relies on the following premise:

(KP) *S* knows that *p* only if *S* may permissibly use *p* as a premise in practical reasoning.

The use of (KP), however, makes the argument dialectically ineffectual, since (KP) entails the falsity of moderate invariantism, as I shall now argue.

Moderate invariantists hold that if two subjects *A* and *B* differ with respect their practical interests, but are otherwise in the same epistemic position with respect to *p* (*p* is true in both of their circumstances, they have the same evidence for *p*, their beliefs were formed by the same equally reliable processes, etc.) then *A* knows that *p* just in case *B* knows that *p*. Furthermore, moderate invariantists are not skeptics. Since placing a certainty constraint on knowledge leads to skepticism, moderate invariantists hold that much of what we know is known in the absence of complete certainty. Let *p* be a proposition that *A* knows with less than complete certainty. Since (practical interests aside) *B*'s epistemic position with respect to *p* is the same as *A*'s, it follows from moderate invariantism that regardless of his practical situation, *B* knows that *p*. It then follows from (KP) that regardless of *B*'s practical situation, *B* may use *p* as a premise in practical reasoning. So it follows from moderate invariantism and (KP) that *B* may rationally bet on *p* no matter what the stakes—\$1 vs. the destruction of the entire world!—even though *p* is less than completely certain for *B*. Since that is obviously false, (KP) and moderate invariantism are inconsistent.¹⁴

Since (KP) entails the falsity of moderate invariantism, moderate invariantists need not accept an argument containing (KP) as a premise unless they have absolutely no alternative.

Though (KP) is initially quite plausible, I will try to show that moderate invariantists are not rationally obligated to abandon their position and endorse (KP): first, because (KP) is not supported by any really convincing arguments; second, because (KP) faces apparent counterexamples.

Probably the best argument for (KP) is inspired by Cohen's (1999, 59) observation that it seems improper to say things like 'I know that p , but still I need to check further'. If (KP) is false, then there are situations in which a subject S knows that p but may not use p as a premise in practical reasoning. In other words, S knows that p but he needs to check further. If the impropriety of this sentence indicates that it is false, then we have reason to believe that (KP) is true. Sosa has argued, however, that the impropriety of saying 'I know that p , but I need to check further' stems from the fact that saying 'I know that p ' conversationally implicates that one does not need to check further, and this implicature contradicts the claim that one does need to check further.¹⁵ If there is a pragmatic explanation for the impropriety of the sentence, then the argument for (KP) fails.

In response, Cohen follows Grice in claiming that conversational implicatures are *cancelable*: one can avoid the implicature p that would normally be conveyed by uttering a sentence s by saying something like ' s , but not p '. For example, if saying 'Jones is an above average soccer player' conversationally implicates that Jones is not a great soccer player, this implicature can be cancelled by saying 'Jones is an above average soccer player; in fact, he's a great soccer player'. Therefore, Cohen argues, it should *not* be improper to utter 'We know that p , but we need to check further' if 'We know that p ' merely implicates that we do not need to check further, since following up with 'but we need to check further' should cancel that

implicature. Since the utterance *is* improper, Cohen concludes that ‘We know that *p*’ *semantically* entails that we do not need to check further.

Consider, however, the impropriety of utterances like ‘I know that *p*, but I’m not completely certain’. We had better not say that ‘I know that *p*’ *semantically* entails that the speaker is completely certain that *p*, or skepticism follows on short notice. However, if saying ‘I know that *p*’ merely implicates that one is completely certain that *p*, then by Cohen’s reasoning following up with ‘I’m not sure’ should cancel the implicature, and the utterance should sound fine. It does not, however, and this suggests that there is something wrong with Cohen’s cancelability test.

There is much more that could be said, on either side, about this argument for (KP). The important thing to note here is that the issue of whether this argument succeeds is not settled. Thus it need not force the moderate invariantist to accept (KP), or the argument for the lottery conviction that uses (KP) as a premise.

The other reason why moderate invariantists need not accept (KP) is that there are at least putative counterexamples to the principle. In an influential paper, Radford (1966) argues through a series of examples that knowledge is consistent with a very low degree of psychological confidence; in particular, one can know that *p* without being at all sure of *p*. Here is Radford’s first example:

Man: Look, I *know* I locked the car. Still—I’ll go back and make absolutely sure.

Wife (irritated): Aren’t you sure?

Man: Well—*yes*, I *am* sure. I’d bet money on it. Still, I could be mistaken. It’s possible, isn’t it, darling? And this is a tough neighborhood.

Wife (surprised): Oh!

Man: Yes. And since it would be disastrous if I hadn’t locked it, I might as well go and check. I won’t be long.... (Radford 1966, 1)

Radford goes on to comment:

We may safely assume that, providing he has locked his car, the man knows that he has done so. (Radford 1966, 1)

If Radford is right that the man in the case knows that he locked his car, then this is a case in which a subject knows that p but needs to check further, and a counterexample to (KP). Now, it may not be as clear-cut as Radford wants to present it that the man in the case really does know that he locked his car. But it is certainly not clear that the man in the case does *not* know that he locked his car. Given that cases like Radford's *might* be counterexamples to (KP), it seems fair to say that it is unclear whether (KP) is true. So we seem to be under no obligation to accept (KP) as a premise in the argument for the lottery conviction.¹⁶

The third and final consideration in support of the lottery conviction that I will discuss involves a version of the lottery paradox involving knowledge rather than rational belief.¹⁷ Suppose you can know that your ticket l_1 is a loser on the basis of the odds alone. Then, if l_2 – l_n are all the other losing tickets, it seems that you can also know that l_2 is a loser on the same grounds, and know that l_3 is a loser, ..., and know that l_n is a loser. This might seem bad enough, but if multi-premise closure (the principle that if you know p_1 and you know p_2 and ... you know p_n , then you can deduce and thereby come to know the conjunction p_1 and p_2 and ... p_n) is correct, then it seems that you can also deduce and thereby come to know that l_1 and l_2 and ... l_n are *all* losers. This is an extremely unpalatable result, since in typical lotteries the proposition that l_1 and l_2 and ... l_n are all losers is highly *improbable*, and we should not be able to come to know truths if they are highly improbable.

The natural response to this version of the lottery paradox would be to reject the premise that you can know that your ticket is a loser on the basis of the odds alone, but of course this would be to accept the lottery conviction. On the other hand, if this premise is not rejected, it is unclear whether the paradox's unacceptable consequences can be avoided.

The consequence that one can know that l_1 – l_n are all losers can be avoided by denying multi-premise closure. This might seem like a desperate move, but in fact there are good reasons to deny multi-premise closure even if one *endorses* the lottery conviction.¹⁸ Here are some plausible claims: (i) skepticism is false—we know many things; (ii) if skepticism is false, then we know many propositions without complete certainty; (iii) propositions can be more or less certain; (iv) there is some threshold of certainty below which a proposition cannot be known because it is too doubtful; and (v) if two logically independent propositions are each known with less than complete certainty, then their conjunction is less certain than either proposition. (i)–(v) are all plausible, but together they entail the falsity of multi-premise closure, since by repeatedly conjoining known but less than completely certain logically independent propositions, one will eventually end up with a big conjunction that is too doubtful to be known. Since (i)–(v) are plausible whether or not the lottery conviction is true, they seem to provide a principled way of avoiding the lottery paradox’s second unpleasant consequence.

But what of the lottery paradox’s first unpleasant consequence, that one can know, of each of the losing tickets, that it is a loser? Instead of trying to evade this consequence, I suggest that we accept it, while trying to explain why the consequence *seems* so abhorrent. One possible explanation is that this response reflects a misguided commitment to multi-premise closure. When we try to imagine a subject who knows that l_1 is a loser, and knows that l_2 is a loser, ..., and knows that l_n is a loser, we reflexively apply multi-premise closure and infer that the subject can know that l_1 – l_n are all losers. Since this is clearly false, we judge that at least one of the premises from which we inferred the false conclusion is itself false. But we fail to consider that the fault might instead lie with our use of multi-premise closure.

Here is another possible explanation for why the first unpleasant consequence seems so abhorrent. When we imagine a subject knowing, of each of the losing tickets, that it is a loser, we imagine the subject having all these propositions in mind at the same time. A thoughtful subject who was considering these propositions all at once would recognize that it is very unlikely that they are all true. However, the subject has no reason to think that one of these proposition is more likely to be false than any of the others. Now, it may well be that something like the following epistemic principle is correct:

If S consciously considers propositions p_1, p_2, \dots, p_n , and recognizes that it is improbable that p_1-p_n are all true, but has no reason to think that one of p_1-p_n is more likely to be false than any of the others, then S does not know any of p_1-p_n .

When we imagine a subject knowing, of every losing ticket, that it is a loser, we imagine her having all these propositions in mind at the same time and recognizing that it is highly improbable that they are all true. If the epistemic principle above is about right, the subject cannot know any of those propositions in that case. Perhaps the abhorrence we feel towards the lottery paradox's first unpleasant consequence results from acknowledging this epistemic principle while imagining a fairly specific scenario in which the subject considers all the propositions at once and recognizes that they are unlikely to all be true. However, the epistemic principle above does not preclude a subject from knowing, of a single lottery ticket, that it is a loser; furthermore, it is consistent with a subject coming to know, of each losing ticket, that it is a loser, provided that he acquires each bit of knowledge discretely, without ever reflecting on those propositions all at once.¹⁹

I have examined three important considerations in support of the lottery conviction, and found that none are conclusive. It appears that the prospects for denying the lottery conviction are not quite so grim as one might have thought.

Not everyone will be convinced by the preceding responses, and some will continue to feel, quite strongly perhaps, that the lottery conviction must be correct. Even if one persists in feeling that the lottery conviction must be correct, one should be cautious about rejecting moderate invariantism on the grounds that it cannot provide an adequate answer to lottery arguments without denying the lottery conviction. Consistency requires that if one rejects moderate invariantism on those grounds, then on the same grounds one must also reject EC and SSI.

We saw in section 5 that Hawthorne's SSI based treatments of the lottery requires rejecting the lottery conviction. With respect to EC, the situation is slightly more complex. As we saw in section 4, the EC-based response to lottery arguments defended by Lewis only rejects one version of the lottery conviction, that involving *m*-knowledge. It seems to me, however, that anybody who was firmly convinced that one cannot know a lottery proposition would be just as firmly opposed to rejecting the version of the lottery conviction involving *m*-knowledge. If a person strongly rejects the claim that one can know that a lottery proposition is true, it will not assuage her one bit if it is made clear to her that it is only in the ordinary, everyday sense of 'know' that lottery propositions can be known. But suppose I am wrong about that. If Lewis's denial of the version of the lottery conviction involving *m*-knowledge *would* be acceptable to the critic, then my moderate invariantist reply to lottery arguments also ought to be acceptable to the critic. After all, like Lewis I hold that one can know lottery propositions in the ordinary sense of

‘know’; I simply believe, *pace* Lewis, that the ordinary sense of ‘know’ is the only sense of ‘know’ there is.

Moderate invariantism, SSI, and EC are all in the same boat: if one is going to reject moderate invariantism on the grounds that it cannot answer lottery arguments without denying the lottery conviction, one will also have to reject EC and SSI. So much the worse, it might be replied, for SSI and EC. But moderate invariantism, EC, and SSI exhaust the views that are committed to both anti-skepticism and some version of epistemic closure. So if one rejects all these views, one must reject either anti-skepticism or epistemic closure. It may well be, however, that one’s reasons for endorsing anti-skepticism and epistemic closure are even more powerful than one’s reasons for endorsing the lottery conviction. In deciding which of these views to retain and which to reject, one needs to balance all of these competing considerations; even if one is reluctant to deny the lottery conviction, it may be ultimately be the least unacceptable of one’s options.

¹ Kyburg’s original lottery paradox concerned rational belief, not knowledge.

² DeRose has not offered a contextualist response to lottery arguments in print. He does discuss lotteries in DeRose (1996), but there his goal there is to explain why we tend to judge that one cannot know that a lottery ticket is a loser on the basis of the odds, but that one can know that a lottery ticket is a loser by reading the results in a newspaper, even though there is some chance that the paper contains a misprint. One could bring together the claims in this article and DeRose’s (1995) response to AI to construct a parallel response to lottery arguments on DeRose’s behalf; the central idea would be that when one utters ‘I don’t know that my ticket will lose’, the sphere of worlds relevant to determining the truth of this utterance expands, if need be, to include the closest possible worlds in which the lottery ticket is a winner. In those worlds, one continues to believe that it is a loser; consequently, it follows from the rule of sensitivity that ‘I don’t know that my ticket is a loser’ is true. Since the expanded sphere of worlds remains in place throughout the argument (due to the “stubbornness” of high standards), the lottery argument is sound. Maintaining the parallel with DeRose’s response to AI, DeRose might claim that one *can* truly utter ‘I know I will never be rich’ in a different context, where nobody has made knowledge claims involving lotteries and the sphere of relevant worlds has not expanded as above, because the closest possible world in which I will be rich is one in which I don’t believe that I will never be rich. The problem for this response, however, is that it is plausible that the closest possible world in which I will be rich is precisely a world in which I win the lottery, and consequently don’t believe (until I win) that I will ever be rich. So it appears that a response that is strictly parallel to DeRose’s response to the AI would not adequately answer this lottery argument. Of course, DeRose could offer a different kind of response, or (my choice) simply accept that one cannot truly claim to know that one will never be rich if one possesses a lottery ticket.

³ DeRose (1996) criticizes this solution because it relies on supposing that there is a winning ticket; as DeRose points out, plenty of actual lotteries don’t have a winner; however, recognizing that a lottery did not have a winner in no way inclines us to judge that one could have known before the drawing that one’s ticket would lose.

⁴ It might be thought that the claim that a sentence can be true in a context even though it cannot be truly uttered in that context is incoherent; consider, however, the sentence ‘I am not speaking’. This sentence is true relative to my context (in which I am not speaking), but it would be false if I were to utter it.

⁵ Hawthorne (2004) calls SSI ‘sensitive moderate invariantism’.

⁶ Later, Hawthorne (2004, 169) cashes out salience in a different way: the possibility that p is salient to S if and only if it appears to S that, for all she knows, p .

⁷ This is in fact an accurate description of my state of mind.

⁸ One way to soften the blow here would be for Hawthorne to continue to deny that one can come to know a lottery proposition q via *deduction* from a known proposition p that entails q , on the grounds that knowledge that p requires antecedent knowledge of q . This move, however, does nothing to avoid the conclusion that one can know (somehow or another) a lottery proposition.

⁹ I say that lottery arguments involving lottery propositions whose probability is greater than or equal to n are *typically* unsound because it is still possible for one to fail to know the lottery proposition for reasons having nothing to do with the lottery proposition’s probability; for example, one might not know the proposition because one does not believe it.

¹⁰ An exception is Williamson (1994, 185-215), who argues that vague terms have precise applicability conditions, but that we cannot know what those applicability conditions are.

¹¹ It is also in keeping with the orthodox position on vagueness to hold that there is no sharp cutoff between the borderline cases of knowledge and the clear cases of knowledge: in other words there are borderline cases of borderline cases, or higher-order vagueness.

¹² In discussing the sharp cutoff view, I noted that holding that there is a probability n such that one can know a lottery proposition just in case it is at least as probable as n is compatible with a huge amount of error and ambivalence when it comes to judging lottery arguments to be sound or unsound, since we may be mistaken or unsure both about the value of n , and about the actual probability of the lottery proposition we are considering. If, as I am now suggesting, there is no sharp cutoff, and in some cases it is vague or indeterminate whether a certain lottery proposition is sufficiently probable to be known, it seems that this will only have a tendency to make our assessments of the soundness of lottery arguments even more fallible and tentative.

¹³ Hawthorne (2004, 20) discusses a similar case. He claims that the intuition that one can know that one will not win the lottery every year for the next thirty years can be disrupted by considering one big lottery where the odds of winning are the same as the odds of winning the lottery thirty years in a row. Presumably the idea is that since we recognize that the odds are the same in both cases, and we judge that one cannot know that one will lose the big lottery, one comes around to the view that one cannot know that one will not win the thirty in a row. However, this comparison has the opposite effect on my intuitions: I become more convinced that I can know that I will lose the big lottery.

¹⁴ Cohen (2004, 487) makes a similar point.

¹⁵ This argument of Sosa’s was apparently made at a conference and is reported by Cohen (1999, 60).

¹⁶ At the 2004 Pacific APA, Baron Reed offered a case along the following lines: suppose a participant in an extreme game show is asked a question that she clearly knows the answer to—say, the year in which Pearl Harbor was bombed. If she gets the answer wrong, she receives one hundred lashes. Suppose further that she has a free, one-time opportunity to consult an authoritative history book before answering. Given the stakes, it seems rational for the contestant to check the history book. But it is quite plausible that she retains the knowledge that Pearl Harbor was bombed in 1941 even as she is checking the book. Again, this is not an utterly clear-cut counterexample to (KP), but it is certainly enough to raise some doubt as to (KP)’s truth.

¹⁷ Knowledge versions of the lottery paradox are discussed by Bonjour (1985, 53-56) and Nelkin (2000).

¹⁸ Hawthorne (2004, 181-185) is quite pessimistic about the prospects for multi-premise closure, though he wants to retain it if possible.

¹⁹ One might object that even ordinary folk find it implausible that one could know, of each of the losing tickets, that it is a loser, but that the two suggested explanations for this reaction require a degree of philosophical sophistication that ordinary folk do not possess. I agree that ordinary folk would not explicitly reason as described in these explanations. It seems more plausible, though, that they have an implicit grasp of the epistemic principles involved in the reasoning, and that they can engage in a non-conscious reasoning-like process, using the implicitly grasped epistemic principles, that culminates in the same intuitive judgment.

CHAPTER 4

THE AIRPORT CASE (I): INTUITION-DENYING REPLIES

1. Introductory remarks

In addition to the contextualist replies to skeptical arguments considered in chapters 2 and 3, EC is supported by so-called *context-shifting arguments*. Generally speaking, a context-shifting argument begins with a case (or cases) in which utterances containing a certain expression e occur in distinct contexts. On the basis of our intuitions about these utterances, it is then argued that e is semantically context-sensitive. A context-shifting argument in support of EC, then, would have to begin with utterances containing ‘know’ occurring in distinct contexts, and, on the basis of our intuitions about those utterances, argue that ‘know’ is semantically context-sensitive.

The context-shifting argument for EC that I will focus on in this chapter begins with Cohen’s famous Airport Case:¹

Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, “Yes I know—it does stop in Chicago.” It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, “How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute.” Mary and John agree that Smith doesn’t really *know* that the plane will stop in Chicago. They decide to check with the airline agent. (Cohen 1999, 57)

To keep things clear in what follows, I will slightly alter and flesh out the case. Suppose, then, that Smith utters (S1) in context c_1 and at time t_1 , and that Mary utters (S2) in context c_2 and at time t_2 :

(S1) I know that the plane stops in Chicago.

(S2) Smith does not know that the plane stops in Chicago.

Finally, for concision let ‘*STOPS*’ denote the proposition that the plane stops in Chicago.

The Airport Case typically elicits both the intuition that Smith is saying something true by uttering (S1), and the intuition that Mary is saying something true by uttering (S2). But how, exactly, are these intuitions supposed to lead to the conclusion that EC is true? There are (at least) two important obstacles to using the intuitions to argue for EC.

First, intuitions can be misleading. If it could be made plausible that the intuitions elicited by the Airport Case are unreliable or suspect, then they would clearly be of little use in arguing for EC. Those involved in the debate over EC are very aware of this issue, and most of the invariantist responses to the Airport Case and other real-world cases have been devoted to explaining why at least one of the intuitions elicited by the case is mistaken. Call such replies *intuition-denying replies*.

The second obstacle to using these intuitions to argue for contextualism is rarely discussed. On the face of it, the intuitions elicited by the Airport Case concern what Smith and Mary are claiming to be true; that is, they are about the truth-values of the propositions *asserted* by Smith and Mary, where assertion is understood as the speech-act by which one makes a claim, states what is the case, or says that something is true.² EC, however, is a semantic thesis, about the context-sensitivity of ‘know’ and sentences in which it occurs. The contextualist must therefore establish a link between the propositions *asserted* by Smith and Mary and the propositions semantically expressed by their utterances in order to use the intuitions elicited by the Airport Case to argue for EC.

Contextualists rarely dwell on the relation between asserting that *p* and uttering a sentence that semantically expresses *p*; they often appear to assume that they amount to the same

thing, at least in ordinary contexts in which one is speaking literally. Still, the case of metaphor reminds us that asserting a proposition and semantically expressing it are not identical: when an overworked clerk says, ‘I’m going back to the salt mines’ after his lunch break, he is not actually claiming to be going to a facility that extracts sodium chloride from the earth. So it is possible to utter a sentence semantically expressing p without asserting that p . In addition, we shall see in chapter 5 that Soames (2002) and Cappellen and Lepore (2005) have argued, quite plausibly, that a sentence can be used to assert propositions that are not the sentence’s semantic content.

Since it is possible to utter a sentence with semantic content p without asserting that p , and (if the arguments I just mentioned are successful) it is possible to assert that p by uttering a sentence whose semantic content is not p , the notion of the proposition asserted by a sentence utterance and the notion of the proposition semantically expressed by a sentence utterance are two distinct notions. A context-shifting argument for EC based on the Airport Case, which begins with intuitions about what Smith and Mary assert, but ends with a conclusion about the semantic contents of their utterances, must therefore explicitly state and defend a premise linking assertion with semantic content.

In light of these points, here is one stab at a context-shifting argument for EC based on the intuitions elicited by the Airport Case:

The Airport Argument

- (P1) The Airport Case elicits the intuition that both Smith and Mary are asserting truths in the Airport Case.
- (P2) The intuitions elicited by the case are correct.
- (C1) Therefore, both Smith and Mary are asserting truths in the Airport Case.

- (P3) The proposition asserted by Smith is the proposition semantically expressed by (S1) in c_1 , and the proposition asserted by Mary is the proposition semantically expressed by (S2) in c_2 .
- (C2) Therefore, the propositions semantically expressed by (S1) in c_1 and (S2) in c_2 are both true.
- (P4) Knowledge attributions and denials claim or deny that some 2-place epistemic relation obtains between a person and a proposition at the time of utterance; that is, the proposition semantically expressed by (S1) in c_1 can be represented as $\langle \text{At } t_1, \text{Smith } K_1 \text{ STOPS} \rangle$, and the proposition semantically expressed by (S2) in c_2 can be represented as $\langle \text{At } t_2, \sim \text{Smith } K_2 \text{ STOPS} \rangle$.
- (P5) If Smith's epistemic position with respect to *STOPS* is identical at t_1 and t_2 , then for any epistemic relation K , Smith bears K to *STOPS* at t_1 if and only if Smith bears K to *STOPS* at t_2 .
- (P6) In the Airport Case, Smith's epistemic position with respect to *STOPS* is identical at t_1 and t_2 .
- (C3) So, $K_1 \neq K_2$; that is, the epistemic relation semantically expressed by 'know' varies with the context of utterance.

In the Airport Argument, the premise linking assertion with semantic content, which I argued any context-shifting argument for EC based on the Airport Case must provide, is (P3). In chapter 5, I will pursue a response to the Airport Argument that involves questioning (P3). As I have mentioned, however, the vast majority of invariantist replies to the Airport Case take (P3) for granted, and instead offer intuition-denying replies, which deny (P2). In this chapter, I will try to make the case that the prospects for an effective intuition-denying reply to the Airport Case

are very poor. My plan is to canvas every distinctive invariantist response to the Airport Case (or similar cases) that I know of. For each of these accounts, I will argue that it either fails to adequately respond to the Airport Case, or, if it does show some promise to adequately respond to the case, this is because it denies (P3), not (P2).

In section 2, I discuss *warranted assertability maneuvers* (WAMs), as characterized by DeRose (2002), who coined the term. In section 3, I discuss invariantist replies that appeal to propositions conversationally implicated by the utterances in the Airport Case; these *implicature based replies* purport to be WAMs, but may not neatly square with DeRose's original characterization of a WAM, and therefore must be treated separately. In sections 4-6, I discuss other replies to the Airport Case, by Bach (2005), Williamson (2005), and Stanley (2005), respectively.³

2. Warranted assertability maneuvers

Invariantists who adopt an intuition-denying reply to the Airport Case have an initial choice to make: whether to deny the intuition that Smith asserts a truth by uttering (S1), or to deny the intuition that Mary asserts a truth by uttering (S2). This decision will be dictated by how demanding the invariantist takes the standards for knowledge to be. Those who hold that the standards for knowledge are uniformly high will deny that Smith is asserting a truth by uttering (S1). Call this the 'high-standards response', and those who offer it 'high-standards invariantists'. Those who hold that the standards for knowledge are relatively low across-the-board will deny that Mary is asserting a truth by uttering (S2). Call this the 'low-standards response', and those who offer it 'low-standards invariantists'.⁴ It would be profoundly unsatisfying, however, for the invariantist to simply deny one of the intuitions and leave it at

that. The invariantist owes us some explanation for how we came to have the allegedly misleading intuition. The purpose of a WAM is to provide such an explanation.

DeRose, who first introduced WAMs, characterizes them as follows:

A WAM involves explaining why an assertion can seem false (or at least not true) in certain circumstances in which it is in fact true by appeal to the fact that the utterance would be improper or unwarranted in the circumstances in question. Going the other way, an intuition that an assertion is true can be explained away by means of the claim that the assertion, while false, is warranted, and we mistake this warranted assertability for truth. Either way, the maneuver is based on the correct insight that truth and warranted assertability are quite different things, but that we can easily mistake one for the other. (DeRose 2002, 171-172)

As the phrase “improper or unwarranted” makes clear, what DeRose means by ‘warranted assertability’ is nothing more than conversational propriety. A WAM, then, is a device for explaining away theoretically inconvenient intuitions about an assertion’s truth-value, by claiming that we confuse conversational impropriety with falsity, or conversational propriety with truth. An intuition that an assertion is false can be explained away by claiming that the assertion is conversationally improper, and that we mistakenly judge it to be false because we confuse conversational impropriety with falsity. An intuition that an assertion is true can be explained away by claiming that the assertion is conversationally proper, and that we mistakenly judge it to be true because we confuse conversational propriety with truth.

It is important to note that any invariantist WAM against the Airport Case, low-standards or high-standards, must explain away an intuition that an assertion is *true*. Low-standards invariantists hold that Mary asserts a falsehood by uttering (S2); this assertion, however, is intuitively true. A low-standards WAM must therefore hold that Mary’s assertion is false but conversationally proper, and that we confuse this conversational propriety with truth. Similarly, a high-standards WAM must explain away the intuition that Smith asserts a truth by uttering (S1);

such a WAM would hold that Smith's assertion is false but conversationally proper, and that we confuse this conversational propriety with truth (DeRose 2002, 191-194).

I will argue that no invariantist WAM against the Airport Case can succeed. My argument makes use of the *knowledge account of assertion*, according to which assertion is governed by the knowledge rule:

(KR) Assert that p only if you know that p .

In what follows, I will have very little to say about what it takes for assertion to be governed by a rule. My argument that an invariantist WAM cannot succeed against the Airport Case relies on only one assumption, which I take to be uncontroversial: that if assertion is governed by a rule, then any assertion that violates the rule is conversationally improper.

The version of the knowledge account presented here is significantly weaker than other accounts relating knowledge and assertion that have been articulated in the literature. Williamson (2000, 238-241), for example, defends the position that KR is the *constitutive* rule of assertion. By this he means that assertion's being governed by KR makes it the speech act that it is, just as chess's being governed by its rules makes it the game that it is. A corollary of Williamson's view is that no speech act that is not a kind of assertion is governed by the knowledge rule. As I've presented it, the knowledge account does not assume that KR is a constitutive rule of assertion, or that assertion is the only speech act type that is governed by KR.

DeRose (2002, 180) also relates assertion and knowledge in a stronger way than I wish to here. He says that KR is the *only* rule governing assertion having to do with how strong one's epistemic position must be in order to assert that p .⁵ My formulation of the knowledge account leaves open the possibility that in addition to KR, other rules having to do with the strength of one's epistemic position govern assertion.

The version of the knowledge account presented here claims only that assertion is governed by KR. This entails only one substantive commitment: that a speaker may not properly assert that p unless she knows that p . This is consistent with the stronger accounts discussed above, without bringing on board their (perhaps questionable) additional commitments.

The knowledge account is well supported by diverse linguistic evidence. A brief summary follows.

First, the knowledge account explains the conversational impropriety of Moore-paradoxical sentences like (S3):

(S3) It is snowing, but I don't know that it is snowing.⁶

It follows from the knowledge account that asserting (S3) cannot be conversationally proper.⁷

Williamson (2000, 253-254) argues along the following lines. Suppose, for *reductio*, that the knowledge account is true and that asserting (S3) is conversationally proper. Then (since knowledge of a conjunction implies knowledge of the conjuncts) the speaker must know that it is snowing and know that she doesn't know that it is snowing. However, if the speaker knows that she doesn't know that it is snowing, then (since knowledge is factive) the speaker doesn't know that it is snowing. So the speaker both knows and doesn't know that it is snowing. Contradiction. So the knowledge account explains the conversational impropriety of asserting (S3), since if the knowledge account is true then it is impossible to properly assert (S3).

Second, the knowledge account explains the impropriety of certain assertions involving lotteries (Williamson 2000, 253-254). For example, it seems improper for a speaker to say (S4) to her friend if the speaker's only evidence that her friend's lottery ticket did not win is that the odds of winning are extremely low:

(S4) Your ticket did not win.

The impropriety of asserting (S4) in the envisaged situation is explained by the knowledge account as follows. One cannot know that a lottery ticket is a loser on the basis of purely statistical evidence. Therefore, in the envisaged situation the speaker asserts that her friend's ticket did not win without knowing that her friend's ticket did not win—and this violates KR. Since it is conversationally improper to violate a rule governing assertion, it follows from the knowledge account that asserting (S4) is conversationally improper in the envisaged situation.

Third, it is commonplace to challenge assertions by saying 'You don't know that' or 'How do you know that?' If assertion is not governed by KR, it would seem that such challenges are inappropriate, and that one could appropriately respond with 'I didn't *say* that I knew it' or 'So what?' Since the challenges seem appropriate and the responses absurd, however, it must be conversationally improper to assert what one doesn't know—as the knowledge account predicts (Unger 1975, 263-264; Williamson 2000, 252-253).

If the knowledge account is correct, then no invariantist WAM against the Airport Case can succeed. To see this, consider the example of a low-standards invariantist WAM. Such a WAM must explain away the intuition that Mary's knowledge denial is true, by holding that the denial is conversationally proper but false and that we confuse the assertion's conversational propriety with its literal truth. It follows from the knowledge account, however, that Mary's denial is conversationally proper only if it is true. Here's why: If assertion is governed by KR, then Mary's knowledge denial is conversationally proper only if Mary knows that Smith doesn't know *STOPS*; but Mary knows that Smith doesn't know *STOPS* only if it is true that Smith doesn't know *STOPS*. In other words, it follows from the knowledge account and the factivity of knowledge that Mary's knowledge denial is conversationally proper only if it is true. Since the

low-standards WAM claims that Mary's knowledge denial is conversationally proper but false, the knowledge account is inconsistent with a low-standards WAM.

More generally, any WAM that denies an intuition that an assertion is true is inconsistent with the knowledge account. Such a WAM must hold that the assertion in question is conversationally proper but false; it follows from the knowledge account, however, that an assertion is conversationally proper only if it is true. Since any invariantist WAM (low or high-standards) against the Airport Case must deny an intuition that an assertion is true, no invariantist WAM against the Airport Case is consistent with the knowledge account.

As we have seen, the knowledge account is well supported by diverse linguistic evidence, and the knowledge account is inconsistent with an invariantist WAM against the Airport Case. These facts constitute a *prima facie* plausible argument that no invariantist WAM against the Airport Case can succeed. The knowledge account, however, is not wholly uncontroversial; philosophers have recently begun to argue that the evidence that motivates the knowledge account can be accommodated by supposing that different, weaker, norms govern assertion.

For example, Weiner (2005) argues that the knowledge account is false, and that assertion is governed by the truth rule:

(TR) Assert that p only if p is true.

Weiner argues that the evidence that motivates the knowledge account can be explained by appealing exclusively to TR and general conversational principles of the sort defended by Grice (1989).

Lackey (2007), on the other hand, argues that neither KR nor TR governs assertion. Instead, she holds that assertion is governed by the reasonable to believe rule, roughly:

(RTBR) Assert that p only if it is reasonable for you to believe that p .⁸

Like Weiner, Lackey argues that the evidence that motivates the knowledge account can be accounted for by RTBR and general conversational principles.

While Weiner's and Lackey's arguments are interesting and certainly worthy of close scrutiny, I will not discuss them here; instead, I will argue that an invariantist WAM against the Airport Case cannot succeed even if assertion is governed by TR or RTBR.

If assertion is governed by TR, then an assertion is conversationally proper only if it is true. As we have seen, however, an invariantist WAM against the Airport Case must hold that there are false but conversationally proper assertions. It follows that no invariantist WAM against the Airport Case can succeed if assertion is governed by TR.

When it comes to RTBR, the argument is more involved. If assertion is governed by RTBR, then asserting that p is conversationally proper only if it is reasonable for the speaker to believe that p . An invariantist WAM against the Airport Case holds that some assertion is false but conversationally proper; therefore, if assertion is governed by RTBR then a WAM entails that it is reasonable for the speaker to believe the content of this false assertion. I will argue, however, that it would not be reasonable for the speaker to believe the content of this false assertion.

Consider first the case of a low-standards WAM against the Airport Case. Such a WAM must hold that Mary's knowledge denial is conversationally proper; hence, if RTBR is a norm governing assertion, it must be reasonable for Mary to believe that Smith doesn't know *STOPS*. A natural thought is that reasonably believing that p requires having good evidence for p ; let's adopt this as a provisional assumption, to be examined more closely later. Given this assumption, if assertion is governed by RTBR and a low-standards WAM is to succeed, Mary must have good evidence that Smith doesn't know *STOPS*. The low-standards invariantist, however, holds

that Smith really does know *STOPS*. According to the low-standards invariantist, consulting an airline schedule—provided that it is in fact accurate and reliable—is sufficient to confer knowledge of *STOPS* on Smith. The fact that Smith consulted the airline schedule is therefore excellent evidence that Smith knows *STOPS*. Furthermore, Mary possesses this evidence: she knows full well that Smith has consulted the airline schedule. Since Smith’s evidence for *STOPS* is good evidence for Mary that Smith knows *STOPS*, if Mary is to have good evidence that Smith *doesn’t* know *STOPS*, it must be of an exceptionally powerful kind, sufficient to override (or undermine) the good evidence Mary has that Smith does know *STOPS*, and make it all-things-considered reasonable for Mary to believe that Smith doesn’t know *STOPS*. Such evidence must be such that the low-standards invariantist would take it to be excellent evidence for Mary that Smith doesn’t know *STOPS*. For example, if Mary had excellent evidence that *STOPS* is false—say having overheard the pilot say that the plane flies directly to New York—or excellent evidence that Smith’s evidence for *STOPS* is misleading—say being told by an airline employee that the schedule is full of typos—then the low-standards invariantist could plausibly claim that Mary has good evidence that Smith doesn’t know *STOPS*. It is clear from the case that Mary has no evidence of this sort, however. It is Mary’s practical concerns and skeptical worries that lead her to deny that Smith knows *STOPS*; the low-standards invariantist, however, denies that such factors can be relevant to whether Smith knows *STOPS*. The nature of the low-standards invariantist position, then, precludes Mary from reasonably believing that Smith doesn’t know *STOPS* in the Airport Case. It follows from this and the claim that assertion is governed by RTBR that Mary’s knowledge denial is conversationally improper; hence, a low-standards WAM against the Airport Case is inconsistent with the claim that assertion is governed by RTBR.

The case of a high-standards invariantist WAM against the Airport Case is analogous. The high-standards invariantist holds that the standards for knowledge are demanding enough that simply consulting a printed schedule is not enough to confer knowledge of *STOPS* on Smith. The high-standards invariantist therefore holds, contrary to intuition, that Smith's knowledge claim is false. A high-standards WAM attempts to explain away the powerful intuition that Smith's claim is true by holding that Smith's claim is conversationally proper, and that we confuse this conversational propriety with truth. If RTBR governs assertion, then for a high-standards WAM against the Airport Case to succeed, it must be reasonable for Smith to believe that he knows *STOPS*. According to high-standards invariantism, however, merely consulting an airline schedule is not sufficient to confer knowledge of *STOPS* on Smith. Therefore, according to high-standards invariantism Smith's evidence for *STOPS*—having consulted the schedule—is not good evidence that he knows *STOPS*. Smith has no evidence that the high-standards invariantist would hold constitutes good evidence that Smith knows *STOPS*. On the other hand, he has excellent evidence—still according to the high-standards invariantist—that he doesn't know *STOPS*. The high-standards invariantist holds that Smith must *at least* rule out the possibility that the schedule contains a misprint before being in a position to know *STOPS*. The fact that Smith has not ruled out this possibility is excellent evidence (by the high-standards invariantist's lights) that he doesn't know *STOPS*. Furthermore, Smith possesses this evidence: if he knows anything, he knows that he hasn't ruled out the possibility that the schedule contains a misprint. Since Smith has no good evidence that he knows *STOPS*, and excellent evidence that he doesn't know *STOPS*, it is not reasonable for him to believe that he knows *STOPS*. If assertion is governed by RTBR, it follows that Smith's knowledge claim is conversationally

improper. So a skeptical invariantist WAM against the Airport Case is inconsistent with the claim that assertion is governed by RTBR.

The preceding arguments rely on the assumption that reasonably believing p requires having good evidence for p . But this assumption does not seem quite right. It is plausible that one can reasonably believe p even if one does not have good evidence for p , provided that one reasonably believes a false theory and one's evidence for p *would* be good evidence for p if the false theory were true. So consider again the example of a low-standards WAM. According to this WAM, Mary's knowledge denial is false but conversationally proper; if assertion is governed by RTBR, it follows that it is reasonable for Mary to believe that Smith doesn't know *STOPS*. I argued that by the low-standards invariantist's lights it is *not* reasonable for Mary to believe that Smith doesn't know *STOPS*, because she has excellent evidence that Smith knows *STOPS*, and no evidence that undermines or overrides this evidence that Smith knows *STOPS*. Suppose, however, that Mary reasonably but falsely believes that the standards for knowledge are high enough that Smith couldn't know *STOPS* merely by checking an airline schedule. Then the fact that Mary knows that Smith hasn't done more than check the airline schedule gives her what is, according to this false theory, excellent evidence that Smith doesn't know *STOPS*. Plausibly, then, it is reasonable for Mary to believe that Smith doesn't know *STOPS*, and proper for her to deny that Smith knows *STOPS*, even if assertion is governed by RTBR.

So far so good, but what is the low-standards invariantist who adopts this line supposed to say about *Smith's* claim to know *STOPS* in the Airport Case? Smith's claim is conversationally proper. If RTBR governs assertion, then it must be reasonable for Smith to believe that he knows *STOPS*. If, however, Smith were like Mary and believed that the standards for knowledge are high, then it would not be reasonable for him to believe that he knows *STOPS*.

So the low-standards invariantist who adopts this line must claim that unlike Mary, Smith does not believe that the standards for knowledge are high; either he believes that the standards for knowledge are relatively low, or (perhaps more plausibly) he simply doesn't have any particular thoughts about what knowledge requires. Either way, the suggestion under consideration explains the fact that both Mary's and Smith's assertions in the Airport Case are conversationally proper by claiming that they think about knowledge differently; Mary has the reasonable but mistaken belief that the standards for knowledge are high, and Smith doesn't.

This suggestion on behalf of the low-standards invariantist is open to the charge that it illegitimately alters the Airport Case in order to avoid an unwelcome result. When we read the unadulterated Airport Case, we are led to believe that Smith and Mary are just ordinary people who happen to be in quite different practical situations. In order to save a low-standards invariantist WAM against the Airport Case (on the supposition that RTBR governs assertion), we must now suppose that Mary, but not Smith, mistakenly believes that the standards for knowledge are high. If the low-standards invariantist WAM can be saved only by introducing this additional feature, then it will fall prey to the Airport Case supplemented by an explicit stipulation: that Mary and Smith, if they think about knowledge at all, both think about it in the same way.

It might be replied that even if the case starts out with characters who (when they think about it at all) think about knowledge in the same way, they will end up thinking about knowledge differently by the time they make their crucial assertions, *because* they are in different situations. The thought is this. Before they get to the airport, Smith and Mary think about knowledge in the same way; since they are ordinary people with better things to do, they don't spend a lot of time reflecting on what knowledge requires. However, whereas Smith

continues not to think about the requirements of knowledge, Mary's practical situation causes her to reflect on various possibilities in which *STOPS* is false in spite of Smith's evidence for *STOPS*. These reflections then lead Mary to adopt a false belief about the requirements of knowledge.

I leave it open whether this reply is psychologically plausible. Even if it is, it will not help save a low-standards WAM if RTBR governs assertion. The trouble is that for the WAM to succeed, Mary must not only adopt a false belief about the requirements for knowledge—it must be *reasonable* for her to do so. But if reflecting on various possibilities makes it reasonable for Mary to believe that the standards for knowledge are high, then it should also make it reasonable for *Smith* to believe that the standards for knowledge are high; Smith, after all, is in just as good a position as Mary to think about the possibilities. Mary and Smith not only have the same evidence that Smith knows, or doesn't know, *STOPS*; they also have the same evidence for or against believing that the standards for knowledge are high. If it is reasonable for Smith to believe that the standards for knowledge are high, however, then it will not be reasonable for him to believe that he knows *STOPS*.

Might one object that what it is reasonable to believe can depend on one's practical situation, and thus, since Smith and Mary are in different practical situations, that it could be reasonable for Smith to believe that the standards for knowledge are low, but reasonable for Mary to believe that the standards for knowledge are high? Note that this move is available only if those who would be doing the objecting can incorporate the view that what it is reasonable to believe can depend on one's practical situation into their overall theory of knowledge. Since this response would be made in defense of a low-standards WAM against the Airport Case, those who might offer it are either moderate invariantists who hold that the standards for knowledge

are uniformly low, or defenders of SSI who hold that the standards for knowledge vary depending on one's practical situation.

Moderate invariantism and the position that what it is reasonable to believe depends on one's practical situation are, if not outright incompatible, at least in great tension with one another. Recall that moderate invariantism denies intellectualism, the view that whether *S* knows that *p* depends in part on *S*'s practical interests. So according to moderate invariantism, the epistemic standards required to know that *p* will not vary with one's practical interests. If moderate invariantism is true but the epistemic standards required to be reasonable in believing that *p* *do* depend on one's practical situation, then there ought to be situations in which *S* knows that *p* (because *S* meets the epistemic standards required to know that *p*) but it is not reasonable for *S* to believe that *p* (because in *S*'s practical situation the epistemic standards required for reasonable belief are too high for *S* to meet). But it does not seem possible that *S* could know that *p* in cases where it is not reasonable for *S* to believe that *p*. Given the tension between moderate invariantism and the view that what it is reasonable to believe depends on one's practical situation, the response under consideration does not seem available to the moderate invariantist.

The response in question might seem more congenial to the invariantist who defends SSI and a WAM against the Airport Case. SSI at least holds that practical differences can be epistemically relevant. However, the way in which SSI holds that practical differences are epistemically relevant seems at odds with the response we are considering.

Abstracting away from the details of particular versions, SSI holds that the strength of the epistemic position required in order for *S* to know that *p* (described in a way that factors out practical interests; that is, in terms of evidence, justification, the set of alternatives to *p* that *S* can rule out, and so on) varies with *S*'s practical situation. This makes it plausible, if not inevitable,

that whether or not S is reasonable to believe that p also depends on S 's practical situation. But nothing in SSI suggests that there could be two subjects, A and B , who are in the same epistemic position (practical interests aside) with respect to p , but that, in virtue of differences in their practical situations, it is reasonable for A to believe p and it is reasonable for B to believe not- p . In fact, if SSI is correct then there are reasons for thinking that this is impossible. If A and B differ only with respect to their practical situations, and in virtue of these differences A —but not B —both knows and is reasonable to believe that p , then B must have substantial evidence for p , and little or no evidence for not- p , since otherwise A would fail to know that p . So B , while not in a position to reasonably believe that p , is certainly not in a position to reasonably believe that not- p .

Let's apply this to the Airport Case. Suppose Smith and Mary differ only with respect to their practical situations, and it is reasonable for Smith to believe that the standards for knowledge are low. It might nevertheless be the case, if SSI is correct, that in virtue of differences between Smith's and Mary's practical situations, it is not reasonable for Mary to believe that the standards for knowledge are low. This is a far cry, however, from the claim that the response we are considering requires: that it be reasonable for Mary to believe that the standards for knowledge are high. Since Mary is in the same epistemic position (factoring out practical interests) as Smith with respect to the proposition that the standards for knowledge are low, Mary must have good evidence that the standards for knowledge are low, and little or no evidence that they are high. So, even if SSI is correct, it is not reasonable for Mary to believe that the standards for knowledge are high if it is reasonable for Smith to believe that the standards for knowledge are low.

Since neither moderate invariantists nor those who defend SSI can avail themselves of the response we have been considering, they all seem forced to deny what is required in order for a low-standards WAM to succeed, assuming that RTBR governs assertion: that it be reasonable for Mary to believe that Smith does not know *STOPS*, and that it also be reasonable for Smith to believe that he does know *STOPS*.

The knowledge account of assertion is inconsistent with an invariantist WAM against the Airport Case, and the knowledge account is very well supported by diverse linguistic evidence. Even if Weiner and Lackey are correct that other rules of assertion suffice to account for the evidence that motivates the knowledge account, no invariantist WAM against the Airport Case can succeed if one of the rules they suggest governs assertion. I conclude that no invariantist WAM against the Airport Case can succeed.

3. Implicature based replies

The preceding argument against WAMs relied entirely on DeRose's general characterization of a WAM—I did not examine any specific proposed WAMs in order to show how they succumb to the argument. Rysiew (2001), Black (2005), and Brown (2005a) have all responded to the Airport Case (or related cases) with what at least *purport* to be WAMs. When one examines these purported WAMs more carefully, however, it is not obvious that they really are WAMs, as characterized by DeRose. Since these purported WAMs may not really be WAMs at all, and since they all rely on Grice's theory of conversational implicature, it will be best to call them *implicature based replies* (or IBRs for short). If IBRs are not really WAMs at all, then they may well be immune to the argument against WAMs. Furthermore, even if the IBRs on offer really are genuine WAMs, it may be possible to modify them so that the resulting IBRs are

not WAMs and are immune to the argument. For these reasons, it is important to examine IBRs separately, in order to determine whether they can avoid the preceding argument and provide an adequate reply to the Airport Case.

In what follows, I do not want to get hung up on the issue of whether the IBRs to be found in the literature really are genuine WAMs, or whether the philosophers who offered them always intended to be presenting something somewhat different from a WAM, as characterized by DeRose. The argument from section 2 shows that if the IBRs *are* genuine WAMs, then they fail to adequately reply to the Airport Case. What is important, then, is to find one or more interpretations of IBRs according to which they turn out *not* to be genuine WAMs.

Although they differ in detail, the IBRs that have been offered in response to the Airport Case (and similar cases) share the same structure. They are all low-standards invariantist replies; that is, they all hold that (S2) semantically expresses a false proposition in c_2 . They all rely on Grice's (1989) account of conversational implicature and his maxims of conversation, and they all explain the problematic intuitions elicited by the utterances in the Airport Case by appealing to propositions that they argue are implicated by the utterances in the cases. They explain away intuitions that an utterance is false by arguing that the utterance, while semantically (literally, strictly) expressing a true proposition, implicates a false proposition. They explain away intuitions that an utterance is true by arguing that the utterance, while semantically expressing a false proposition, implicates a truth.

The argument from Section 2 relied on the fact that a low-standards WAM against the Airport Case holds that although Mary falsely asserts that Smith does not know *STOPS*, this assertion is nevertheless conversationally proper. An interpretation of IBRs that avoids the argument, then, must hold either (i) that, while Mary does assert that Smith does not know

STOPS, this assertion is *not* conversationally proper, or (ii) that Mary does not even *assert* that Smith does not know *STOPS*.

Consider first an IBR according to which Mary's assertion that Smith does not know *STOPS* is conversationally improper. Such an IBR would avoid the argument against WAMs from Section 2. However, the IBR can no longer explain the intuition that what Mary asserts is true by claiming that we mistake the assertion's conversational propriety for its literal truth. Some other explanation for the intuition will have to be found. Such an alternative explanation is suggested in Brown's defense of an IBR:

Suppose that a false utterance pragmatically conveys a truth. If speakers concentrate on what the utterance pragmatically conveys rather than on what is literally said, or if speakers mistake what the statement pragmatically conveys for what it literally says, then the utterance will seem correct even though it is literally false. (Brown 2005a, 284)

The explanation suggested by this passage is that one's intuitions concerning the truth-value of what Mary is asserting are sensitive to the truth-conditions of the implicated proposition, either because we mistakenly think that the implicated proposition just *is* the proposition semantically expressed by Mary's utterance, or because we are completely focused on the implicated proposition. One advantage of an IBR along these lines is that it can also explain why we (mistakenly) find Mary's utterance of (S2) to be conversationally proper—we just don't attend to the content of what she is asserting, and so we just don't notice that she is violating a rule governing assertion.

There is, however, a serious problem with an IBR of this sort. Briefly put, the problem is that the implicature posited by the IBR is possible only if hearers infer the implicated proposition from the utterance's semantic content, or at least are in a position to work out such an inference. As I will argue, this makes Brown's suggested explanations for the intuition that Mary asserts a

truth by uttering (S2) very dubious. First, however, a brief digression about Grice's theory of conversational implicature.

The crucial idea underlying Grice's theory of conversational implicature is that it is often reasonable to infer that a speaker intends to convey a proposition beyond the proposition semantically expressed by her utterance, since conversations are generally cooperative endeavors, and it is therefore reasonable to presume that the speaker is abiding by the *Cooperative Principle* (CP): "Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged" (Grice 1989, 26).

Grice uses four 'conversational maxims', most of which encompass several more specific maxims, to spell out what abiding by CP amounts to. The maxims are:

Quantity	Make your contribution as informative as is required. Do not make your contribution more informative than is required.
Quality	Do not say what you believe to be false. Do not say that for which you lack adequate evidence.
Relation	Be relevant.
Manner	Be perspicuous -Avoid obscurity. -Avoid ambiguity. -Be brief. -Be orderly.

Here is a sketch of how Grice thinks implicature works. Suppose a speaker utters a sentence s with semantic content p , and if the speaker were *only* conveying p by his utterance, then he would be violating one of the conversational maxims. Then it follows from the presumption that the speaker is abiding by CP that he intends to convey another proposition q such that he is abiding by all the maxims of conversation by conveying both p and q .

Grice thought that the hearer could often work out exactly what q is. In one of his more famous examples, a philosophy professor writes a letter of recommendation for a student that says only, “Dear Sir, Mr. X’s command of English is excellent, and his attendance at tutorials has been regular. Yours, etc.” (Grice 1989, 33) If the writer were conveying only what is semantically expressed by his words, then he would be violating the first maxim of Quantity, by not providing information that his readers require. Grice argues that since the writer is in possession of the required information, it is reasonable to conclude that he does not want to put the information on paper. One can conclude from there that the information is highly prejudicial to the student, and hence that the student is no good at philosophy. Given that the writer knew perfectly well that his readers would draw just this inference, it can be reasonably concluded that the writer intended to convey to the readers that the student is no good at philosophy. This, then, is what the writer non-semantically conveys, or *implicates*.

The important point to take away from this brief discussion of Grice’s theory is that in order to infer that a speaker is implicating a proposition q , she must first grasp the semantic content p of the speaker’s utterance. The hearer is reasonable to believe that the speaker intends to convey another proposition q only because she recognizes that if the speaker were conveying p alone, he would be violating a conversational maxim.

Grice is quite clear that in order for a proposition q to be implicated by an utterance with semantic content p , the hearer must be capable of working out that the speaker intended to convey q , in the manner sketched above (Grice 1989, 33). Therefore, the hearer must be capable of grasping that the utterance conveys p in order to grasp that it conveys q . Furthermore, since p is not itself grasped via an inference, were the hearer to actually perform the inference, she would recognize that the conclusion that the utterance implicates q is inferred, in part, from the

premise that the utterance conveys p , but that this premise is not itself the product of an inference.

Let's temporarily adopt an assumption that will turn out to be controversial, but makes the argument clear. Suppose that in order to grasp that an utterance with semantic content p implicates q , a hearer must consciously work out (and not just be capable of working out, or work out without being aware that one is doing so) that the utterance implicates q . Now consider whether Brown's suggested explanations for the intuition that Mary asserts a truth by uttering (S2) are at all plausible in light of this assumption.

Consider first the suggestion that we are so completely focused on the implicated proposition that we simply judge the utterance to be true. There are actually two ways to interpret this suggestion. The idea might be that although we are aware both of the proposition semantically expressed by Mary's utterance and of the implicated proposition, only the implicated proposition is of interest to us, so we assign the truth-value of the implicated proposition as the truth-value of what Mary asserted.

So interpreted, Brown's suggestion is not plausible. In standard cases of implicature where we are only interested in the implicated proposition but we recognize both the implicated proposition and the (uninteresting) proposition semantically expressed by the utterance, we have no tendency to assign the truth-value of the implicated proposition as the truth-value of what was asserted. Consider Grice's example of the letter-writer. Suppose that in addition to his philosophical ineptitude, the subject of the letter is chronically late and has terrible spelling. The letter-writer knows this, but writes that the subject is punctual and has good handwriting in order to implicate the true proposition that the student is no good at philosophy. If they became aware of, or already knew, these facts, the readers would judge that what the professor asserted in the

letter was false, in spite of the fact that they are only interested in the implicated proposition, and the implicated proposition (that the student is no good at philosophy) is true.

It might be objected that the argument above involves a case in which the hearers have evidence that the semantically expressed proposition is false; it could then be suggested that when hearers have either no evidence or conflicting evidence for the semantically expressed proposition, *then* they tend to assign the truth-value of the implicated proposition as the truth-value of what was asserted. To test this suggestion, consider another example of Grice's:

- A*: Where does *C* live?
B: Somewhere in the South of France. (Grice 1989, 32)

In Grice's original example, *A* and *B* are planning a trip and are considering visiting *C* when *A* asks the question. According to Grice, *B*'s utterance implicates that *B* does not know which city *C* is living in, roughly because if *B* knew more he should have said it. It seems clear in this case that, if *A* were to judge that *B* is asserting a truth, then this is because *B* judges the proposition semantically expressed—that *C* lives somewhere in the south of France—to be true, and not because *B* judges the implicated proposition—that *B* does not know which city *C* lives in—to be true. As it stands, however, this case tells us nothing about how a hearer assigns a truth-value to the proposition asserted by a speaker's utterance when the hearer has no evidence, or conflicting evidence, for the semantically expressed proposition. This is because by asserting that *C* lives somewhere in the south of France, *B*, a reliable informant, provides *A* with evidence that *B* lives somewhere in the south of France. So this is not a case where *A* has no evidence or conflicting evidence for the proposition semantically expressed, and cannot be used to test our suggestion.

For an adequate test, we need to modify Grice's example. Suppose, then, that *C* is a spy whose cover story is that he lives in the south of France. *A* knows about the cover story, and knows that *C* might be living anywhere in Europe, including, possibly, the south of France. So *A*

does not have evidence for or against the proposition semantically expressed by *B*'s utterance before she asks the question, and *B*'s answer, repeating as it does *C*'s cover story, does not provide *A* with new evidence that the semantically expressed proposition is true. Suppose that *B* is sincere and cooperative, and that *A* recognizes this, so that *B*'s utterance implicates, as in Grice's original case, that *B* does not know what city *C* is living in. Suppose further that *A* is not interested in where *C* lives, but is interested in what *B* *knows* about where *C* lives (perhaps *A* wants to assess *B*'s usefulness as an informant). Here we have a case in which *A* has no evidence for or against the proposition semantically expressed by *B*'s utterance, has every reason to judge the proposition implicated by *B*'s utterance to be true, and is only interested in the truth-value of the implicated proposition. If the suggestion above were correct, then *A* ought to judge that *B* asserted a truth. But this is not a judgment that *A* would, or should, make; the natural explanation for this is that *B* asserted that *C* lives somewhere in the South of France, that *A* knows what *B* asserted, and that *A* is well aware that she has no grounds for judging *this* proposition to be true. This case confirms that when we are aware of both the proposition semantically expressed by an utterance and the proposition implicated by the utterance, we have no tendency to assign the truth-value of the implicated proposition as the truth-value of what was asserted by the utterance, even when we don't have evidence for or against the semantically expressed proposition and we are only interested in the implicated proposition.

A different interpretation of Brown's suggestion is that we are so focused on the implicated proposition that we simply do not notice that the utterance conveys any proposition other than the implicated proposition. In those circumstances it is plausible that we would assign the truth-value of the implicated proposition as the truth-value of what *Mary* asserted. So interpreted, however, Brown's suggestion does not square with the assumption that in order to

grasp the implicated proposition, the hearer must consciously work it out from (in part) the proposition semantically expressed by the utterance. We can't be completely unaware of the proposition semantically expressed by Mary's utterance if we have to use this proposition in order to grasp the implicated proposition, and be aware of doing so.

Finally, consider Brown's other suggestion, that we mistakenly think that the implicated proposition is the semantically expressed proposition. Given our working assumption, in order to make this mistake we have to be aware of both p , the proposition semantically expressed, and q , the implicated proposition, and we have to be aware that whereas p is immediately, non-inferentially grasped, q is grasped only by performing an inference from p , together with various claims about the speaker's communicative intentions and the conversational context. Given this awareness, it would be bizarre for us to judge that q , not p , is the proposition semantically expressed by Mary's utterance.

Brown's suggested explanations for the intuition that what Mary asserts is true, then, do not square with our working assumption that we grasp an implicated proposition by consciously working out an inference from the utterance's semantic content. This assumption, however, seems unrealistic in two respects. First, it is plausible that we can sometimes grasp the implicated without performing any inference at all. Second, it is plausible that when we do perform the inference, we are often not fully aware of doing so. Let's consider each of these problems in turn.

Rysiew (2001) thinks that we often grasp the proposition implicated by an utterance without performing any inference from the utterance's semantic content:

Because the relevant sentences are *standardly* used in the manner being described, far from having to be conscious and/or explicit, the hearer's inference to what is being conveyed is liable to be of the default variety. And precisely because, in such cases, hearers will 'hear' what's conveyed in the sentences themselves, the present account actually predicts that the audience will most often end up regarding the original utterance as true. (Rysiew 2001, 492)

It is not clear from this passage what an inference of the “default variety” is supposed to be.

Rysiew cites Bach in this connection:

Our seemingly semantic intuitions are especially unreliable when there is a recurrent pattern of nonliterality associated with particular locutions or forms of sentence ... It seems that the hearer’s inference to what the speaker means is short-circuited, compressed by precedent (though capable of being worked out if necessary), so that the literal content of the utterance is apparently bypassed. (Bach 2000, 263 n.4)

The idea Rysiew takes from Bach, then, is that if a certain sentence s is standardly used to implicate a proposition q other than its semantic content p , then upon hearing s those who are familiar with this standardized implicature may grasp q without ever grasping p . In those circumstances the hearer would grasp only one proposition upon hearing s uttered, and this proposition would appear not to be the product of an inference. In such circumstances, it is plausible enough that the hearer would judge the utterance to be true, even though p is false.

Bach explicitly states, however, that hearers must at least be in a position to work out the inference from p to q , even if they ordinarily employ a compressed inference directly from s to q . It is not entirely clear what circumstances would prompt a hearer to eschew a compressed inference and instead work out the full inference from p to q . Presumably, though, if a hearer were presented with a suggested semantic content for s and a suggested inference from that suggested content to a suggested implicated proposition, the hearer should be in a position to determine whether this is the correct inference, and to work out the correct inference if the suggested inference is incorrect. This, however, is precisely the sort of prompting that is offered by the various IBRs in the literature. So, even if we initially judge Mary’s knowledge denial to be true, because we perform a compressed inference from (S2) to a true implicated proposition which we take to be the proposition semantically expressed by (S2) in c_2 , we should abandon

that judgment upon reading a philosophy article in which an IBR is defended, because the article would prompt us to recognize that (S2)'s *real* semantic content is false. That, however, does not happen. Even after reading articles defending IBRs, many philosophers continue to have the intuition that Mary asserts a truth in the Airport Case. Therefore, although Rysiew's appeal to standardized implicatures and compressed inferences may explain our *initial* intuition that Mary asserts a truth, it fails to explain why this intuition persists even after we consider the IBRs offered by Rysiew and others.

The other problem for our working assumption is the possibility that when we do perform an inference from an utterance's semantic content to an implicated proposition, we are often not fully aware of doing so. This is troublesome, since my arguments against Brown relied on the assumption that the hearer performing the inference would be aware of doing so. There are two ways of modifying the assumption in order to eliminate this unrealistic feature. According to the first, while we often perform such inferences without noticing that we are doing so, we can become aware of performing the inference when we are suitably prompted. For the reasons discussed above in connection to Rysiew, modifying the assumption in this way will not ultimately help Brown. While Brown may now at least be in a position to explain the *initial* intuition that Mary is asserting a truth by uttering (S2), she will not be able to explain why this intuition persists even after we are suitably prompted and become aware of performing the inference. One might modify the assumption more radically, and claim that there are certain inferences from semantic content to implicated proposition that we can *never* become aware of performing, regardless of how we are prompted. Notice, however, that if this claim is true then one can have no evidence that it is, since acquiring evidence of the claim's truth would require becoming aware of an inference that, if the claim is true, one cannot become aware of. So,

depending on such a claim in order to respond to my arguments against Brown would put the invariantist in the very weak position of relying on a claim for which there can be no evidence.

Let me sum up this discussion of the working assumption that I relied on in arguing against Brown. The assumption *is* unrealistic in two respects: it is plausible that the inference from semantic content to implicated proposition is sometimes bypassed in the manner described by Rysiew and Bach, and it is plausible that we sometimes perform such inferences without being fully aware of doing so. In both cases, however, one can become aware of the inference with suitable prompting. Once one becomes aware of the inference, Brown's suggestion that we assign the truth-value of the proposition implicated by Mary's utterance as the truth-value of the proposition Mary asserts again comes to seem deeply implausible.

So far, we have been considering IBRs that hold that Mary asserts the proposition semantically expressed by (S2) in c_2 . Such IBRs have been found wanting, either because they are *bona fide* WAMs as characterized by DeRose, and succumb to the argument of Section 2, or because they do not adequately explain the intuition that Mary's knowledge denial is true. We have not yet considered the possibility that Mary does not assert the proposition semantically expressed by (S2) in c_2 at all.

There are two ways of construing such a reply, depending on whether the implicated proposition counts as asserted. If it is held that the implicated proposition is not asserted, then it will turn out that *no* proposition is asserted when Mary utters (S2). This is implausible—it is quite clear that Mary is making some kind of claim when she utters (S2), even if the content of that claim is disputed. An IBR that denies that (S2)'s semantic content is asserted, then, has to hold that the implicated proposition *is* asserted.

This is an interesting suggestion. Note, however, that even if an IBR along these lines could be made to work, its success would not undermine the thesis of this chapter. According to such an IBR, when we hear the Airport Case we *correctly* recognize that the proposition Mary is asserting is true—the mistake in the preceding argument for contextualism is to assume that Mary is asserting the utterance’s semantic content, rather than the implicated proposition. So an IBR along these lines denies (P3), not (P2). Consequently, such an IBR does not threaten the thesis that all intuition-denying replies to the Airport Case fail.

This section has been devoted to considering the prospects for IBRs that are not *bona fide* WAMs, as characterized by DeRose. It has been argued that such IBRs either fail to adequately explain the intuitions elicited by the Airport Case, or must hold that the intuitions elicited by the Airport Case are correct. Therefore, IBRs do not provide an adequate intuition-denying reply to the Airport Case.

4. Bach and Belief-Removal

Like most invariantist replies to the Airport Case, Bach’s (2005) response holds that while Smith asserts a truth by uttering (S1), Mary asserts a falsehood by uttering (S2). In order to adequately respond to the Airport Case, then, Bach must explain the intuition that Mary asserts a truth by uttering (S2). Bach, however, never gets around to explaining this intuition. If one tries to construct such an explanation out of what Bach does say, the explanation turns out to be essentially the same as Williamson’s, to be discussed in section 5.

Bach’s treatment of the Airport Case depends on the idea that in high-stakes situations in which a lot turns on whether *p* is true, a subject may quite appropriately require a higher degree of justification for *p* in order to confidently believe that *p* than she would in a situation where *p* is

of little practical importance. A subject in a high-stakes situation may therefore fail to know that p (by failing to confidently believe that p) even though in a low-stakes situation she would confidently believe that, and so know that, p .

With respect to the Airport Case, Bach holds that Mary is in a high-stakes situation relative to *STOPS*. Mary therefore does not have the degree of justification for *STOPS* that she requires in order to confidently believe *STOPS*; so she neither confidently believes nor knows *STOPS*. Bach argues that this explains why it would be improper for Mary to attribute knowledge of *STOPS* to Smith—if Smith knows *STOPS* then *STOPS* is true, so asserting that Smith knows *STOPS* would in effect be asserting *STOPS*, which is not proper given that Mary does not confidently believe *STOPS* to be true.⁹

Explaining why it would be improper for Mary to attribute knowledge of *STOPS* to Smith, however, is a far cry from explaining why it seems true for Mary to deny that Smith knows *STOPS*. When Bach addresses the issue of Mary's knowledge denial, he focuses on explaining why Mary denies that Smith knows *STOPS*. According to Bach, Mary mistakenly thinks that the degree of justification that she requires, in her high-stakes context, in order to confidently believe *STOPS*, just is the degree of justification required of anybody, in any context, in order to *know* that *STOPS*. Rightly judging that Smith is no more justified in believing *STOPS* than she, Mary concludes that Smith does not know *STOPS*. In fact, however, Smith does know *STOPS*, because the degree of justification required for him to know *STOPS* is less than the degree of justification that Mary requires in order to confidently believe *STOPS*. Mary's knowledge denial, then, issues from her false (though perhaps excusable) belief that Smith does not know *STOPS* (Bach 2005, 76-77).

That is the substance of what Bach has to say about Mary's knowledge denial. The problem, of course, is that even if we accept Bach's explanation in its entirety, it is not clear how this helps explain the intuitions of the theorists who read the Airport Case and judge Mary's knowledge denial to be true.

The only way (that I can think of) of parlaying Bach's explanation for Mary's knowledge denial into an explanation of the desired sort is to suppose that when theorists read the Airport Case, they make the same mistake as Mary, and judge the degree of justification that she requires in order to confidently believe *STOPS* to be the degree of justification that anybody must have in order to know that *STOPS*. According to this explanation theorists would mistakenly judge on this basis that Smith does not know *STOPS*, and consequently mistakenly judge Mary's knowledge denial to be true. An explanation along these lines is not substantively different from Williamson's response to the Airport Case, to which I now turn.

5. Williamson and psychological bias

Williamson (2005) explains the intuition that Mary's knowledge denial is true by claiming that when we consider the Airport Case we empathize with Mary's plight to the extent that we, like her, overestimate the power of the error possibilities she is considering to prevent Smith from knowing *STOPS*:

When we as theorists contemplate Mary's position, we are struck by the disastrous consequences of believing falsely (in counterfactual circumstances) that the plane stops in Chicago ... That makes salient to us ... the weaknesses of the epistemic position that Mary and Smith share with respect to the proposition that the plane stops in Chicago. Consequently, when we consider from Mary's point of view whether Smith knows that the plane stops in Chicago, we give more weight to considerations that favour a negative answer. It therefore appears that Mary can truly assert 'Smith does not know that the plane stops in Chicago.' But the appearance is deceptive, the effect of psychological bias. (Williamson 2005, 234)¹⁰

Note, however, that if we have managed to get ourselves in a frame of mind in which we judge that Smith does not know that *STOPS* and consequently that Mary asserts a truth by uttering (S2), then in that frame of mind we will be just as prone to judge that Smith asserts a falsehood by uttering (S1). The Airport Case, however, elicits the intuition that Smith asserts a truth. Clearly, then, Williamson has to say that we theorists are not in a ‘pro-Mary’ frame of mind when we judge that Smith asserts a truth. Instead we are in a ‘pro-Smith’ frame of mind, in which either we are not thinking about the error-possibilities that Mary is contemplating, or we are no longer as impressed with their power to thwart knowledge.

Since one cannot be in a pro-Mary and a pro-Smith frame of mind at the same time, Williamson is committed to holding that theorists who consider the Airport Case do not *simultaneously* judge that both Smith and Mary assert truths in the Airport Case. Instead, he has to say that our judgment about whether Smith knows *STOPS* oscillates from true to false depending on whether we are in a pro-Mary or a pro-Smith frame of mind. There are, however, good reasons to think that we do simultaneously judge that both Smith and Mary assert truths.

First, the following intuitive report seem to be a perfectly normal way of characterizing one’s intuitions about the Airport Case:

(S5) It appears to me that Smith and Mary are both asserting truths in the Airport Case.

An utterance of (S5) is clearly reporting that at the time of the utterance, it appears to the speaker that both Mary and Smith are both asserting truths. Williamson would therefore have to claim that this seeming unobjectionable intuitive report is false.

Further, when we are in a pro-Mary frame of mind, we will tend to judge that she asserts a truth and that Smith asserts a falsehood; when we are in a pro-Smith frame of mind, we will tend to form the opposite judgments. So given Williamson’s explanation, the following reports

ought to properly describe theorists' intuitive judgments with respect to the Airport Case, at one time or another:

(S6) It appears to me that Smith asserts a truth, and that Mary asserts a falsehood.

(S7) It appears to me that Smith asserts a falsehood, and that Mary asserts a truth.

However, neither (S6) nor (S7) appear to accurately describe the intuitions typically elicited by the Airport Case.¹¹

Suppose that Williamson soldiers on, denies that (S5) accurately captures the intuitions elicited by the Airport Case, and holds that (S6) and (S7) really do capture our intuitions at one time or another. He would then have to contend with the following intuitive report, which seems even more normal and uncontroversial than (S5):

(S8) Both Smith's and Mary's assertions appear to be conversationally proper.¹²

Williamson holds that the knowledge account of assertion is the correct; let us suppose he is right about that. Since an assertion will be conversationally improper if it violates KR, presumably part of judging that an assertion is conversationally proper involves judging that the speaker knows the proposition being asserted, and hence that this proposition is true; one would therefore not utter (S8) unless one simultaneously judged that Smith and Mary are both asserting truths. One would therefore not utter (S8) unless (S5) accurately described one's intuitions at the time of utterance. Given Williamson's commitment to the knowledge account of assertion, then, it seems that he would also have to deny that (S8) accurately captures our intuitions when we read the Airport Case.

Williamson's response to the Airport Case explains our intuition that Mary asserts a truth by uttering (S2), but this explanation comes at too high a cost, since it commits him to holding that we never simultaneously judge both Smith's and Mary's assertions to be true or

conversationally proper. Since this seriously misdescribes the intuitions typically elicited by the case, I conclude that Williamson's reply to the Airport Case fails.

6. Stanley and salient questions

In contrast to the preceding responses to the Airport Case, which have all been offered by moderate invariantists, Stanley's response makes use of SSI (Stanley 2005, 85-96). A rough characterization of Stanley's version of SSI is that as the practical costs of being wrong about p get worse for S , the degree of justification that S must have in order to know that p gets higher. So, in the Airport Case Stanley holds that Smith knows *STOPS*, because he has little at stake with respect to *STOPS* and consequently consulting the airline schedule is enough justification for him to know *STOPS*; Mary, on the other hand, does not know *STOPS*—given her raised stakes with respect to *STOPS*, more justification than an airline schedule is required for her to know *STOPS*.

Those who defend SSI typically hold that knowledge is normatively related to assertion and practical reasoning; more specifically, they hold that one can assert that p without qualification and use p as a premise in practical reasoning if and only if one knows that p (Hawthorne 2004, 21-31). On this way of looking at things, the question of whether you know that p can clearly be of great practical significance to you, since answering this question can give you insight into what you should (or shouldn't) do and say. However, since whether one knows that p depends in part on one's own practical interests with respect to p , no conclusions about what *you* ought to do follow from the fact that somebody else, who may have more or less at stake with respect to p than you, knows that p . From a practical point of view, the salient

question is not whether somebody else knows that p , but rather whether that other person *would* know that p if she were in your practical situation.

The practically salient question for Mary, then, is not whether Smith knows *STOPS*, but rather whether Smith would know *STOPS* if he were in Mary's practical situation. Stanley holds that the intuitions of theorists considering the Airport Case are explained by their recognition that this is the salient question for John and Mary:

When *we* are asked about our intuitions about the case, we intuitively recognize that what John and Mary *really* care about is whether Smith would know, were he in John and Mary's practical situation. That is, we recognize that the proposition that John and Mary *really* want answered—*would Smith know were he in our practical situation*—is in fact false. So, we are strongly inclined to go along with John and Mary's judgments, since we recognize that they are perfectly correct about the information in which they are really interested. (Stanley 102-103)¹³

To facilitate discussion, let us grant some of Stanley's key points: first, that the salient question for John and Mary is whether Smith would know that *STOPS* if he were in their practical situation; second, that theorists who consider the Airport Case recognize recognize that this is the salient question for John and Mary; third, that John and Mary correctly judge that Smith would not know that *STOPS* if he were in their practical situation; fourth, that theorists recognize this judgment to be correct. Even granting all these points, it hardly follows that theorists will judge that Mary asserts a truth by uttering (S2).

It is easy enough to come up with cases in which we would judge that a person makes a false assertion even though the person correctly answers the question that is most salient to her and we recognize that she has correctly answered that salient question. Suppose, for example, that a construction foreman wants to ensure that there are no people left in a building before it is demolished. She correctly believes that there is only one person left in the building, but incorrectly believes that the janitor is the only person left in the building. In fact the explosives

expert is the only person left in the building. If, upon seeing the explosives expert leaving the building, the foreman were to say, ‘The janitor is leaving the building’, we would judge that assertion to be false, in spite of the fact that the foreman has correctly answered the salient question—whether the last person has left the building—and we recognize as much.

The fact that Mary has correctly answered the question that is most salient to her and that we recognize as much does not by itself explain why we judge that she asserts a truth by uttering (S2). Provided that the content of what Mary is asserting is different from the content of the answer to her salient question, and provided that we recognize these distinct contents, there is no reason to think that our judgments about the truth-value of what Mary asserts will be influenced by our judgment about whether Mary has correctly answered the question that is most salient to her.

Unfortunately, Stanley does not say more to explain why he thinks that our judgment that Mary has correctly answered the most salient question leads to our judgment that Mary asserts a truth by uttering (S2). Let me suggest two directions he might pursue.

First, Stanley might claim that although Mary asserts a falsehood, she thereby conveys the important truth that the answer to the practically salient question is ‘no’. So developed, Stanley’s response is an intuition-denying reply, since it holds that the intuition that Mary asserts a truth by uttering (S2) is false. Such an account would need to be fleshed out with a specific account of how this important truth gets conveyed; it seems likely, however, that fleshed-out in this way Stanley’s reply would end up looking very much like a WAM or an IBR, and would be subject to the criticisms of Sections 2 or 3.

Second, and more promisingly, Stanley might hold that by uttering (S2) Mary does *not* assert the proposition semantically expressed by (S2) in c_2 , that Smith does not know *STOPS*.

Instead, the reply would hold that Mary asserts the proposition that Smith would not know that *STOPS* if he were in Mary's practical situation. To convince, a reply along these lines would have to specify a mechanism via which Mary can use (S2) in c_2 to assert a proposition other than (S2)'s semantic content in c_2 ; merely pointing out that the asserted proposition answers the question of greatest salience to Mary will not suffice. Setting that issue aside, however, such a reply would nicely explain why we judge that Mary asserts a truth by uttering (S2): we correctly recognize that the proposition she *asserts* by uttering (S2) *is* true.

Unfortunately, Stanley's reply to the Airport Case fails to provide a complete explanation for our intuition that what Mary asserts when she utters (S2) is true. The most promising way to flesh out his explanation, I have suggested, involves claiming that by uttering (S2) Mary asserts a conditional claim that is not the proposition semantically expressed by (S2) in c_2 . This discussion of Stanley's reply therefore supports my contention that intuition-denying replies fail to adequately respond to the Airport Case, and that replies denying (P3) are more promising.

¹ Another famous context-shifting argument for EC begins with DeRose's (1992) Bank Cases.

² I take these different formulations to be synonymous, or nearly so—at any rate the differences in meaning, if there are any, will not concern us here.

³ I will not consider those replies to the Airport Case that fail to take on the intuitions elicited by the case as I have described them; that is, that both Smith and Mary are asserting truths in the Airport Case. For example, Hawthorne (2005, 160) argues that it would be improper for Mary to say 'Smith knows that the plane will stop in Chicago', on the grounds that properly asserting that p requires knowing that p . Thus Mary may properly say 'Smith knows that the plane will stop in Chicago' only if she knows that Smith knows *STOPS*, and thus (since knowledge is factive) only if she knows *STOPS*. It follows from Hawthorne's version of SSI, however, that Mary does not know *STOPS*. So—granting that Hawthorne's version of SSI is correct—Mary may not properly say, 'Smith knows that the plane will stop in Chicago'. This argument, however, utterly fails to address the relevant intuition, which is that Mary is stating a truth when she says 'Smith does *not* know that the plane stops in Chicago'. (It clearly does not follow, from the fact that Mary may not properly say, 'Smith knows that the plane stops in Chicago', that it would be either true or proper for her to say, 'Smith does *not* know that the plane stops in Chicago'.) Black and Murphy's (2005) article also does not address the relevant intuitions. They focus on responding to Cohen's (1999) objection that, if low-standards invariantism is correct, then it ought to be correct for Mary to say, 'We know that the plane stops in Chicago, but still we need to check further', even though this is clearly infelicitous. Black and Murphy argue that this infelicity is merely pragmatic; consequently, Cohen's objection fails. What Black and Murphy never address is the felicity—indeed, the apparent truth of—Mary saying, 'Smith does not know that the plane stops in Chicago.'

⁴ I want to distinguish low-standards invariantists, those who would respond to the Airport Case by denying that Mary is asserting a truth, from moderate invariantists. Moderate invariantism is an anti-skeptical position, in the sense that it holds that we really do know much of what we ordinarily take ourselves to know. But this is compatible with conceding that there is much else that we take ourselves to know that we do not in fact know. So it is possible

for a moderate invariantist to also be a high-standards invariantist, and to adopt the high-standards response to the Airport Case.

⁵ DeRose (2002) uses this strengthened version of the knowledge account of assertion in an argument against invariantism. The argument appears to fail, however, because the claim that KR is the *only* epistemic rule governing assertion is not itself well supported. For a detailed critique of DeRose's argument, see Brown (2005a).

⁶ For arguments using Moore paradoxical sentences in support of the knowledge account, see Moore (1962, 277), Unger (1976, 256-260), Williamson (2000, 253-254), and DeRose (2002, 180-181).

⁷ Strictly speaking, we assert propositions, and utter sentences. By 'asserting *s*', where *s* is a sentence, I mean asserting *s*'s semantic content by uttering *s*.

⁸ Lackey's formulation of RTBR (which she calls 'RTBNA', for 'reasonable to believe norm of assertion') includes a second necessary condition on proper assertion, namely: "if one asserted that *p*, one would assert that *p* at least in part because it is reasonable for one to believe that *p*" (Lackey 2007, 608). Since the conclusion that an invariantist WAM against the Airport Case cannot succeed if assertion is governed by RTBR follows from the first condition of RTBR alone, I will ignore the second condition in what follows.

⁹ Brown (2005b, 78-84) offers an illuminating discussion of Bach's position.

¹⁰ I have altered the original quote by systematically replacing Williamson's 'Hi' with 'Mary' and 'Lo' with 'Smith'.

¹¹ Incidentally, if (S6) or (S7) did accurately describe the intuitions typically elicited by the Airport Case, then the case would be completely useless in an argument for contextualism, anyway.

¹² Williamson cannot explain the apparent truth of (S8) by appealing to the sort of bias that, according to him, explains why Mary's utterance of (S2) seems true. Though this bias would explain why we find Mary's utterance of (S2) to be not only true, but conversationally proper, it should also lead us to judge that Smith's utterance of (S1) is not only false, but conversationally *improper*. So the bias Williamson posits cannot explain why we simultaneously find both Smith's and Mary's utterances to be conversationally proper.

¹³ The case actually discussed by Stanley is a fusion of DeRose's bank cases and the Airport Case, with characters named by Stanley. In the quoted passage, I have systematically replaced 'Hannah and Sarah' with 'John and Mary' and 'Bill' with 'Smith'. The unaltered quote is: "When *we* are asked about our intuitions about the case, we intuitively recognize that what Hannah and Sarah *really* care about is whether Bill would know, were he in Hannah and Sarah's practical situation. That is, we recognize that the proposition that Hannah and Sarah *really* want answered—*would Bill know were he in our practical situation*—is in fact false. So, we are strongly inclined to go along with Hannah and Sarah's judgments, since we recognize that they are perfectly correct about the information in which they are really interested."

CHAPTER 5

THE AIRPORT CASE (II): SPEECH ACT PLURALISM

1. Introductory remarks

Invariantist replies to the Airport Case have typically involved denying the intuition that Smith and Mary both assert truths by uttering (S1), in context c_1 , and (S2), in context c_2 , respectively:

(S1) I know that the plane stops in Chicago.

(S2) Smith does not know that the plane stops in Chicago.

In chapter 4 I argued that such intuition-denying replies fail across the board. In this chapter I will consider the prospects for a response to the Airport Case that concedes that the intuitions are true, but denies that this leads inevitably to EC. This response, briefly suggested by Cappelen and LePore (2005), relies on their thesis of *speech act pluralism*, according to which many propositions are typically asserted by a particular utterance of a sentence.¹ In section 2, I articulate the response to the Airport Argument based on speech act pluralism, and discuss how speech act pluralism is justified. In section 3, I develop a new argument for EC based on the Airport Case whose premises are consistent with speech act pluralism. This argument relies on the premise that Mary and Smith are both in *normal contexts*—contexts in which they are speaking literally and not making any relevant errors. In section 4, I argue that either Mary or Smith is not in a normal context, and hence that the new argument for EC also fails.

2. The response from speech act pluralism

Speech act pluralism is the view that whenever a sentence is uttered, a variety of distinct propositions are simultaneously asserted:²

No one thing is said (or asserted or claimed, or ...) by any utterance: rather, indefinitely many propositions are said, asserted, claimed, stated. What is said (asserted, claimed, etc.) depends on a wide range of facts other than the proposition semantically expressed. It depends on a potentially indefinite number of features of the context of utterance and of the context of those who report on (or think about) what was said by the utterance. (Cappelen and LePore 2005, 4)

If correct, speech act pluralism offers a potential invariantist response to the Airport Argument that does not involve denying the truth of the intuitions elicited by the Airport Case.

As Cappelen and LePore (henceforth, 'C&L') put it:

An epistemic corollary of [speech act pluralism], one that cannot be emphasized enough, is that intuitions about, and other evidence for, speech act content are not direct evidence for semantic content: an intuition to the effect that an utterance *u* said that *p* is not even prima facie evidence that *p* is the proposition semantically expressed by *u*. (C&L 2005, 145)

Even if the intuitions elicited by the Airport Case are true, EC does not follow directly, since the intuitions are about the truth-values of propositions that Smith and Mary *assert*, whereas EC is a thesis about the *semantic contents* of Smith and Mary's utterances. In order to argue for EC on the basis of the truth of the intuitions, the contextualist must defend a principle that appropriately relates the proposition or propositions asserted by an utterance and the utterance's semantic content.

In the Airport Argument, the premise linking assertion to semantic content is:

- (P1) The proposition asserted by Smith is the proposition semantically expressed by (S1) in c_1 , and the proposition asserted by Mary is the proposition semantically expressed by (S2) in c_2 .³

Given the definite descriptions ‘the proposition asserted by Smith’ and ‘the proposition asserted by Mary’, (P1) is true only if Mary and Smith each assert exactly one proposition with their respective utterances.⁴ If speech act pluralism is true, however, then Mary and Smith each assert multiple propositions with their utterances; hence, there is no such thing as *the* proposition asserted by Smith and *the* proposition asserted by Mary. So if speech act pluralism is correct, then (P1) is not true, and the Airport Argument fails.

The contextualist might respond to this objection by replacing (P1) with a premise that does not presuppose that Smith and Mary each assert exactly one proposition. For example, it might be suggested that however many propositions are asserted by Smith and Mary, when we intuitively judge that, say, Smith asserts a truth by uttering (S1) in c_1 , we are focusing on just one of these propositions. Further, it might be claimed, the proposition we focus on when we form this judgment is the utterance’s semantic content. So (P1) might be replaced by something along the lines of (P2):

(P2) When one intuitively judges that Smith is asserting a truth by uttering (S1) in c_1 , one is focusing on (S1)’s semantic content in c_1 and judging that it is true, and when one intuitively judges that Mary asserts a truth by uttering (S2) in c_2 , one is focusing on (S2)’s semantic content in c_2 , and judging that it is true.

It is far from obvious that (P2) is true if Smith and Mary each assert multiple propositions with their utterances. If multiple propositions are asserted by a single utterance, then it seems possible for the most conversationally salient proposition asserted by the utterance to be distinct from the utterance’s semantic content. Furthermore, it is plausible that if one is going to focus on just one of the asserted propositions in order to form a judgment about whether the utterance asserts something true, then one will focus on the proposition of greatest interest; that is, the most

conversationally salient proposition. So (P2) is doubtful, since if one focuses on just one of the asserted propositions, it is quite likely that one will be focusing on a proposition other than the utterance's semantic content.

Suppose the contextualist changes course and holds that when one judges that Smith asserts a truth by uttering (S1) in c_I , one considers *all* of the propositions asserted by Smith's utterance. This does not help, since if just one of the asserted propositions is true, and all the others are false, it is still true that Smith asserts a true proposition. So, provided that at least one of the asserted propositions other than the utterance's semantic content is true, one could correctly judge that Smith's utterance asserts a truth even if the utterance's semantic content is false. Consequently, one cannot infer that (S1)'s semantic content in c_I is true from the fact that when one considers all the propositions asserted by Smith's utterance of (S1), one correctly judges that one of them is true.

In response, a contextualist might note that in addition to having the intuition that Smith and Mary both assert truths in the Airport Case, we do *not* have the intuition that either Smith or Mary asserts any falsehoods in the case. The contextualist might then argue that if we do consider all the propositions asserted by Smith and Mary's utterances, and if one of these asserted propositions were false, then we *would* have the intuition that a falsehood is asserted; since we do not have the intuition that a falsehood is asserted, all of the propositions asserted by Smith and Mary's utterances, including the utterances' semantic contents, must be true.

This contextualist response is not convincing, however. If speech act pluralism is correct and each utterance asserts multiple propositions, and if in judging whether an utterance asserts a truth we evaluate all of the asserted propositions, then this process of evaluation must be largely non-conscious. Certainly, I do not find myself consciously sorting through numerous

propositions whenever I try to determine whether somebody is speaking the truth. If the process of evaluation is non-conscious, then it is an open question how conscious intuitive judgments of truth or falsity emerge from this non-conscious process of evaluation. One possible answer, seemingly assumed by the contextualist response just articulated, is that if one non-consciously evaluates one of the asserted propositions as false, then one will consciously judge that the utterance asserts a falsehood. There are other possible answers, however. For example, it might be that if the most salient proposition asserted by an utterance is true, then one consciously judges that the utterance asserts a truth, but forms no corresponding conscious judgment of falsity. Or, one might consciously judge that the utterance asserts a truth yet form no corresponding conscious judgment of falsity provided that *any* of the asserted propositions is true. Finally, it might be that one forms a conscious judgment of falsity only if one of the asserted propositions is both false and sufficiently salient. As far as I know, the question how our conscious intuitive judgments about truth and falsity are formed given the thesis of speech act pluralism is completely unexplored. The contextualist may therefore not adopt one theoretically convenient answer to this question without justification.

If speech act pluralism is true, then (P1) is not true and the Airport Argument fails. One might try to replace (P1) with a premise that is consistent with speech act pluralism. As we have seen, however, the natural replacement premises are not obviously true and therefore require justification that contextualists have not provided. Therefore, if speech act pluralism is true then the Airport Argument is inadequate as it stands.

What has been established so far is (at best) a conditional claim: *if* speech act pluralism is true, then the Airport Argument fails and it is not clear how it can be patched up. In order to pose a genuine threat to the Airport Argument, however, speech act pluralism itself stands in need of

justification. Such justification has been provided in the form of cases involving a particular utterance of a sentence and a seemingly correct indirect report of the utterance whose *that*-clause has different semantic content from the semantic content of the utterance.

Consider first the following example of C&L's:

(S3a) *A*: I bought a pair of Bruno Magli shoes and then I ate lunch.

(S3b) *B*: *A* said that he bought a pair of Bruno Magli shoes. (C&L 1997, 282)

Clearly, (S3b) is true; therefore, by uttering (S3a) *A* asserted *q*, the semantic content of (S3b)'s *that*-clause. Now, *q*, the proposition that *A* bought a pair of Bruno Magli shoes, is clearly distinct from *p*, the proposition that *A* bought a pair of Bruno Magli shoes and then ate lunch; however, it is clear that by uttering (S3a) *A* also asserted *p*. Therefore, by uttering (S3a) *A* asserted both *p* and *q*—two distinct propositions.

C&L offer two more examples along the same lines:

(S4a) *A*: I own a very expensive pair of Bruno Magli shoes.

(S4b) *B*: *A* said that he owns a pair of Bruno Magli shoes.

(S5a) *A*: At around 11 p.m., I put on a white shirt, a blue suit, dark socks and my brown Bruno Magli shoes.

(S5b) *B*: *A* said that he dressed around 11 p.m. (C&L 1997, 282-283)

It seems undeniable that in the preceding examples, *A* really does assert multiple propositions with a single utterance. However, in each of these cases the asserted proposition *q* that is not the semantic content *p* of the uttered sentence is nevertheless entailed by *p*—a conjunction entails each of its conjuncts, one cannot own a very expensive pair of Bruno Magli shoes without owning a pair of Bruno Magli shoes, and one cannot put on a shirt, a suit, and shoes without dressing. What these examples show is that multiple propositions may be asserted

by a single utterance, provided that the propositions that are not the utterance's semantic content are entailed by the utterance's semantic content.⁵

C&L offer other examples in which such an entailment relation does not obtain:

(S6a) François: Chartreuse is Maria's favourite colour.

(S6b) François said that *the color of that dress* [demonstrating a Chartreuse dress] is Maria's favourite color. (C&L 1997, 283; italics are C&L's)

Once again, (S6b) seems correct even though the semantic content q of (S6b)'s *that*-clause is distinct from the semantic content p of (S6a). In this case, however, there is no entailment from p to q , since the demonstrated dress might have been red, and it is therefore possible for p to be true and q to be false.⁶ Since François asserts both p and q , he asserts two distinct propositions, neither of which entails the other.

Here's another example of C&L's:

(S7a) *A*: Do you like that car? [pointing at a pink car]

(S7b) *B*: I hate pink cars.

(S7c) *A*: *B* said that he doesn't like that car [pointing at the same car]. (C&L 197, 284)

If (S7c) is correct, as it appears to be, then by uttering (S7b) *B* asserted q , the proposition that he doesn't like the car in question. But q is clearly distinct from p , the semantic content of (S7b).

Nevertheless, *B* also asserted p by uttering (S7b), since he did say that he hates pink cars.

Therefore, *B* asserted distinct propositions by uttering (S7b).⁷

I will conclude this discussion of the case for speech act pluralism with a slightly more complex example of C&L's. In this example, *A* is convinced that Stanley is Smith's murderer.

Looking right at Stanley, *A* says to *B*:

(S8a) *A*: Smith's murderer didn't comb his hair today.

Now suppose that *B* wants to report what *A* said to *C*, who does not know that *A* believes that Stanley is Smith's murderer, and in fact believes that somebody other than Stanley murdered Smith. It would be incredibly misleading for *B* to report *A*'s utterance with '*A* said that Smith's murderer didn't comb his hair today.' It would be less misleading, and still true, to report what *A* said as follows:

(S8b) *B*: *A* said that Stanley didn't comb his hair today. (C&L 1997, 284)

Once again, *p*, the semantic content of (S8a), and *q*, the semantic content of (S8b)'s *that*-clause, are distinct; nevertheless, if (S8b) is true then *A* asserts both *p* and *q* with (S8a).

3. A new argument for EC

The preceding examples and others like them present quite compelling evidence for speech act pluralism.⁸ However, they also suggest a reply on behalf of the contextualist, who can point to a common feature of the examples that leaves it unclear whether speech act pluralism can really help the invariantist respond to the Airport Case.

In order for speech act pluralism to be of use to the invariantist, it must not only be plausible that multiple propositions are asserted by one of the utterances in the Airport Case—it must also be plausible that one of the utterances semantically expresses a false proposition. The idea behind the appeal to speech act pluralism is to put pressure on the move from the correctness of an intuition that an utterance asserts a truth to the conclusion that the utterance's semantic content is true. The thought is that if there is a true proposition other than the utterance's semantic content that is nevertheless asserted by the utterance, then the intuition that the utterance asserts a truth is correct even if the utterance's semantic content is false. This strategy is clearly hopeless unless it is at least somewhat plausible that the utterance's semantic

content is false. As I shall explain, however, the examples C&L offer in support of speech act pluralism are cases where *all* the propositions asserted by the utterance, including the utterance's semantic content, would be true in any normal context.

By a *normal context*, I mean a context in which the speaker is sincere, is speaking literally, and is not making any relevant linguistic or factual errors. (An error is *relevant* when its occurrence might result in the speaker saying something false, as when one has a false belief about a name's referent, or a false belief about the facts being discussed.) In each of the preceding examples, we should expect both p , the semantic content of the utterance, and q , the proposition asserted by the utterance that is not its semantic content, to be true in any normal context. We would expect p to be true in any normal context, since if the speaker is sincere and is speaking literally, she will aim to utter a sentence that accurately represents the world. If she is in a normal context, she is not mistaken about how the world is and she is not mistaken about how the sentence represents the world; so, whenever she utters a sentence in a normal context, the world will be as the sentence represents it; that is, p will be true.

Some of C&L's examples indicate that a speaker may assert a proposition that she is not aware of asserting; for instance, it is doubtful that François knew that by uttering (S6a) he was asserting that the color of a particular dress (which might have been sewn after he uttered the sentence) is Maria's favorite color. One might worry, then, that speakers in normal contexts can unwittingly assert false propositions that are not the semantic contents of the sentences they utter. However, in all of the preceding examples, q (the proposition that is not the utterance's semantic content) either follows directly from p (the utterance's semantic content) as in (S3)-(S5), follows from p together with facts about the world, as in (S6) and (S7), or follows from p together with other beliefs of the speaker's, as in (S8). Therefore, in normal contexts, in which

the speaker is not making a relevant error, p and q will both be true, whether or not the speaker is aware of asserting q . (With respect to (S8), both p and q will be true provided that A 's belief that Stanley is Smith's murderer is true; if that belief is false, then A is making a relevant factual error, and is thus not in a normal context.)

The upshot is that C&L's examples should lead us to expect *all* of the propositions asserted by Smith and Mary in the Airport Case, including the semantic contents of their utterances, to be true provided that their utterances occur within normal contexts. Since speech act pluralism does not help the invariantist if the semantic contents of Smith and Mary's utterances are both true, invariantism is tenable only if there is some plausibility to the claim that either Smith's or Mary's utterance does not occur in a normal context; that is, it must be somewhat plausible that either c_1 or c_2 is not a normal context.

As we have seen, speech act pluralism is of little use to the invariantist unless she can make it plausible that either c_1 or c_2 is not a normal context. There is still the matter of fixing up the Airport Argument, since (P1) remains false—or at least not true—and no viable replacement premise has yet been found. Taking a cue from the preceding discussion, I suggest that the Airport Argument be reformulated as follows:

New Airport Argument

- (P3) In the Airport Case, c_1 and c_2 are normal contexts; that is, Smith and Mary are speaking sincerely and literally, and are not making any relevant factual or linguistic errors.
- (P4) If a speaker utters a sentence s in a normal context c , then the semantic content of s in c is true.
- (C1) Therefore, the semantic contents of (S1) in c_1 and (S2) in c_2 are both true.

From (C1), the argument proceeds as in the original Airport Argument, to the conclusion that EC is true.

The New Airport Argument's premises, (P3) and (P4), are plausible, consistent with speech act pluralism, and jointly entail EC. So the New Airport Argument is a *prima facie* compelling argument for EC that invariantists must reckon with. In the following section I present an objection to the New Airport Argument, by arguing against (P3).

4. An objection to the New Airport Argument

The thrust of the New Airport Argument is that since Smith and Mary are in normal contexts, the semantic contents of their utterances are true. In this section I will provide reasons for thinking that either Smith or Mary is not in a normal context.

Consider again the examples that C&L use to motivate speech act pluralism. If we imagine these utterances occurring in normal contexts, where the speaker is sincere, is speaking literally, and is not making any relevant factual or linguistic errors, we would expect *all* of the propositions asserted by the utterance to be true. We would expect the utterance's semantic content to be true, since in a normal context the speaker will aim to utter a sentence whose semantic content correctly depicts the world, and will succeed in this aim. We would expect the asserted propositions that are not the utterance's semantic content to be true, because in the cases C&L use to support speech act pluralism, a proposition that is asserted by an utterance but is not the utterance's semantic content nevertheless follows from the truth of the utterance's semantic content together with facts about the world or beliefs of the speaker's that would have to be true in a normal context.

Since we should expect all of the propositions asserted by an utterance to be true in any normal context, evidence that a false proposition is asserted by an utterance is evidence that the utterance does not occur in a normal context. In what follows I will argue that in the Airport Case, at least one false proposition is asserted by either Smith or Mary, and hence that either c_1 or c_2 is not a normal context.

By a *collective indirect speech report*, I mean an indirect speech report (a speech report that does not simply quote the words that were uttered) that reports what was said by more than one speaker. One makes a collective speech report when one assertively utters sentences of the form ‘ A said that p and B said that q ’ or ‘ A and B both said that p ’, among others. I will argue that accepting the intuitions elicited by collective indirect speech reports leads to the conclusion that at least one of the propositions asserted by Smith and Mary is false; for the reasons stated above, this is evidence that either c_1 or c_2 is not a normal context.

Consider the following extension of the Airport Case. Suppose that a bystander, Brown, overhears both Smith’s and Mary’s utterances. When Brown gets home to New York, he tells his wife the story of what happened at the airport. Brown reports what was said by Smith and Mary as follows:

(S9) Smith claimed to know that the plane stops in Chicago, but Mary said that Smith did not know that the plane stops in Chicago.

Intuitively, this collective indirect report is true.⁹ Let us suppose, as seems plausible, that ‘know’ expresses the same relation, which we will call ‘ r -knowledge’, throughout Brown’s report.

Granting the intuition that Brown’s report is true, it follows that by uttering (S1) Smith asserted that he r -knows *STOPS*, and that by uttering (S2) Mary asserted that Smith does not r -know *STOPS*. Since Smith’s epistemic position with respect to *STOPS* remains fixed from t_1 , when

Smith utters (S1), to t_2 , when Mary utters (S2), it follows that one of these assertions is false. Thus, the intuition that (S9) is true leads to the conclusion that either Smith or Mary asserts a falsehood in the Airport Case. Now, this does not show that the semantic content of either Smith's or Mary's utterance in the Airport Case is false, since given speech act pluralism Smith and Mary each asserts more than one proposition with his or her utterance; the false proposition might therefore not be the semantic content of the utterance that asserts it. It does, however, give us reason to believe that either Smith or Mary is not in a normal context. Since there is good reason to doubt (P3), the New Airport Argument fails.

How might the contextualist respond to this argument against (P3)? The most straightforward response is to simply deny that Brown's report is true. A *mere* denial of Brown's report, however, unaccompanied by a compelling explanation for the intuition that it is true, will not do. For, if the contextualist can simply deny Brown's report and leave it at that, then the invariantist should be entitled to simply deny that Smith and Mary are both asserting truths, and leave it at that. However, if either Mary or Smith asserts a falsehood, then he or she must not be in a normal context. So, if the contextualist is entitled to simply deny the truth of Brown's report, then the invariantist is entitled to simply deny the New Airport Argument's (P3). Of course, if the contextualist were to accompany her denial of Brown's report with a compelling explanation for why the report *seems* true, that would be sufficient to renew the threat posed by the New Airport Argument. Such an explanation, however, has yet to be provided.

Instead of denying Brown's report, the contextualist might attempt to interpret the content of the report in a way that avoids the result that either Mary or Smith asserts a falsehood in the Airport Case. I will consider two such alternative interpretations of Brown's report. According to the first, different occurrences of 'know' in (S9) express different epistemic

relations; according to the second, there are unarticulated or implicit quotation marks around at least one occurrence of ‘know’ in Brown’s report.

The argument that either Smith or Mary asserts a falsehood if Brown’s report is true depended on supposing that ‘know’ has the same content throughout (S9). Suppose, however, that a context-shift occurs in the middle of Brown’s utterance of (S9), or, alternatively, that there is a semantic rule according to which when ‘know’ occurs in the context of an indirect speech report, its semantic value is the same as the semantic value it had in the context of the reported speech. Either way, it would turn out that the two occurrences of ‘know’ in (S9) semantically express distinct relations; if so, then the two propositions that are reported by (S9) as being asserted might both be true even though Smith’s epistemic position with respect to *STOPS* does not change from t_1 to t_2 .

Although one might attempt to reply by arguing that every occurrence of ‘know’ in (S9) has the same semantic value, it is more straightforward to simply produce collective indirect speech reports that are immune to this line of response. So suppose that in the preceding story, Brown utters one of (S10)-(S12) instead of (S9):

(S10) Smith claimed to know that the plane stops in Chicago, but Mary denied it.

(S11) Smith claimed to know that the plane stops in Chicago, but Mary disagreed.

(S12) Smith claimed to know that the plane stops in Chicago, but Mary said he didn’t.

(S10)-(S12) are intuitively true, and since ‘know’ occurs only once in each report, responses according to which ‘know’ shifts contents during the course of the report clearly do not apply.

Each report says that there is a certain proposition, having to do with Smith’s epistemic position, that Smith asserted and that Mary denied. Since the same proposition cannot be both true and false at the same time, it follows that either Smith or Mary is asserting a falsehood.¹⁰

The contextualist might deny (S10)-(S12) on the grounds that if EC is correct, then Smith and Mary are not really disagreeing. The intuition that Smith and Mary are disagreeing, it might be argued, is a result of attending only to the superficial form of (S1) and (S2), and of failing to recognize that 'know' is a context-sensitive term. Such a response misconstrues the dialectic, however. (S10)-(S12) are offered as part of an argument against (P3), a premise in an argument for EC. The contextualist cannot reply to this objection by assuming EC without begging the question.

The second way of reinterpreting Brown's report, inspired by Hawthorne's (2006, 445-446) response to a test for context-sensitivity involving indirect speech reports defended by C&L (2005, 88-99), notes that reports beginning with 'said that' often have the feel of *mixed quotation*, where part of the report quotes the very words used by the speaker, and part of the report reports, without quoting, what the speaker said. Mixed quotation is pervasive in journalistic contexts; for example:

Agee said that in his case, he disclosed the identities of his former CIA colleagues to "weaken the instrument for carrying out the policy of supporting military dictatorships" in Greece, Chile, Argentina, Uruguay and Brazil. Those regimes "were supported by the CIA and the human cost was immense: torture, executions, death squads," he said. (*New York Times*, 9 January 2008)

Hawthorne notes that when we use mixed quotation in print we are often quite cavalier about putting quotation marks around the directly quoted material. For example, it is plausible that (S13) involves mixed quotation even though it contains no quotation marks:

(S13) If someone in Oxford in the 1860's said that they were going to see a grinder, they meant that they were going to see someone who prepares students for examinations. (Hawthorne 2006, 445)

It is very plausible that (S13) is a case of mixed quotation where the quotation marks have been dropped from around ‘grinder’. In spite of the lack of quotation marks, it is implausible to think that ‘grinder’ is being used, rather than mentioned, in (S13), since in twenty-first century English ‘grinder’ no longer has the archaic meaning that is being attributed to it in (S13).

When a speech report involving mixed quotation does not use quotation marks to explicitly indicate the quoted material, it can easily be confused with a pure, non-quotational, indirect speech report. The possibility suggested by Hawthorne’s response to C&L, then, is that Brown’s report is really a case of mixed quotation, where the quotation marks that ought to be around ‘knows’ are only implicit in (S9). So, according to this response, Brown’s report is more accurately represented by (S14):

(S14) Smith claimed to ‘know’ that the plane stops in Chicago, but Mary said that Smith did not ‘know’ that the plane stops in Chicago.

This reinterpretation of Brown’s report raises two issues. First, is it really plausible that Brown’s report involves mixed quotation? Second, granting that (S14) is the most accurate representation of Brown’s report, does this defeat the response to the New Airport Argument that I have been defending? I will consider these issues in turn.

We should not deny that Brown *could* correctly report Mary’s and Smith’s utterances using mixed quotation. Ordinarily, a verbal report involving mixed quotation would be accompanied by cues, such as ‘finger-quotes’, a certain intonation, or the insertion of the word ‘quote’, to indicate the quoted material. It does not seem impossible, however, for Brown to make a verbal report involving mixed quotation even without such cues; perhaps Brown’s intention to quote certain words in the report is all it takes to do so. It might be, then, that Brown can report what Smith and Mary said using mixed quotation even though from the hearer’s

perspective, Brown's report is indistinguishable from a pure indirect report. We can grant all this, without conceding that Brown *must* use mixed quotation to report Smith and Mary's utterances. There is no obvious reason why Brown could not, if he wanted to, use (S9) to make a pure indirect report of Smith and Mary's utterances. Supposing the aforementioned cues indicating quotation marks are absent, what distinguishes the case in which Brown uses mixed quotation from the case where Brown makes a pure indirect report? Presumably, the difference lies in Brown's communicative intentions—whether he wants to mention some of the words in Smith and Mary's utterances, or whether he is only interested in what they say, not how they say it. We would expect Brown to intend to quote certain words only when he finds the words used interesting or telling in some way. So, if Brown were a linguist or a linguistically-minded philosopher, we might expect him to use mixed quotation to highlight the word 'know' in Smith and Mary's utterances. In ordinary cases, however, where Brown is not especially interested in the word 'know', we would not expect him to use mixed quotation.

We can just stipulate, then, that Brown is not especially interested in the word 'know', and intends to use (S9) to make a pure indirect report of what Smith and Mary said. The intuition that Brown's report is true, it seems to me, is unaffected. Therefore, the argument against (P3) is unaffected by the fact that Brown *could*, if he had special interests, use mixed quotation to report on what Smith and Mary said, since he could also use (S9) in a pure indirect report, and the report would remain intuitively true.

Suppose that the preceding considerations are somehow wrong and that Brown *must* use mixed quotation to report Smith and Mary's utterances. Even so, it is doubtful that the argument against (P3) is seriously threatened. As C&L (2006, 476-477) point out in their response to Hawthorne, there is a near consensus among those interested in mixed quotation that the

semantic content of a report involving mixed quotation includes at least the semantic content of the sentence obtained by stripping off the quotation marks. So, for example, most would agree that the semantic content of (S15a) includes *at least* the semantic content of (S15b):

(S15a) Agee said that those regimes ‘were supported by the CIA.’

(S15b) Agee said that those regimes were supported by the CIA.

The debate over mixed quotation concerns what semantic content the quoted material contributes *in addition to* the semantic content of the sentence obtained by stripping off the quotes.

If the consensus view is correct, then even if Brown cannot help but report Smith and Mary’s utterances using mixed quotation in a way that is best captured by (S14), part of the semantic content of Brown’s report is the semantic content of (S9). Presumably, then, our intuition that Brown’s report is true is at least in part an intuition that the semantic content of (S9) is true. So, given the standard view about mixed quotation, the argument against (P3) proceeds even if Brown cannot help but use mixed quotation to report on Smith and Mary.

5. Conclusion

I have argued that it follows from our intuitions about collective indirect speech reports that either Smith or Mary asserts a falsehood, and hence is not in a normal context. Consequently, there is reason to doubt premise (P3) of the New Airport Argument. The original Airport Argument, on the other hand, contains a premise, (P1), that is inconsistent with speech act pluralism, a well-supported position. Furthermore, there does not appear to be a justified replacement premise that could do the work that needs to be done in the original Airport Argument and yet is consistent with speech act pluralism. Since the original Airport Argument and the New Airport Argument both appear to fail, and I can see no other way of using the

Airport Case to argue for EC, there is good reason to believe that the Airport Case does not support EC.

Unlike the intuition-denying replies considered in chapter 4, my response to the Airport Case accepts the intuition that both Smith and Mary assert truths in the Airport Case. In order to use this intuition in support of EC, however, the contextualist must draw a semantic conclusion from a claim about what is asserted; this, I have argued, they fail to do.

Something I wish I could do, but cannot, is provide a plausible story that explains *why* at least one of the contexts in the Airport Case is non-normal. There must some sort of non-literality or a relevant factual or linguistic error at work, but I am not prepared to say just what is going on in the case. This is out of step with the standard invariantist approach to responding to the Airport Case, which has typically involved providing a specific account of the sort of error or non-literality that is at work in the case. A thoroughly satisfying response to the Airport Case would include such an account. Nevertheless, if the argument of section 4 is correct, then we have good reason to believe that the Airport Case does involve a non-normal context, even if we cannot yet say why that is.

In light of the fact that I have not provided an account of the nature of the error or non-literality involved in the Airport Case, the reader might be left with the impression that I have not entirely undermined the support provided for EC by the Airport Case, and that the Airport Case still provides us with some reason to think that EC, rather than invariantism, is correct. This strikes me as incorrect. If speech act pluralism is correct—and there is strong evidence that it is—then an argument for EC based on the Airport Case must rely on the premise that c_1 and c_2 are both normal contexts. In order for the Airport Case to provide us with reason to believe EC, then, there needs to be on balance more reason to believe this premise than not. But

contextualists have offered no positive arguments in support of the premise that c_1 and c_2 are both normal contexts.

It might be suggested that it is reasonable to presume, in the absence of contrary evidence, that an utterance occurs in a normal context, provided that the utterance seems perfectly *apropos*, and nothing about the case suggests that the speaker is making a mistake or speaking non-literally. I can easily grant this, since in the Airport Case that presumption is defeated: there *is* evidence that either c_1 or c_2 is a non-normal context, since (S9)-(S12) are intuitively true. Now, granted, the claim that either c_1 or c_2 is a non-normal context would be far stronger if I could provide a plausible account of the sort of error or non-literality at work in the Airport Case; nevertheless, the intuitive truth of (S9)-(S12) does provide *some* reason to believe that either c_1 or c_2 is non-normal. Even if the lack of any obvious error or non-literality in the Airport Case provides some countervailing reason to believe that c_1 and c_2 *are* both normal contexts, it seems a stretch to insist that this reason far outweighs the reason to believe that either c_1 or c_2 is non-normal provided by the intuitive truth of (S9)-(S12). Since the argument for EC based on the Airport Case must rely on a premise that it is not all-things-considered reasonable to believe, the Airport Case provides us with no good reason to believe EC—and this is what I set out to show.

¹ For their discussion of speech act pluralism, see Cappelen and LePore (2005, 4-5 and 190-207). For their suggestion that speech act pluralism might be used to respond to contextualist arguments, see Cappelen and LePore (2005, 4 and 145). In a single-authored paper, Cappelen (2005) uses speech act pluralism to defend a form of skeptical invariantism against certain objections; for example, he responds to the objection that if skeptical invariantism were true, then the majority of our ordinary knowledge attributions would turn out false, by pointing out that if speech act pluralism is true then the majority of knowledge attributions might still assert true propositions even if their semantic contents are false.

² This statement of speech act pluralism should probably be somewhat qualified, since utterances of questions or commands presumably do not *assert* any propositions at all, and one can also utter a declarative sentence without asserting anything, if for example one is a French teacher going over conjugations, or an actor on the stage; it would be more accurate, then, to say that every *assertive* utterance of a *declarative* sentence simultaneously asserts multiple propositions

³ Numbered (P3) in chapter 4.

⁴ Given Russell's theory of descriptions, (P1) is false if Mary or Smith assert more than one proposition; on Strawson's account, (P1) would turn out to lack a truth-value; on neither account would (P1) turn out to be true.

⁵ We should not assume, however, that whenever an asserted proposition p entails another proposition q , that q is also asserted. If I assert that there are two people in the room, the semantic content of this assertion logically entails that the square root of the number of people in the room is an irrational number; but I do not assert this latter proposition when I assert that there are two people in the room.

⁶ Although there is no entailment from p to q , there is a weaker sort of implication at work here: q is entailed by p and contingent facts about the world that do not by themselves entail q (in this case the contingent fact that together with p entails q but that by itself does not entail q is the fact that the demonstrated dress is Chartreuse).

⁷ Here again the semantic content of (S7b) does not entail (S7c)'s *that*-clause—it is possible for A to hate pink cars without disliking a particular pink car, if A has not considered that particular car, or, having considered it, still does not know that it is pink (if it is covered in a thick layer of gray dust, for example.). However, (S7c)'s *that*-clause does follow from (S7b)'s semantic content together with the facts we are naturally inclined to read into the case—that A has considered the demonstrated car and does recognize that it is pink—and psychological facts about human beings—for example, that if one hates all F s and recognizes of x that it is F , then one will hate x .

⁸ Soames (2002) also defends speech act pluralism, and offers several examples in which an utterance asserts a proposition other than its semantic content. Here are two:

Coffee, Please

A man goes into a coffee shop and sits at the counter. The waitress asks what he wants. He says, "I would like coffee, please." The sentence uttered is unspecific in several respects—its semantic content does not indicate whether the coffee is to be in the form of beans, grounds, or liquid, nor does it indicate whether the amount in question is a drop, a cup, a gallon, a sack, or a barrel. Nevertheless, it is obvious from the situation what the man has in mind ... [The waitress] brings the man a cup of freshly brewed coffee. If asked to describe the transaction, she might well say, "He ordered a cup of coffee" or "He said he wanted a cup of coffee," meaning, of course, the brewed, drinkable kind. In so doing, she would, quite correctly, be reporting the content of the man's order, or assertion, as going beyond the semantic content of the sentence he uttered. (Soames 2002, 78)

The Terrorist

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 says, "I will detonate the bomb if my demands are not met" ... The negotiator reports to his superior that the terrorist said that he will kill thousands of people if his demands are not met. This report seems correct. (Soames 2002, 79)

⁹ If this intuition needs support, consider the intuitive falsity of denying (S9), for example:

'Smith claimed to know that the plane stops in Chicago, and Mary did not deny this.'

'Smith did not claim to know that the plane stops in Chicago, but Mary said that Smith did not know that the plane stops in Chicago.'

'Smith did not claim to know that the plane stops in Chicago, and Mary did not say that Smith did not know that the plane stops in Chicago.'

¹⁰ It does not help to suggest that the proposition Smith asserts is about his epistemic position with respect to *STOPS* at t_1 whereas the proposition Mary asserts is about Smith's epistemic position with respect to *STOPS* at t_2 , since Smith's epistemic position with respect to *STOPS* remains unchanged from t_1 to t_2 .

CHAPTER 6

CAPPELEN AND LEPORE'S CRITIQUE OF CONTEXTUALISM

1. Introductory remarks

In addition to suggesting the response to the Airport Case pursued in chapter 5, C&L (2005) present several other arguments directed at various forms of contextualism, EC included. In this chapter I wish to discuss two such arguments that strike me as especially interesting and relevant to EC. According to the first, the methodology of contextualism—that is, arguing for views like EC on the basis of cases like the Airport Case—cannot avoid a collapse into the radical position that every single English expression is context-sensitive. C&L then argue that this radical view is false; hence, they claim, the methodology of contextualism is unsound. According to the second argument, 'know' fails a test for context-sensitivity involving indirect speech reports; consequently, EC is false. Ultimately, I will conclude that these arguments fail; however, they fail for interesting reasons, and so deserve careful consideration. I will discuss the collapse argument in Section 2, and the test for context-sensitivity in Section 3.

2. The collapse argument

In order to adequately present and critique C&L's collapse argument, some initial exposition is required, having to do with C&L's characterizations of the different positions one might adopt with respect to the scope of context-sensitivity, and C&L's characterization of the methodology of contextualism.

The scope of context-sensitivity

Insensitive Semantics (C&L 2005) is a sustained argument for *semantic minimalism*, and against *moderate contextualism* and *radical contextualism*. These three views differ with respect to how pervasive they find the phenomenon of context-sensitivity to be. C&L present a list of linguistic items, based on a similar list of Kaplan's (1989, 489), that they take to be uncontroversially context-sensitive:

The personal pronouns 'I', 'you', 'he', 'she', 'it' ... the demonstrative pronouns 'that' and 'this' ... the adverbs 'here', 'there', 'now', 'today', 'yesterday', 'tomorrow' ... the adjectives 'actual' and 'present' ... words and aspects of words that indicate tense. (C&L 2005, 1)

C&L call this set of uncontroversially context-sensitive expressions the *basic set*. Semantic minimalism is the view that only members of the basic set are context-sensitive (C&L 2005, 2). Moderate contextualism is the view that there are some context-sensitive expressions outside the basic set, but that many other expressions are not context-sensitive (C&L 2005, 7-8). Finally, radical contextualism is the view that *every* expression of English is context-sensitive.

The thesis of radical contextualism applies not only to simple expressions like 'know', 'dog', and 'country', but also to complex expressions, including complete sentences:

RC [radical contextualism] is the view that all sentences are context sensitive. If so, then the propositions expressed by utterances of an arbitrary sentence S change depending on its contexts of utterance. The truth conditions for an utterance of an arbitrary sentence S vary between contexts of utterance. (C&L 2005, 128)

Among the ranks of radical contextualists, C&L include Searle (1980), Travis (1996), Recanati (2002), Bezuidenhout (2002), and the relevance theorists Sperber and Wilson (1995) and Carston (2002). C&L provide ample textual support for the claim that these philosophers and linguists really do endorse radical contextualism, including the thesis that every sentence has context-sensitive truth-conditions. Here are a few representative quotes borrowed from C&L:

In general the meaning of a sentence has application (it only, for example, determines a set of truth-conditions) against a background of assumptions and practices that are not representable as part of the meaning. (Searle 1980, 221)

Any sentence that is produced itself has its content only against a background of assumptions. (Bezuindenhout 2002, 113)

Sentences, by themselves, do not have determinate contents. What gives them the determinate contents they have (in context) is the fact that they are used in performing meaningful actions. (Recanati 2002, 194)

Very plausibly, C&L claim that epistemic contextualists like DeRose, Lewis, and Cohen are moderate contextualists, not radical contextualists: while epistemic contextualists hold that the word 'know' is context-sensitive, they find this to be a surprising and important fact about the word 'know', not a reflection of an absolutely ubiquitous linguistic phenomenon.

The methodology of contextualism

Suppose e is some expression outside the basic set. According to C&L, the entire case for contextualism about e (the claim that e is semantically context-sensitive) rests one or both of two kinds of argument: incompleteness arguments and context-shifting arguments (C&L 2005, 17).

An incompleteness argument begins with the intuition that a sentence containing e only determines a complete (truth-evaluable) proposition relative to a context of utterance; for example, the sentence 'Sam is ready' does not seem to express a complete proposition independent of a context that informs us of what it is that Sam is ready for (C&L 2005, 33-38). A context-shifting argument, according to C&L, is an argument that appeals to intuitions elicited by imagining a sentence s uttered in a range of different contexts. These intuitions about utterances of s in different imagined contexts are then used to draw conclusions about s 's semantic content relative to those contexts (C&L 2005, 10 and 17-18).

This characterization of the methodology of contextualism contains two inaccuracies, one having to do with the structure of context-shifting arguments, the other having to do with the sorts of considerations that have been adduced in support of EC. First, C&L take for granted that the two most important real-world cases that motivate EC, DeRose's (1992) Bank Cases and Cohen's (1999) Airport Case, are examples of context-shifting arguments (C&L 2005, 26-29). Strictly speaking, however, these cases do not conform to C&L's characterization of context-shifting arguments. As Montminy (2006) points out, both Cohen's and DeRose's cases involve a sentence uttered in one imagined context and the sentence's *denial* uttered in another context. In what follows, I will use 'context-shifting argument' to refer both to arguments that imagine the same sentence uttered in various contexts, and to arguments that imagine a sentence and its denial uttered in various contexts. Second, C&L claim that *all* arguments for contextualism about some expression *e* are, at root, either context-shifting or incompleteness arguments. Certainly, C&L marshal a great deal of evidence that shows that contextualists do extensively rely on such arguments (C&L 2005, 17-38). EC, however, is also supported by its alleged ability to solve various epistemological problems, most notably the problem of providing adequate replies to skeptical arguments that arise from radical hypotheses or lottery considerations. Chapters 2 and 3 of this dissertation are intended to evaluate precisely these sorts of arguments for EC. Therefore, C&L's claim that contextualism is *exclusively* supported by context-shifting or incompleteness arguments is somewhat overstated.

The collapse argument

In broad strokes, C&L's collapse argument is this: moderate contextualism is an unstable position which inevitably collapses into radical contextualism; radical contextualism, however, is false; therefore, semantic minimalism is the only tenable position.

Moderate contextualism collapses into radical contextualism, C&L argue, because the incompleteness and context-shifting arguments that support the view that some specific expression is context-sensitive can be used just as effectively to argue, of any expression *e* whatsoever, that *e* is context-sensitive. So one cannot happily argue that this or that expression is context-sensitive while also holding that most English expressions are not context-sensitive, since the arguments that support the former position also support the denial of the latter position.

Since EC is supported by context-shifting arguments but not incompleteness arguments, I am particularly interested in whether it is true that any expression *e* whatsoever is amenable to a context-shifting argument as effective as the context-shifting arguments that support EC. An affirmative answer to this question, together with C&L's claim that radical contextualism is false, would entail that the context-shifting arguments that support EC are unsound.

One might have expected C&L to deny the very strong claim that any expression whatsoever is amenable to a context-shifting argument, and instead hold the somewhat weaker position that any expression whatsoever is subject to either a context-shifting argument or an incompleteness argument; but C&L really do make the stronger claim. Specifically, they argue for GEN:

(GEN) With sufficient ingenuity, a CSA [context-shifting argument] can be provided for any sentence whatsoever, and consequently, for any expression. (C&L 2005, 40)¹

C&L argue for GEN by presenting context-shifting arguments (many borrowed from radical contextualists) built around a fairly short list of sentences, including the following:

(S1) John went to the gym. (C&L 2005, 44-45)

(In one context, John counts as having gone to the gym only if he went inside and exercised. In another context, he counts as having gone to the gym if he walked to the gym but did not go in.)

(S2) Smith weighs 80 kg. (C&L 2005, 43)

(In one context, the conversation is about Smith's diet. The utterance is true if Smith weighs 80 kg with his clothes off. In another context, the conversation is about whether the maximum weight limit of an elevator has been exceeded. The utterance is true only if Smith weighs 80 kg with his clothes on.)

(S3) Jill didn't have fish for dinner. (C&L 2005, 45-46)

(In one context, Jill counts as having had fish only if she ate fish. In another, she counts as having had fish if she ordered fish.)

(S4) Justine destroyed those shoes. (C&L 2005, 46)

(In one context, one counts as having destroyed one's shoes only if one has damaged them structurally. In the other context, one counts as having destroyed one's shoes if one pours paint all over them.)

C&L set up the cases so that each of the preceding sentences is intuitively true when uttered in one of the imagined contexts, and intuitively false in the other imagined context. For example, if Smith weighs 80 kg with his clothes off and 83 kg with his clothes on, an utterance of (S2) is intuitively true in the first, 'diet' context, and intuitively false in the second, 'elevator' context.

C&L's procedure is obviously inductive: they are trying to establish GEN, a fully general claim, by presenting a finite list of sentences that are susceptible to context-shifting arguments. In a good inductive argument, the observed sample is representative of the population that the

claim is about. But the sentences C&L present appear to be carefully selected to be clear cases where context-shifting arguments apply. Since C&L tell us nothing at all about the procedure they used to select the sentences featured in their examples, it is questionable whether the sentences they focus on really are representative of the population of all sentences; consequently, the cogency of C&L's argument for GEN is unclear.

There are two additional problems with C&L's procedure. First, they do not seek out putative counterexamples to GEN, in order to ensure that context-shifting arguments really can be built around them. As we shall see, there are sentences that, on the face of it, are much less susceptible to context-shifting arguments than the sentences C&L consider. C&L should have addressed these tough cases. Second, recall that C&L describe context-shifting arguments as cases in which the same sentence can seem true when uttered in one imaginary context, false in another; EC, however, is primarily supported by cases in which a sentence seems true in one imaginary context and the sentence's *denial* seems *true* in a different imaginary context. In order to show that the methodology of moderate contextualism really does lead to a collapse into radical contextualism, C&L should have argued that *this* sort of context-shifting argument can be constructed around any sentence whatsoever.

Let's consider this latter objection first. Montminy (2006) has argued that the context-shifting arguments C&L provide in support of GEN are less convincing when altered so as to feature an utterance of the sentence and an utterance of the sentence's *denial* in another context.

(He calls such arguments 'context-shifting arguments involving denials', or CSAIDs.) As

Montminy writes:

Although CSAIDs can ... be applied to 'A is tall', 'A is rich', 'It's cold', and other similar sentences containing comparative adjectives, they do not generalize to all sentences. Consider again the sentence 'John went to the gym'. In a conversation about John's exercise routine, it may seem wrong or infelicitous to

utter this sentence if John merely went to the gym to chat with his friends and did not perform any exercise. However, in this context, speakers would not go so far as to assent to ‘John did not go to the gym’. (Montminy 2006, 11)

Montminy’s objection is weakened by the fact that he does not directly take on any of C&L’s actual examples, choosing instead to imagine (S1) uttered in different circumstances than those envisioned by C&L. The point is not nearly as effective when we return to the original circumstances envisioned by C&L, in which John only walks to the vicinity of the gym. Given these circumstances, while it seems true to utter (S1) if one is describing the trajectory of John’s nightly walk, it also seems true to say, ‘John did not go to the gym’, in a context where one is providing an account of the exercise John did that day to John’s personal trainer. Montminy provides a pair of contexts and a world in which it would seem true to utter (S1) in the first context but it would not seem true to utter the denial of (S1) in the second context. This shows that a CSAID cannot be generated from the world and contexts Montminy provides. It does *not* show that (S1) is not susceptible to any CSAID whatsoever; indeed, the world and contexts originally provided by C&L do permit a CSAID. Furthermore, as far as I can tell, the other examples provided by C&L (with worlds and contexts unaltered) also permit CSAIDs.

There is a logical point here that is worth emphasizing. In order to provide an effective CSAID for a sentence s , it is sufficient to provide a world w and contexts c and c' such that s has the same meaning in w as in the actual world and such that s is intuitively true when uttered in w and c and not- s is intuitively true when uttered in w and c' . It is not necessary to argue that s and not- s are intuitively true no matter which world and contexts they are uttered in. Thus, providing a world $w\#$ and contexts $c\#$ and $c\#\#$ such that s is intuitively true when uttered in $w\#$ and $c\#$ but not- s is intuitively false when uttered in $w\#$ and $c\#\#$ does nothing to show that s does not permit a CSAID. Unfortunately, that is how Montminy appears to be arguing.

So Montminy's objection is not convincing. Nevertheless, he raises an important point. The methodology of contextualism, at least insofar as EC is concerned, depends on the use of CSAIDs, and not just the sorts of arguments, involving the *same* sentence uttered in different contexts, considered by C&L. If C&L were only able to show that the use of this latter sort of context-shifting argument leads to a collapse into radical contextualism, then the moderate contextualist could happily reply that since she does not use such arguments, *her* position does not collapse into radical contextualism. Thus, C&L's collapse argument will not be effective against moderate contextualism unless they can establish the following interpretation of GEN:

For any sentence s , there is a possible world w and contexts c and c' such that:

- (i) s has the same meaning in w as in α , the actual world.
- (ii) An utterance of s is intuitively true in w and c , and an utterance of not- s is intuitively true in w and c' .

A couple points of clarification are in order. First, one might wonder whether condition (i) is really part of what C&L intend to be claiming with GEN. I think so, since if (i) were dropped GEN would be obviously true but uninteresting—if we can make s mean whatever we like in w , then it will come as no surprise that s can be true when uttered in one context but false when uttered in another—for example, just let s mean in w what 'I am hungry' means in α , and let c and c' be contexts involving different speakers, only one of which is hungry. Second, one might question whether GEN is only making the relatively weak claim that in *some* possible world (in which s means the same as in α), we assent to s in one context and not- s in another, and not the stronger claim that there are contexts in α in which we assent to s in one context and not- s in another. Given how C&L support GEN, it seems that they must only be making the weaker

claim, since the examples they use involve merely possible, not actual, circumstances in which certain sentences are uttered.

The collapse argument fails if there is at least one sentence s that does not permit a CSAID; that is, if there is no world w and contexts c and c' satisfying conditions (i) and (ii) above. Such counterexamples seem quite easy to come by. Consider the following sentences:

(S5) There are no married bachelors.

(S6) $5 + 7 = 12$.

(S7) Water contains hydrogen atoms.

In our present context, it seems true for me to utter any of (S5)–(S7). Furthermore, I cannot imagine a possible context in which it could seem true for me to utter the denial of any of these sentences. I can imagine a certain context in which it might be *appropriate* for somebody to utter the denial of (S5), ‘There is a married bachelor’; for example, it might be appropriate to utter the denial of (S5) when speaking of a man separated from his wife who lives like a bachelor. But this involves a loose use of ‘bachelor’, according to which ‘bachelor’ means something like, ‘adult man who lives alone and doesn’t pick up’. So, while this utterance of the denial of (S5) can be appropriate, it is literally false.² With respect to (S6) and (S7), I cannot even imagine a situation in which it could be appropriate to utter their denials, but perhaps this is a failure of imagination on my part. What I want to insist on is that there is no context in which the denials of (S6) or (S7) could be literally true. Furthermore, it does not seem to help to imagine (S5)–(S7) uttered in possible worlds other than α , provided that their meanings remain unchanged. Since (S5)–(S7) do not permit CSAIDs, GEN is false; therefore, C&L’s collapse argument fails.

One might question whether C&L could argue for a somewhat weaker principle than GEN and still create trouble for contextualism. For example, it might be suggested that C&L should argue for the following:

(GEN2) For any individual word e , there is a sentence s containing e that permits a CSAID.

An immediate problem with GEN2 is that it appears to be obviously true but uninteresting—for any individual word e , one can always construct a meaningful sentence containing both e and a member of the basic set; therefore, there will always be some sentence containing e that permits a CSAID. For example, the word ‘hippopotamus’ occurs in the sentence ‘that is a hippopotamus’, which clearly permits a CSAID. This, however, tells us nothing whatever about the linguistic features of the word ‘hippopotamus’, since the fact that the sentence permits a CSAID is clearly traceable to the linguistic features of the word ‘that’. More generally, the fact that every individual word is part of a sentence that permits a CSAID does not commit the contextualist to the claim that every individual word is context-sensitive; therefore, C&L cannot even use GEN2 to argue that moderate contextualism collapses into the view that every individual word is context-sensitive, let alone that it collapses into radical contextualism.

C&L *could* argue that moderate contextualism collapses into the position that every individual word is context-sensitive by defending the following principle:

(GEN3) For every individual word e , there is a sentence s containing e that permits a CSAID, where the fact that s permits a CSAID is clearly attributable to e .³

If contextualists are right that CSAIDs license conclusions about semantic context-sensitivity, then it follows from GEN3 that every individual word is context-sensitive. Furthermore, the

cases C&L present in support of GEN seem to support GEN3 just as well as GEN. Finally, unlike GEN, GEN3 does not appear to be susceptible to absolutely clear-cut counterexamples. Arguably, then, C&L *can* show that moderate contextualism collapses into the position that every individual word is context-sensitive. To distinguish it from radical contextualism, let's call this position *rampant contextualism*.

Even if C&L can establish that moderate contextualism collapses into rampant contextualism, this will not help them to expose a weakness in the methodology of contextualism. Recall that C&L's original argumentative strategy was to begin by showing that moderate contextualism collapses into radical contextualism, and then to show that radical contextualism is false. That strategy depended on establishing GEN, which (S5)–(S7) show is false. Thus, C&L's original strategy fails. Now we are investigating the prospects for a parallel strategy, in which C&L first establish that moderate contextualism collapses into rampant contextualism, then show that rampant contextualism is false. Let's grant that C&L can show that moderate contextualism collapses into rampant contextualism. It does not appear that they can carry out the second part of the strategy, since their arguments against radical contextualism do not apply to rampant contextualism.

The reason why C&L's arguments against radical contextualism do not apply to rampant contextualism is that these arguments exploit radical contextualism's claim that every *sentence* is context-sensitive. For example, C&L (2005, 128-130) argue that radical externalism is internally inconsistent, on the ground that if radical contextualism is true, then (S8) is context-sensitive:

(S8) Radical contextualism is true.

If (S8) is context-sensitive, C&L argue, then there are false utterances of (S8) in some contexts.

So there are *true* utterances of (S9):

(S9) Radical contextualism is false.

The thrust of the argument is that radical contextualism is internally inconsistent because its truth entails its falsity. Although there is much more that could be said about this argument—and C&L (2005, 130) do consider a number of possible replies on behalf of the radical invariantist—the details need not concern us here, because the argument clearly doesn't even get off the ground if the target view does not hold that *every* complete sentence is context-sensitive. Thus, an argument of this sort cannot be used against rampant contextualism.

Similarly, C&L argue that radical contextualism makes communication impossible, because in order to understand what somebody else said, the hearer would need to know a great deal about the speaker's context, and hearers do not usually have such information:

If RC [radical contextualism] were true, then what's said by an utterance by a speaker A in a context of utterance C depends, at least in part, on very specific features of C. (C&L 2005, 123)

Rampant contextualism, however, is consistent with some sentences not being context-sensitive; so the hearer might not require any contextual information at all in order to know what the speaker is saying when she utters some sentences. Since rampant contextualism is consistent with the possibility of (some) communication, C&L's argument that radical contextualism makes communication impossible is not effective against rampant contextualism.⁴

I have argued that C&L do not successfully establish GEN, and hence do not show that moderate contextualism collapses into radical contextualism. It is possible—though not certain—that C&L can establish that moderate contextualism collapses into rampant contextualism. However, C&L's arguments against radical contextualism do not apply to rampant contextualism. Since C&L fail to show that moderate contextualism collapses into a false position, they fail to show that context-shifting arguments are unsound.

3. A test for context-sensitivity

As part of their case against radical contextualism, C&L propose a test for context-sensitivity involving indirect speech reports.⁵ They argue that all expressions outside the basic set, including ‘know’, fail this test.⁶ Unfortunately, C&L’s test is unclear in important respects, which makes it difficult to evaluate both their argument and the critical responses to it. For example, Hawthorne (2006, 443-446) has severely criticized C&L’s test when applied to clearly context-sensitive expressions like ‘nearby’, but it is not clear that he is directing his fire at the strongest formulation of the test in the neighborhood of C&L’s test. I will attempt to shed some light on these issues. I will identify those aspects of C&L’s test that are importantly unclear, and formulate two clearer tests that seem to be in the ballpark of what C&L had in mind. Then, I will discuss Hawthorne’s criticism of C&L’s test and argue that the new tests are immune to his criticism. Finally, I will argue against the new tests.

C&L’s test

C&L’s test is motivated by the fact that every context-sensitive expression in the basic set “typically blocks inter-contextual disquotational indirect reports” (C&L 2005, 88). An inter-contextual disquotational indirect report (ICDIR) is defined by C&L as follows:

Take an utterance u of a sentence S by speaker A in context C . An [ICDIR] of u is an utterance u' in a context C' (where $C' \neq C$) of ‘ A said that S .’ (C&L 2005, 88)

For example, if A says ‘I know that Jordan is tall’ in context c , an ICDIR of A ’s speech is an utterance in a context $c' \neq c$ of ‘ A said that I know that Jordan is tall’. An ICDIR is ‘disquotational’ because the reported sentence occurs unaltered in the *that*-clause of the report.

Context-sensitive expressions in the basic set *do* typically “block” ICDIRs—that is, ICDIRs of utterances containing the expression in the basic set are typically false. For example, if I say ‘I am hungry’, most ICDIRs of this utterance (those occurring in contexts in which the speaker is not me) will turn out false. Similarly, if I say ‘It rained yesterday’, most ICDIRs of this utterance (those occurring a day or more after the reported utterance) will turn out false.

Putatively context-sensitive expressions outside the basic set, on the other hand, do not typically block ICDIRs. For example, if I say ‘Jordan is tall’, it would be true to report this utterance with ‘Leo said that Jordan is tall’ in most, if not all, contexts. Similarly, if I say ‘Saul Kripke knows that descriptivism is false’, it would be true to report that utterance with ‘Leo said that Saul Kripke knows that descriptivism is false’ in most, if not all, contexts.

To pave the way for a clear presentation of C&L’s test, two preliminary comments are in order.

First, C&L’s test involves the notion of an inter-contextual disquotational* indirect report (or ‘ICDIR*’ for short). An ICDIR* permits the sentence s^* occurring in the *that*-clause of the report to differ in some respects from the reported sentence s :

To be ‘disquotational*’ just means you can adjust the semantic values of components of S that are generally recognized as context sensitive, i.e., we just test for the controversial components. (C&L 2005, 89)

I will spell this out more fully. An ICDIR* of an utterance by a speaker A of a sentence s in a context c is an utterance occurring in a context $c' \neq c$ of a sentence of the form ‘ A said that s^* ’, where s^* is obtained from s by applying the following procedure: for every uncontroversially context-sensitive expression e in s , if e has distinct semantic values in c and c' , replace e with an expression e' that has the same semantic value in c' that e has in c . So, for example, if A says ‘I am hungry’ in c , an ICDIR* report of that utterance by a reporter $B \neq A$ in c' would be ‘ A said

that he is hungry'; since 'I' refers to different individuals (has different semantic values) in c and c' , we replace 'I' with 'he', which refers to the same individual, A, in c' that 'I' refers to in c .⁷

Second, C&L's test involves the notion of a 'relevantly different context'. Supposing that an utterance containing an expression e occurs in a context c , C&L say that a context relevantly different from c is one that is "different according to the standards significant according to contextualists about e ," (C&L 2005, 89) and differs "with respect to whatever features the contextualist thinks determine content" (C&L 2005, 94 n. 4). Presumably, then, contexts c and c' are relevantly different (with respect to the question of the context-sensitivity of e) just when contextualists about e say that e has different semantic values in c and c' . Contextualists about e , however, do not always agree about the contextual features that determine content, and some might simply hold that e is context-sensitive, without going so far as to provide an account of how context determines e 's content. So the notion of a relevantly different context only makes sense relative to a particular contextualist *theory* of how context relates to content: at a bare minimum, such a theory should specify two contexts in which e has distinct semantic values according to the theory. Note, however, that one might show that a particular contextualist *theory* is false, by showing that e has the same semantic value in the specified contexts, without thereby showing that contextualism about e is false (since contextualism about e just says that there are at least two contexts in which e has distinct semantic values). This point should be kept in mind, since it will lead to serious problems for C&L's test.

With these preliminaries out of the way, here is C&L's test:

Suppose you suspect, or at least want to ascertain whether, e is context sensitive. Take an utterance u of a sentence S containing e in context C . Let C' be a context relevantly different from C ... If there's a true disquotational* indirect report of u in C' , then that's evidence e is context insensitive. (C&L 2005, 89)⁸

By way of example, let's apply this test to 'tall', which is widely thought to be context-sensitive. Imagine that I say 'you are tall' to my three-year-old daughter Millie, in a context c where it is clear that I am comparing Millie to other three-year-olds. Many contextualists about 'tall' would hold that 'tall' has a different semantic value in c than it would in a context in which a different comparison class has been made salient—say, a context in which Millie is now being compared to all the students at her school, and the oldest students at the school are thirteen. So that context, call it c' , is relevantly different, in C&L's sense, from c . A disquotational* report in c' of my utterance in c would be somebody saying, in a context where the salient comparison class is all the students in the school, 'Leo said that Millie is tall'. This report is true in c' (let's suppose), so according to the test we have evidence that 'tall' is not context-sensitive.

Is the test coherent?

C&L's formulation of their test is puzzling in a number of respects. Sorting out these issues will lead to two new tests that are more clearly motivated and easier to assess than C&L's original test.

A first puzzle is that C&L's test is stated as a conditional, of the form 'If p , then that's evidence that q '. On the face of it, however, this conditional is plainly false. C&L's test is supposed to provide evidence that a certain expression e is not context-sensitive. If e is not context-sensitive, then e has the same semantic value in *every* context. C&L's test, however, asks us to consider only two contexts—the context of an utterance containing e , and the context of an ICDIR* of that utterance, where those two contexts are relevantly different, in the sense described above. If the ICDIR* is true, that's supposed to be evidence that e is not context-sensitive. At best, however, considering the context of the reported utterance and the context of

the ICDIR* could tell us that *e* has the same semantic content in those two contexts; how could that possibly tell us, or even provide us with any evidence, that *e* has the same semantic value in every possible context?

A possible reply to this worry is suggested by the way that C&L justify their test:

By definition, for *e* to be context sensitive is for *e* to shift its semantic value from one context of utterance to another. So, if *e* is context sensitive and Rupert uses *e* in context *C*, and Lepore uses it in context *C'*, and the relevant contextual features change, then it will just be an *accident* if their uses of *e* end up with the same semantic value. In particular, if Lepore finds himself in a context other than Rupert's and wants to utter a sentence that matches the content of Rupert's utterance of a sentence with *e*, he can't use *e*, i.e., he can't report Rupert's utterance disquotationally. (C&L 2005, 89)

This passage seems to contain an argument of the following form:

If *p*, then it's just an accident that *q*;

q;

Accidents are improbable;

Therefore, *p* is probably false.

Let's consider what we might substitute for *p* and *q* in order to do justice to the text and produce a strong anti-contextualist argument. Here's one natural attempt:

- (P1) If *e* is context-sensitive, then it's just an accident if there are distinct contexts *c* and *c'* such that a true utterance containing *e* occurs in *c* and a true ICDIR* of that utterance occurs in *c'*.
- (P2) There are distinct contexts *c* and *c'* such that a true utterance containing *e* occurs in *c* and a true ICDIR* of that utterance occurs in *c'*.
- (P3) Accidents are improbable.
- (C1) Therefore, probably, *e* is not context-sensitive.

This argument is unsound, since (P1) is false. No plausible contextualist theory of *e* will deny that for many pairs of contexts, *e* has the same semantic value in both contexts. Therefore, any plausible contextualist theory of *e* will hold that there are many pairs of contexts in which an utterance containing *e* in one context can be truly reported with an ICDIR* in the other context. Furthermore, if *e* had the same semantic value in every possible context but one, *e* would *still* be context-sensitive. In such a situation, however, it would be hardly be an accident if an utterance containing *e* in one context were truly reported by an ICDIR* in another context; in fact, the accident would be if we were to stumble on the one pair of contexts in which the ICDIR* turns out false! The likelihood that *e* has the same semantic value in two arbitrarily selected contexts depends on the degree of context-sensitivity *e* exhibits. But *e* could be context-sensitive even if it were highly probable, and hence no accident, that *e* has the same semantic value in two arbitrarily selected contexts. So (P1) is false.

One reason to think that the preceding argument does not capture C&L's justification for their test is that it drops all mention of *relevantly* different contexts. Here's an argument that builds that back in:

- (P4) If *e* is context-sensitive and *c* and *c'* are relevantly different contexts, then it is just an accident if there is a true ICDIR* in *c'* of an utterance in *c* containing *e*.
- (P5) There is a true ICDIR* in *c'* of an utterance in *c* containing *e*.
- (P3) Accidents are improbable.
- (C2) Therefore, probably, either *e* is not context-sensitive or *c* and *c'* are not relevantly different contexts.

Once again, this argument's first premise appears to be false. As I understand C&L, two contexts are relevantly different if *e* has distinct semantic values in *c* and *c'*. If that's right, then it

is *impossible* for e to have the same semantic value in relevantly different contexts. Now, it is clear from the previously quoted passage that C&L think that if there is a true ICDIR* in one context of an utterance containing e in another context, then e must have the same semantic value in both contexts. So, by C&L's own lights, there is *no* possible situation in which c and c' are relevantly different contexts, an utterance u containing e occurs in c , and a true ICIDR* of u occurs in c' . Since such a situation is impossible, it cannot accidentally occur.

It might be suggested that this problem is easy to fix by simply omitting mention of accidents, as follows:

- (P6) If e is context-sensitive and c and c' are relevantly different contexts, then there is no true ICDIR* in c' of an utterance containing e in c .
- (P5) There is a true ICDIR* in c' of an utterance containing e in c .
- (C3) Therefore, either e is not context-sensitive or c and c' are not relevantly different contexts.

This change, however, does nothing to remedy another problem, which affects both arguments equally. The problem is that the conclusions of the arguments no longer cast doubt on the claim that e is context-sensitive; instead, they cast doubt on a particular contextualist theory of e , according to which e has distinct semantic values in c and c' . So these arguments cannot adequately justify C&L's test, which concludes that (there is evidence that) e is not context-sensitive.

C&L's test tries to draw a universal conclusion, that (there is evidence that) e is not context-sensitive, on the basis of a single pair of contexts in which e has the same semantic value. On the face of it, this is a very dubious move. As we have seen, no credible interpretation of C&L's justification for their test succeeds in addressing this worry. It appears, then, that as it

is stated C&L's test is simply unworkable. In what follows I want to suggest two replacements for C&L's test. Each is closely related to the original test, and each is supported by a clear rationale. In what follows I will present and defend these tests, and consider how they might be used against EC.

C&L's test fails because it tries to draw a universal conclusion on the basis of a single instance. Surely, however, each new observation of two contexts in which *e* has the same semantic value provides *some* inductive confirmation for the claim that *e* is not context-sensitive, just as each new observation of a black raven provides *some* inductive confirmation for the claim that all ravens are black. If we observe a sufficiently large and varied set of pairs of contexts, and observe that *e* has the same value in every pair, then that would provide strong inductive support for the claim that *e* is not context-sensitive.⁹ Now, if C&L are right in assuming that *e* has the same semantic value in *c* and *c'* provided that there is an utterance containing *e* in *c* and a true ICDIR* of that utterance in *c'*, then we can verify that *e* has the same semantic value in two distinct contexts *c* and *c'* by imagining an utterance containing *e* occurring in *c* and an ICDIR* of that utterance occurring in *c'*. If the ICDIR* is true, then we know that *e* has the same semantic value in those two contexts. Repeated application of this procedure with respect to a large and varied enough set of pairs of contexts would provide strong inductive support for the claim that *e* is not-context-sensitive. (On the other hand, a single observation of a pair of contexts in which there is a *false* ICDIR* in one context of an utterance containing *e* in another would show that *e* is context-sensitive.) This suggests the following test:

Inductive Test (IT)

Consider a pair of distinct contexts *c* and *c'*. Imagine an utterance *u* containing *e* occurring in *c* and an ICDIR* of *u* occurring in *c'*. If the ICDIR* would be true,

that provides some confirmation that *e* is not-context sensitive. Repeated application of the procedure with respect to a sufficiently large and diverse set of pairs of contexts, with no observation of a pair in which the ICDIR* would be false, provides strong inductive support for the claim that *e* is not context-sensitive.

Let's apply IT to the case of 'know'. Suppose that after considering a large and varied set of pairs of contexts, we cannot find a pair of contexts such that a knowledge attribution occurs in one context, and a false ICDIR* report of the knowledge attribution occurs in the other. Then, if IT is correct, there is strong inductive support for the claim that EC is false. At this relatively late stage in the debate over EC, it is quite plausible that this is the position we are actually in. Epistemologists have imagined knowledge attributions occurring in all sorts of contexts. If there were contexts in which ICDIR*s of knowledge attributions occurring in other contexts turned out false, it is very likely that defenders of EC would have noticed them and used them to argue for EC. No such arguments have been proposed, however, and this suggests either that there are no such pairs of contexts, or at least that they are not easy to find. So, it appears that we do have strong inductive support for the falsity of EC, provided that IT is a sound test.¹⁰

The second test is inspired by the argument discussed above that makes explicit reference to relevantly different contexts and whose conclusion succeeds only in casting doubt on a particular contextualist theory of *e* rather than the general claim that *e* is context-sensitive. It might be quite useful to argue against a particular contextualist theory of *e*, especially if for some reason only a few contextualist theories of *e* have emerged as plausible contenders; if one could apply the same type of argument to all of these contenders, one would have built, in a

roundabout way, a strong case against the general thesis that e is context-sensitive. So consider the following test, aimed at a specific contextualist theory:

Theory Test (TT)

Let T be a contextualist theory of e , determining a set C of pairs of relevantly different contexts in which, according to T , e has distinct semantic values. If $\{c, c'\}$ is a member of C , and there is a true ICDIR* in c' of an utterance u in c containing e , then T is false.

The justification for TT again relies on C&L's assumption, to be examined later, that if there is a true ICDIR* in c' of an utterance in c containing e , then e has the same semantic value in c and c' . Given this assumption, the justification is straightforward: if it follows from T that e has distinct semantic values in c and c' , then it follows from T and the aforementioned assumption that any ICDIR* in c' of an utterance in c containing e is false. So, if there is a true ICDIR* in c' of that utterance, T is false.

Let's now apply TT to the case of 'know'. Since TT applies to a specific contextualist theory, we need such a theory as a target. According to Cohen's version of EC, in the Airport Case 'know' expresses different relations in c_1 , the context in which Smith says, 'I know that the plane stops in Chicago', and c_2 , the context in which Mary says, 'Smith does not know that the plane stops in Chicago'. Presumably, if Mary (with her practical situation unchanged, still worrying about whether the plane stops in Chicago, and still imagining possible scenarios in which it does not) were to report Smith's utterance with 'Smith said that he knows that the plane stops in Chicago', that would be an ICDIR* in c_2 of an utterance in c_1 containing 'know'. If, as intuition suggests, this report is true, it follows from TT that Cohen's version of EC is false.

Although C&L's original test for context-sensitivity is profoundly flawed, IT and TT have emerged as plausible replacements. Each is so closely related to the original test that it is not implausible to think that C&L had a test of this sort in mind all along; each is supported by a clear, if not conclusive, rationale; and each can generate *prima facie* compelling arguments against EC—IT can be used in an inductive argument that EC is simply false, whereas TT can be used to argue that specific versions of EC, like Cohen's, are false. It is therefore important to evaluate IT and TT and determine whether these tests really are sound. Soon I will present some new criticisms of IT and TT. First, however, I want to consider how these tests fare against the criticism that has been directed at C&L's original test. In the following section I will discuss Hawthorne's criticism of C&L's original test, and assess whether the same criticism applies to IT and TT.

Hawthorne's critique

Hawthorne's critique of C&L's test boils down to this: if the test is correct, then it would follow that 'nearby' is not context-sensitive; since this is absurd, the test is no good. Hawthorne characterizes C&L's test as follows:

If *u* utters a sentence *S* in a context, an inter-contextual Disquotational report in a different context *C'* will be an utterance of 'U said that *S*' in *C'*, where 'U' refers to *u*. The authors' idea is that if an expression is context-dependent then it will not in general be amenable to this style of reporting. (Hawthorne 2006, 444)

For an expression *e* to "not in general be amenable to this style of reporting" presumably means that there is at least one false ICDIR* of an utterance containing *e*. Reformulated in the terminology of the preceding section, C&L's test as characterized by Hawthorne is:

If an expression *e* is context-sensitive, then there are contexts *c* and *c'* such that an utterance *u* containing *e* occurs in *c* and there is a false ICDIR* of *u* in *c'*,

or, equivalently:

If there are no contexts c and c' such that an utterance u containing e occurs in c and there is a false ICDIR* of u in c' , then e is not context-sensitive.

How could we go about applying this test? There are infinitely many possible conversational contexts, so we cannot consider every single pair of possible contexts. The best we can hope for is to consider a set of pairs of contexts that is sufficiently large and diverse, and verify that there is no pair of contexts in that set that would falsify the antecedent of the test. Since this provides at best strong inductive support for the antecedent of the test, it provides at best strong inductive support for the consequent. As characterized by Hawthorne, then, C&L's test is not substantively different from IT.

If the preceding fairly portrays Hawthorne's characterization of C&L's test, then his criticism misses the mark. The meat of Hawthorne's criticism occurs in the following passage:

Suppose Ernie is in New York City and I am in Birmingham. Ernie says 'A nearby restaurant has good Vietnamese food'. I can report this by saying 'Ernie said that a nearby restaurant has good Vietnamese food,' even though I am far away from him ... By the lights of the test, then, 'A nearby restaurant has good Vietnamese food' [is] context-invariant, at least with respect to the contribution of ... 'nearby'. (Hawthorne 2006, 444)

As Hawthorne acknowledges, an ICDIR* in Brighton of Lepore's utterance in New York is not always felicitous; there are contexts in which such an ICDIR* can seem very wrong, though whether the ICDIR* is actually false is unclear. For example, if Hawthorne had just been asked whether there is any good Vietnamese food in Brighton, the ICDIR* 'Ernie said that a nearby restaurant has good Vietnamese food' would seem very wrong, though one might argue that it is only misleading, not strictly speaking false. On the other hand, there are clearly other contexts in which the ICDIR* would seem both proper and true. Consider, for example, the following dialogue:

Lepore (speaking to Hawthorne on Monday in New York City): A nearby restaurant has good Vietnamese food. It's just down the street from my apartment.

Stanley (speaking to Hawthorne on Friday in Brighton): How is Ernie liking his new neighborhood in New York? I bet he doesn't have to walk far to get good Vietnamese food.

Hawthorne (speaking to Stanley on Friday in Brighton): Yes, I spoke to him a few days ago and he said that a nearby restaurant has good Vietnamese food.

What Hawthorne shows, then, is that there are *some* pairs of contexts such that there is a true ICDIR* in one context of an utterance containing 'nearby' in the other context. Hawthorne concludes that it follows from C&L's test that 'nearby' is not context-sensitive. But, given Hawthorne's own characterization of C&L's test, one could not reasonably conclude that 'nearby' is not context-sensitive on the basis of the test unless we could determine that for *every* member of large and varied set of pairs of contexts, an utterance containing 'nearby' occurring in one of the contexts can always be reported with a true ICDIR* in the other context. Since Hawthorne does not consider a large and varied set of pairs of contexts, he fails to derive the absurd consequence that 'nearby' is not context-sensitive from the test.

This reply to Hawthorne may seem unsatisfactory, since it leaves unanswered an important question that Hawthorne's critique raises: does IT in fact deliver the absurd result that 'nearby' is not context-sensitive? If we *were* to consider a sufficiently large and diverse set of pairs of contexts, would it turn out that every pair of contexts in the set is such that when an utterance containing 'nearby' occurs in one member of the pair, an ICDIR* of that utterance is true in the other?

It is surprisingly difficult to find distinct contexts c and c' such that an utterance containing ‘nearby’ occurs in c and there is a *clearly false* ICDIR* of that utterance in c' . This is because it may always be objected that an apparently false ICDIR* is in fact just radically misleading, but not strictly speaking false. Consider, for example, the following dialogue:

Lepore (Monday in New York City, speaking to Hawthorne): There is a good Vietnamese restaurant nearby. It’s the Saigon Café, just down the street.

Stanley (Tuesday in Brighton, speaking to Hawthorne): I could do with some good Vietnamese food. But there probably isn’t a good Vietnamese restaurant for miles.

Hawthorne (Tuesday in Brighton, speaking to Stanley): Lepore said that there is a good Vietnamese restaurant nearby.

Hawthorne’s ICDIR* of Lepore’s utterance certainly sounds wrong, and one might conclude that the ICDIR* is simply false. If so, then IT clearly does not yield the absurd result that ‘nearby’ is not context-sensitive. It might be replied, however, that Hawthorne’s ICDIR* in the example above is only radically misleading, not strictly speaking false. It is reasonable for Stanley to presume that Hawthorne is trying to make a relevant conversational contribution. Since the issue at hand is whether there is a good Vietnamese restaurant near Hawthorne and Stanley, in Brighton, Hawthorne’s ICDIR* would not be relevant unless Lepore had been speaking about Vietnamese restaurants in Brighton. So it is reasonable for Stanley to conclude that ‘nearby’ in Hawthorne’s ICDIR* refers to somewhere near them. Hawthorne’s ICDIR* is therefore highly misleading, since it would lead a reasonable hearer to form a false belief about the content of what Lepore said. However, the reply goes, strictly speaking Hawthorne’s ICDIR* is true—‘nearby’ in Hawthorne’s report refers to the same location as ‘nearby’ in Lepore’s utterance.

The preceding case, then, is not a *clear* case of a pair of contexts in which an utterance containing ‘nearby’ occurs in one context, and a false ICDIR* of that utterance occurs in the other. However, it is also not a clear case of a pair of contexts in which an utterance containing ‘nearby’ occurs in one context, and a *true* ICDIR* of that utterance occurs in the other. It is simply not clear *what* is going on in this case: there are two reasonably plausible but incompatible ways of interpreting the case, and, as far as I can tell, there is no compelling reason to think that one of these interpretations is correct. This suffices to show that Hawthorne’s critique is not effective against IT—since there are some pairs of contexts in which the truth-value of the ICDIR* is unclear, IT does not deliver a clear verdict on ‘nearby’; so Hawthorne’s overall strategy, of arguing that C&L’s test delivers the absurd result that ‘nearby’ is not context-sensitive, fails against IT.

As we have seen, Hawthorne’s criticism does not fare well against IT. How does it fare against TT? In order to apply TT, we need a contextualist theory of ‘nearby’. Consider a ‘naïve’ contextualist theory of ‘nearby’, according to which ‘nearby’ refers to a location close to the speaker whenever it is uttered. This theory predicts that ‘nearby’ has distinct semantic values when Ernie says, ‘A nearby restaurant has good Vietnamese food’, in New York City, and when Hawthorne says, ‘Ernie said that a nearby restaurant has good Vietnamese food’, in Brighton. Given that this latter utterance of Hawthorne’s is a true ICDIR* of Ernie’s utterance, it follows from TT that the naïve theory is false. This, however, in no way invalidates TT, since the data Hawthorne presents should lead us to exactly this conclusion. Hawthorne agrees with this:

On a very natural view, ‘nearby’ is context-dependent but much more amenable to disquotational reporting than ‘I’. And there is an easy explanation of this: ‘Ernie is nearby’ is true relative to a contextually relevant location that is supplied by the context, but the *relevant location need not have anything to do with the location where the utterance is made.* (Hawthorne 2006, 445; Hawthorne’s italics)

It follows from TT and Hawthorne's data that the naïve theory is false. It does *not* follow from TT that 'nearby' is not context-sensitive. Since TT applied to Hawthorne's data leads to a plausible conclusion that Hawthorne agrees with, Hawthorne's data does not show that TT is a flawed test.

As I have argued, Hawthorne's criticism of C&L's test leaves two plausible clarifications of C&L's test, IT and TT, unscathed. Thus, the threat these tests pose to EC remains. In what follows I will present some different criticisms of IT and TT.

Two problems for the new tests

I will consider two problems that afflict TT and IT alike. These problems are serious, but perhaps not fatal. Here I will be content to draw out the problems, without considering every possible solution.

Consider the intuitive rationale for IT I provided above. According to that line of thinking, if there are contexts c and c' such that an utterance u of a sentence s containing e occurs in c and a true ICDIR* of u occurs in c' , then e has the same semantic value in c and c' . So, the thinking goes, if one considers a sufficiently large and varied set of such pairs of contexts, one has strong inductive support for the claim that e has the same semantic value in *every* context. The idea is that one can take any two possible contexts one likes, imagine an utterance of a sentence containing e in one context, note that an ICDIR* of that utterance in the other context would be true, and conclude that e has the same semantic value in both contexts. Repeating the procedure with more and more arbitrarily selected pairs of contexts would gradually build inductive support for the claim that e is not context-sensitive.

The error in this thinking is the tacit assumption that one can apply the procedure to any arbitrarily selected pair of contexts. Applying the procedure to a pair of contexts involves imagining an ICDIR* occurring in one of the contexts; so, the assumption entails that for any context, there is a possible situation in which an ICDIR* occurs in that context. This is true only if the occurrence of an ICDIR* cannot *shift* the context, so that the context in which the ICDIR* occurs is different from the context that would have been in place had the ICDIR* not occurred. Contextualists typically hold that the occurrence of kinds of utterances can result in a context-shift. For example, as we saw in chapter 2, DeRose (1995) holds that any utterance of ‘I don’t know that I am not a brain in a vat’ results in a shift from an ordinary context to a skeptical context in which ‘know’ expresses *h*-knowledge, a relation that is not expressed by ‘know’ in ordinary contexts. Thus, on DeRose’s view, one can never utter ‘I know that I am not a brain in a vat’ in an ordinary context, which explains why one can never utter it truly—it is false when uttered in a skeptical context, and it simply cannot be uttered in an ordinary context. If contextualists are right that the occurrence of certain types of utterance can shift the context, then one cannot assume that for any context *c* and utterance *u*, there is a possible situation in which *u* occurs in *c*. For *u* might be the kind of utterance that results in a shift from *c* to some different context *c*’. The reasoning that justified IT therefore depends on supposing that the occurrence of an ICDIR* cannot result in a context-shift. We have no reason to think this is so, however.

If an ICDIR* can result in a context-shift, then we cannot apply IT to any arbitrarily chosen pair of contexts, since some of these contexts might be such that an ICDIR* cannot occur in them. Therefore, unless it could be shown that an ICDIR* cannot shift the context, IT cannot provide strong inductive support for the claim that a certain expression is not context-sensitive.

Consider how a contextualist might respond to the attempt to use IT to establish the falsity of EC. In order to apply IT to EC, the invariantist would have to claim that after having considered a large and diverse set of pairs of contexts, it turns out that for every pair in the set, if a knowledge attribution occurs in one of the contexts in the pair, then there is a true ICDIR* of that knowledge attribution in the other context. Consequently, there is no pair of contexts in the set such that ‘know’ has distinct semantic values in the two contexts, and this provides strong inductive support for the claim that EC is false. The contextualist can reply that for many of the contexts in the set under consideration, there are no true ICDIR*s in those contexts, because there can be no ICDIR*s *at all* in those contexts. (Attempting to perform an ICDIR* in one of those contexts would simply result in a shift to a different context.) Consequently, the invariantist is wrong to conclude that for every pair of contexts in the set, ‘know’ has the same semantic value in both contexts.

TT faces a similar problem. Suppose a certain contextualist theory T of e says that e has distinct semantic values in two distinct contexts. To apply TT to T , we imagine an utterance containing e occurring in one of those contexts, and an ICDIR* of that utterance occurring in the other. If the ICDIR* is true, this shows that e has the same semantic value in both contexts, and consequently that T is false. For example, Cohen’s version of EC says that ‘know’ takes on distinct semantic values in c_1 , Smith’s context, and c_2 , Mary’s context. So we imagine Mary saying ‘Smith said that he knows that the plane stops in Chicago’ in c_2 . Since that ICDIR* is true, we conclude that ‘know’ has the same semantic value in c_1 and c_2 , and that Cohen’s version of EC is false. Cohen, however, could reply as follows: you cannot *really* imagine Mary’s ICDIR* occurring in c_2 , because performing the ICDIR* would shift the context from c_2 to some

other context c_3 . Sure, ‘know’ has the same semantic value in c_1 and c_3 ; that hardly shows that ‘know’ has the same semantic value in c_1 and c_2 .

As I have argued, IT and TT cannot be used in effective arguments against EC unless it could somehow be shown that ICDIR*s cannot result in context-shifts. Let’s suppose this problem is solved somehow. A second problem for IT and TT is that their justifications depend on denying speech act pluralism.

The rationales for IT and TT provided above each relied on the following assumption:

- (A) If an utterance of a sentence s containing e occurs in context c and there is a true ICDIR* of that utterance in a context $c' \neq c$, then e has the same semantic value in c and c' .

Why think that (A) is true? Well, assume that the ICDIR* in c' is reporting on the *semantic content* of the utterance occurring in c . Given this assumption, the ICDIR* is true only if the semantic content of s^* (the ICDIR*'s *that*-clause) in c' is identical to the semantic content of s in c . (Otherwise, the ICDIR* would be misreporting the utterance’s semantic content.) If e had distinct semantic values in c and c' , then the semantic content of s in c would be distinct from the semantic content of s^* in c' . So e must have the same semantic value in c and c' if the ICDIR* is to truly report the semantic content of the utterance.

An ICDIR*, however, need not report an utterance’s semantic content; an ICDIR* is a report of what a speaker asserted, and if speech act pluralism is correct, then a speaker asserts propositions other than the semantic content of her utterance. So an ICDIR* in c' of an utterance of a sentence s containing e in c , where $c \neq c'$, might correctly report that the speaker asserted a proposition q , where q is distinct from p , the semantic content of s in c . If the semantic content of s^* is q , and $p \neq q$, then given the manner in which s^* was obtained from s , it follows that e has

distinct semantic values in c and c' . Provided that speech act pluralism is correct, then, it does not follow from the fact that there is a true ICDIR* in c' of an utterance containing e in c , that e has the same semantic values in c and c' . Further, we saw in chapter 5 that there is quite compelling evidence for speech act pluralism. Thus, (A), on which both IT and TT depend for their justification, is at best unsupported and at worse plain false.

In light of this criticism, consider again the attempt to use TT against Cohen's version of EC. Let's grant that in c_2 Mary can truly say, 'Smith said that he knows that the plane stops in Chicago.' The invariantist would like to argue that 'know' in this ICDIR* expresses m -knowledge, the relation expressed by 'know' in ordinary contexts, since the semantic value of 'know' in Smith's utterance is m -knowledge, and the ICDIR* correctly reports what Smith said. The epistemic contextualist, however, could reply as follows: it is possible that in addition to asserting the semantic content of his utterance in c_1 —that he m -knows that the plane stops in Chicago—Smith also asserts that he h -knows that the plane stops in Chicago. If so, then Mary's ICDIR* *correctly* reports what Smith asserted, even though the semantic value of 'know' in Mary's ICDIR* is h -knowledge. Thus there is no inconsistency in claiming both that 'know' has distinct semantic values in c_1 and c_2 , and that Mary's ICDIR* of Smith's utterance is true.¹¹

I have presented two criticisms of IT and TT. The first is that IT and TT are both unsound if ICDIR*s are capable of shifting contexts. The second is that the rationales for IT and TT rely on (A), a dubious assumption. The tests are perhaps not dead in the water, since it might be argued that ICDIR*s cannot shift contexts, and new rationales for IT and TT that do not rely on (A) might be provided. As it stands, however, we should not rely on these tests in arguing against EC, or any other form of contextualism.

In chapter 5, I used intuitions about the truth-values of indirect speech reports (of the collective variety) to argue that one of the utterances in the Airport Case occurs in a non-normal context, and consequently that the New Airport Argument, which relies on the premise that Smith's and Mary's utterances both occur in normal contexts, is unsound. But now I am critiquing C&L's test for context-sensitivity (or, rather, two plausible clarifications of C&L's test; I will dispense with this qualification in what follows) on the grounds that it draws unwarranted conclusions from intuitions about the truth-values of indirect speech reports. One might worry, then, that my critique of C&L's test applies to my own use of indirect speech reports in chapter 5.

Fortunately, this worry is unfounded. My critique of C&L's test reduces to the fact that the test relies on two claims that are far from being clearly true, yet are left entirely undefended by C&L: first, that an indirect speech report cannot itself result in a context-shift; second, that if an indirect speech report correctly reports that *S* said that *p* with some utterance *u*, then *p* is *u*'s semantic content. The argument of chapter 5, however, is entirely independent of both of these claims, as the following outline of the argument makes clear:

- (P6) The New Airport Argument relies on the premise that Smith and Mary are both in normal contexts in the Airport Case.
- (P7) No false propositions are asserted in normal contexts.
- (P8) The indirect speech report, 'Smith claimed to know that the plane stops in Chicago but Mary said that Smith did not know that the plane stops in Chicago', is true.
- (P9) If this indirect speech report is true, then there is a proposition *p* such that Smith asserted *p* and Mary asserted not-*p*.

(C4) Therefore, if the indirect speech report is true, then either Smith or Mary asserted a false proposition.

(C5) So either Smith or Mary is in a non-normal context, and the New Airport Argument is unsound.

It is hard to see how any of the premises (P6)–(P9) could be relying, even implicitly, on the questionable claims upon which C&L’s test relies. The notion of semantic content is not involved in the argument at all, and the argument does not require assuming that the indirect speech report it focuses on occurs in the one of the contexts involved in the Airport Case. So my critique of C&L’s test does not apply to my own use of indirect speech reports to argue against the New Airport Argument.

Although I have argued that C&L’s use of indirect speech reports to argue against EC is unsuccessful, let me stress that the general strategy of using indirect speech reports to argue against EC continues to strike me as very promising indeed. The key to successfully implementing such a strategy is to undertake a thorough investigation of indirect speech reports: to determine the conditions (if such ever obtain) under which the use of an indirect speech report results in a context-shift, and to defend principles that tell us when we are permitted to draw conclusions about utterances’ semantic contents on the basis of indirect speech reports.

Deepening our understanding of indirect speech reports, in addition to being an independently laudable philosophical and linguistic pursuit, might ultimately lead to a decisive refutation of EC. This is a project I eventually hope to undertake—but that is work for a different context.

¹ Notice that in addition to the claim that a context-shifting argument can be provided for any sentence whatsoever, GEN also states that it follows from this that a context-shifting argument can be provided for any expression whatsoever. This secondary claim is neither completely clear, nor, once its meaning has been elucidated, obviously true. What does it mean for a context-shifting argument to be provided for an expression *e* that is not a sentence? It can’t mean that in different contexts we form differing judgments with respect to the truth and falsity of *e*, since

expressions like ‘my hand’ and ‘knows’ are not capable of being either true or false. On the other hand, the mere fact that we judge a sentence to be true when uttered in one context, false in another, says nothing about which of the sentence’s constituent expressions is responsible for these differing judgments. Presumably, what C&L have in mind is a case involving a sentence *s* containing *e*, a world *w*, and a pair of contexts *c* and *c'*, such that an utterance of *s* in *c* and *w* is intuitively true, and utterance of *s* in *c'* and *w* is intuitively false, and every expression in *s* other than *e* has the same semantic value in *c* and *w* that it has in *c'* and *w*. If it is appropriate to draw any conclusion about semantic context-sensitivity from such a case, it would have to be the conclusion that *e* is context-sensitive. Given this understanding of what it means to provide a context-shifting argument for an expression that is not a sentence, it does not obviously follow, from the claim that a context-shifting argument can be provided for any sentence whatsoever, that a context-shifting argument can be provided for any expression whatsoever. In order to provide a context-shifting argument for an expression *e* that is not a sentence, one must provide a sentence *s*, a world *w*, and contexts *c* and *c'* such that every constituent of *s* other than *e* has the same semantic value when uttered in *c* and *w* and when uttered in *c'* and *w*. Given these constraints, it is not obvious that for any *e*, one can always produce an *s*, *w*, *c*, and *c'* such that *s* is intuitively true when uttered in *c* and *w* and intuitively false when uttered in *c'* and *w*. Certainly, that such a sentence, world, and contexts can be produced does not follow from the first part of GEN, according to which there must be *some* world *w#* and *some* pair of contexts *c#* and *c##* such that *s* is intuitively true when uttered in *c#* and *w#* but intuitively false when uttered in *c##* and *w#*; for we have no guarantee that every expression in *s* other than *e* has the same semantic value when uttered in *c#* and *w#* as it does when uttered in *c##* and *w#*.

² In claiming that one would be speaking loosely if one were to appropriately deny (S5), I am not relying on an explicit test for distinguishing literal speech from mere loose use. I’m not convinced I need such a test, since it seems to me that competent speakers are quite good at distinguishing cases of loose use from cases of literal speech, and that if it is ever appropriate to deny (S5), this is so only because ‘bachelor’ is being used loosely. At any rate, if the absence of a test distinguishing loose use from literal speech really is a problem for me, then it is a problem that afflicts C&L’s collapse argument as well. In order for C&L to establish GEN, they must show that every English sentence is susceptible to a context-shifting argument. As C&L (2005, 42) grant, however, a context-shifting argument for a sentence *s* is one in which a *literal* use of *s* is intuitively true in one context, and a *literal* use of *s* is intuitively false in a different context. But C&L do not articulate a test distinguishing loose use from literal use. So, if their lack of such a test means that C&L are unreliable at discriminating loose use from literal speech, then there is reason not to accept the examples of context-shifting arguments that form the basis for their collapse argument.

³ Presumably, one would attribute the fact that *s* permits a CSAID to *e* when it is clear that the contents of all the words in *s* other than *e* remain constant across the contexts involved in the CSAID.

⁴ The arguments C&L deploy against radical contextualism also show why a different weakening of GEN will not allow them to carry out their argumentative strategy. It might be noted that (S5)-(S7) all express necessary truths; so one might argue that any sentence expressing a *contingent* proposition permits a CSAID. I don’t know if this claim is correct, but even if it is, it won’t help C&L. Arguably, if radical contextualism is true then it is necessarily true. So C&L would not be able to use the suggested weakening of GEN to argue that radical contextualism is internally inconsistent, since they could not assume that ‘radical contextualism is true’ is context-sensitive. Furthermore, they could not argue that radical contextualism makes communication impossible, since they would not be able to show that sentences expressing necessary truths cannot be communicated.

⁵ In addition to the test I focus on here, C&L offer two other tests for context-sensitivity in *Insensitive Semantics*. The first of these states, roughly, that context-sensitive expressions block collective descriptions—that is, if *v* is a context-sensitive verb phrase, the inference from ‘A *v*’s’ and ‘B *v*’s’ to ‘A and B *v*’ will typically be invalid (C&L 2005, 99-104). Stanley (2005, 49-51) offers a compelling critique of this. The central example is that if Jill and Mary are sisters we can move from ‘Jill loves her mother’ and ‘Mary loves her mother’ to ‘Each sister loves her mother’, but of course this does not show that ‘loves her mother’ is context-invariant. The other proposed test states, roughly, that context-sensitive expressions typically allow us to say thing like: ‘There is a false utterance of “I am hungry” even though I am hungry’ (C&L 2005, 104-112). As Aidan McGlynn points out on his blog *The Boundaries of Language* (<http://aidanmcglynn.blogspot.com/2008/01/testing-for-context-sensitivity-part-1.html>), this test appears to count *any* expression as context sensitive. However, it is not obvious how to reformulate the test in order to avoid this problem. In light of these criticisms, these tests strike me as far less compelling than the test involving speech reports I focus on here, and I will not discuss them further.

⁶ More troublingly, they argue that comparative adjectives like ‘tall’ and ‘flat’ also fail the test and so are not context-sensitive.

⁷Notice that ‘am’ also had to be replaced with ‘is’ in order to keep the sentence grammatical, so strictly speaking the procedure described above should be supplemented with something like: If the resulting sentence is ungrammatical, make whatever semantically innocuous adjustments are necessary to restore grammaticality.

⁸In the quoted passage I correct what appears to be a typo—the original passage concluded “that’s evidence that S is context insensitive,” which makes little sense, since (i) the test can be applied to sentences that are *obviously* context-sensitive, since they may contain indexicals in the basic set, and (ii) the conclusion of the test in its original form is about the context insensitivity of a sentence, but the test is supposed to help us determine whether *e*, a sentential constituent, is context-sensitive.

⁹It might even be suggested that the qualified nature of the conclusion of C&L’s test, “that’s evidence that *e* is context-insensitive,” suggests that they had something like an inductive procedure in mind all along.

¹⁰Notice that unlike C&L’s test as originally stated IT makes no reference to relevantly different contexts. This is important, since the test would clearly be inadequate if it required that the pairs of contexts being considered must be relevantly different. Relevant difference, recall, is determined by a particular contextualist theory of *e*, which specifies certain contexts in which *e* has distinct semantic values. The contextualist theory can be represented as a set of pairs of contexts, $C = \{\{c_1, c_2\}, \{c_2, c_3\}, \dots, \{c_{n-1}, c_n\}\}$, where for any pair of contexts in the set, it follows from the theory that *e* has distinct semantic values in those two contexts. Clearly, for any remotely plausible contextualist theory of *e*, there will be some pairs of contexts that are *not* members of *C*; again, no plausible contextualist theory of *e* will deny that there are some distinct contexts in which *e* has the same semantic value. If IT were restricted to pairs of contexts in *C*, then it could not form part of an adequate inductive procedure, since it could not be used to test pairs of contexts outside of *C*; in order to inductively confirm that *e* is not context-sensitive, however, one must be able to test any pair of contexts whatsoever.

¹¹It might be objected that it is rather far fetched to think that Smith asserts both that he *m*-knows and that he *h*-knows that the plane stops in Chicago. I think this misses the point, however. My purpose in exploring this possibility is not to provide a plausible story about what is going on in the Airport Case, but to provide a more concrete illustration of how (A) might turn out to be false.

WORKS CITED

- Bach, K. 2000. Quantification, Qualification, and Context: A Reply to Stanley and Szabo. *Mind and Language* 15: 262-283.
- _____. 2005. The Emperor's New 'Knows'. In Preyer and Peter, 51-89.
- Bezuindenhout, A. 2002. Truth Conditional Pragmatics. *Philosophical Perspectives* 16: 105-134.
- Black, T. 2005. Classical Invariantism, Relevance, and Warranted Assertability Maneuvers. *Philosophical Quarterly* 55: 328-336.
- Black, T. and P. Murphy. 2005. Avoiding the Dogmatic Commitments of Contextualism. *Grazer Philosophische Studien* 69: 165-182.
- Blome-Tillmann, M. 2008. The Indexicality of 'Knowledge'. *Philosophical Studies* 138: 29-53.
- BonJour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- _____. 2002. *Epistemology: Classic Problems and Contemporary Responses*. Lanham: Rowman and Littlefield.
- Brown, J. 2005a. Adapt or Die: The Death of Invariantism? *Philosophical Quarterly* 55: 263-314.
- _____. 2005b. Comparing Contextualism and Invariantism on the Correctness of the Contextualist Intuitions. *Grazer Philosophische Studien* 69: 71-99.
- Burge, T. 1986. Intellectual Norms and Foundations of Mind. *Journal of Philosophy* 83: 699-720.
- Cappelen, H. 2005. Pluralistic Skepticism: Advertisement for Speech Act Pluralism. *Philosophical Perspectives* 19: 15-39.
- Cappelen H. and E. LePore. 1997. On an Alleged Connection Between Indirect Speech and the Theory of Meaning. *Mind and Language* 12: 278-296.
- _____. 2005. *Insensitive Semantics*. Oxford: Blackwell Publishing.
- _____. 2006. Reply to Hawthorne. *Philosophy and Phenomenological Research* 73: 473-481.
- Carston, R. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell Publishing.
- Cohen, S. 1988. How to be a Fallibilist. *Philosophical Perspectives* 2: 91-123.

- _____. 1998. Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery. *Australasian Journal of Philosophy* 76: 289-306.
- _____. 1999. Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives* 13: 77-79.
- _____. 2002. Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research* 65: 309-329.
- _____. 2004. Knowledge, Assertion, and Practical Reasoning. *Philosophical Issues* 14: 482-491.
- Conee, E. 2005. Is Knowledge Contextual? Contextualism Contested. In Steup and Sosa, 47-56.
- DeRose, K. 1992. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research* 52: 913-929.
- _____. 1995. Solving the Skeptical Problem. *Philosophical Review* 104: 1-52. Reprinted in DeRose and Warfield, 183-219 (page references are to the reprinted article).
- _____. 1996. Knowledge, Assertion, and Lotteries. *Australasian Journal of Philosophy* 74: 568-581.
- _____. 1999. Introduction: Responding to Skepticism. In DeRose and Warfield, 1-24.
- _____. 2002. Assertion, Knowledge, and Context. *Philosophical Review* 111: 167-203.
- _____. 2004. Single Scoreboard Semantics. *Philosophical Studies* 119: 1-21.
- DeRose, K. and Warfield, T., eds. 1999. *Skepticism: A Contemporary Reader*. Oxford: Oxford University Press.
- Dretske, F. 1970. Epistemic Operators. *Journal of Philosophy* 67: 1007-1023. Reprinted in DeRose and Warfield, 131-144 (page references are to the reprinted article).
- _____. 2005. Is Knowledge Closed Under Known Entailment? The Case Against Closure. In Steup and Sosa, 13-26.
- Gettier, E. 1963. Is Justified True Belief Knowledge? *Analysis* 23: 121-123.
- Grice, P. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.

- _____. 2005. Is Knowledge Closed Under Known Entailment? The Case for Closure. In Steup and Sosa, 26-43.
- _____. 2006. Testing for Context-Dependence. *Philosophy and Phenomenological Research* 73: 443-450.
- Heller, M. 1999. The Proper Role for Contextualism in an Anti-Luck Epistemology. *Philosophical Perspectives* 13: 115-129.
- Kaplan, D. 1989. Demonstratives. In *Themes from Kaplan*, eds. J. Almog, J. Perry, and H. Wettstein. Oxford: Oxford University Press.
- King, J. 2008. Structured Propositions. In *The Stanford Encyclopedia of Philosophy*, ed. E. Zalta. Spring 2008 ed. URL=< <http://plato.stanford.edu/archives/spr2008/entries/propositions-structured/>>
- Kvanvig, J. 2007. Epistemic Closure Principles. *Philosophy Compass* 1: 256-267.
- Kyburg, H. 1961. *Probability and the Logic of Rational Belief*. Middletown: Wesleyan University Press.
- Lackey, J. 2007. Norms of Assertion. *Nous* 41: 594-626.
- Lewis, D. 1996. Elusive Knowledge. *Australasian Journal of Philosophy* 74: 549-567. Reprinted in DeRose and Warfield, 220-239 (page references are to the reprinted article).
- Moore, G. 1962. *Commonplace Book*. London: Allen and Unwin.
- Montminy, M. 2006. Semantic Content, Truth Conditions and Context. *Linguistics and Philosophy* 29: 1-26.
- Nelkin, D. 2000. The Lottery Paradox, Knowledge, and Rationality. *Philosophical Review* 109: 373-409.
- Neta, R. 2003a. Contextualism and the Problem of the External World. *Philosophy and Phenomenological Research* 66: 1-31.
- _____. 2003b. Skepticism, Contextualism, and Semantic Self-Knowledge. *Philosophy and Phenomenological Research* 67: 396-411.
- Nozick, R. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Preyer, G. and G. Peter, eds. 2005. *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford: Oxford University Press.

- Pritchard, D. 2005. Neo-Mooreanism versus Contextualism. *Grazer Philosophische Studien* 69: 1-24.
- Radford, C. 1966. Knowledge-by Examples. *Analysis* 27: 1-11.
- Recanati, F. 2002. The Limits of Expressibility. In *John Searle*, ed. B. Smith. Cambridge: Cambridge University Press.
- Rysiew, P. 2001. The Context-Sensitivity of Knowledge Attributions. *Nous* 35: 477-514.
- _____. 2008. Epistemic Contextualism. In *The Stanford Encyclopedia of Philosophy*, ed. E. Zalta. Spring 2008 ed. URL=<<http://plato.stanford.edu/archives/spr2008/entries/contextualism-epistemology/>>
- Sainsbury, M. 1995. *Paradoxes*. 2d ed. Cambridge, UK: Cambridge University Press.
- Schiffer, S. 1996. Contextualist Solutions to Scepticism. *Proceedings of the Aristotelian Society* 96: 317-333.
- Searle, J. 1980. The Background of Meaning. In *Speech Act Theory and Pragmatics*, eds. J. Searle, F. Kiefer, and M. Bierwisch. Dordrecht: Reidel.
- Soames, S. 2002. *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. Oxford: Oxford University Press.
- Sosa, E. 2000. Skepticism and Contextualism. *Philosophical Issues* 10: 1-18.
- _____. 2004. Relevant Alternatives, Contextualism Included. *Philosophical Studies* 119: 35-65.
- Sperber, D. and D. Wilson. 1995. *Relevance: Communication and Cognition*. 2d ed. Oxford: Blackwell Publishing.
- Stalnaker, R. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Steup, S. and E. Sosa, eds. 2005. *Contemporary Debates in Epistemology*. Oxford: Blackwell Publishing.
- Stanley, J. 2004. On the Linguistic Basis for Contextualism. *Philosophical Studies* 119: 119-146.
- _____. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Travis, C. 1996. Meaning's Role in Truth. *Mind* 100: 451-466.
- Unger, P. 1975. *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.

- Vogel, J. 1990. Are there Counterexamples to the Closure Principle? In *Doubting: Contemporary Perspectives on Skepticism*, eds. M. Roth and G. Ross. Dordrecht: Kluwer.
- _____. 1999. The New Relevant Alternatives Theory. *Philosophical Perspectives* 13: 155-180.
- Weiner, M. 2005. Must We Know What We Say? *Philosophical Review* 111: 227-251.
- Williamson, T. 1994. *Vagueness*. New York: Routledge.
- _____. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- _____. 2005. Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *Philosophical Quarterly* 55: 213-235.