

CONTENTS

Contributors 7

I Introduction: The Limits of Contextualism I
GERHARD PREYER AND GEORG PETER

Part I *Contextualism in Epistemology*

2 Contextualism and the New Linguistic Turn in Epistemology II
PETER LUDLOW

3 The Emperor's New 'Knows' 5I
KENT BACH

4 Knowledge, Context, and the Agent's Point of View 9I
TIMOTHY WILLIAMSON

5 What Shifts? Thresholds, Standards, or Alternatives? II5
JONATHAN SCHAFFER

6 Epistemic Modals in Context 13I
ANDY EGAN, JOHN HAWTHORNE,
AND BRIAN WEATHERSON

Part II *Compositionality, Meaning, and Context*

7 Literalism and Contextualism: Some Varieties 17I
FRANÇOIS RECANATI

8	A Tall Tale: In Defense of Semantic Minimalism and Speech Act Pluralism	197
	HERMAN CAPPELEN AND ERNIE LEPORE	
9	Semantics in Context	221
	JASON STANLEY	
10	Meaning before Truth	255
	PAUL M. PIETROSKI	
11	Compositionality and Context	303
	PETER PAGIN	
12	Presuppositions, Truth Values, and Expressing Propositions	349
	MICHAEL GLANZBERG	
	<i>Index</i>	397

I

Introduction: The Limits of Contextualism

GERHARD PREYER AND GEORG PETER

In contemporary epistemology, the thesis that epistemic vocabulary is context-sensitive has been adopted by a number of authors. Since the thesis is a semantic claim, evaluation of its truth has interested those working in philosophy of language. In particular, serious consideration of the thesis requires some account of when a linguistic construction is context-sensitive. This book will look at contextualism in epistemology and its linguistic underpinnings, and at related general issues in the philosophy of language. The fruitful interaction between empirical questions about language and philosophical issues in epistemology that occurs in the essays in this book can perhaps be viewed a considerably more empirically grounded return to the linguistic turn in epistemology offered by Wittgenstein, Austin, and Malcolm.

The motivations of epistemic contextualism, until recently, had chiefly to do with its supposedly enabling good responses to scepticism, showing where—at least a certain form of—sceptical arguments get their seeming strength, without actually endorsing (acontextually, anyway) the conclusions thereof (see for example the work of Keith DeRose, Stewart Cohen, David Lewis). Thinking about our use of epistemic terms, DeRose advertises contextualism as at least as much motivated by a sort of ‘ordinary language’. In his essay, Peter Ludlow uses the test case of contextualism to illustrate this new linguistic turn in epistemology, and raises some of the many subtle questions that we need to consider when we take up topics such as the nature of gradable predicates, the different forms that implicit arguments can take, and the way in which different definitions of contextualism will yield different predictions about bound variable anaphora in knowledge reports. He

concludes that linguistic resources, handled carefully, can be a valuable tool, but that they cannot be deployed in isolation.

Contextualists try to resolve sceptical paradoxes not by refuting sceptical arguments but by confining them to contexts in which far-fetched possibilities are raised. In so doing, they assume that the conflicting intuitions that generate these paradoxes relate to the truth conditions of knowledge and are not merely vacillating responses to sceptical considerations. Contextualists reject invariantism about knowledge attributions and claim that a given knowledge-ascribing sentence can express different propositions in different contexts, which implies that there are many knowledge relations, not just one. However, as Kent Bach argues in his essay, this thesis is not as dramatic as it sounds, for even if it were correct, those propositions themselves would not be context-bound. More importantly, the fact that it can vary from one context to another how strictly we apply 'know' does not require a contextualist explanation. It also does not require dubious warranted assertibility arguments (WAMs), whose use DeRose imputes to invariantism. The fact that people use words with varying degrees of strictness and looseness does not show that the words themselves have semantic contents that come in various degrees. It could well be, Bach suggests, that with 'know' we often attribute knowledge to people who do not have it and often resist attributing it to people who do. Sometimes we are extra cautious, and sometimes we are even taken in by seductive sceptical arguments. Either way, we cannot attribute knowledge to someone, even if he has it, when he believes something on grounds that leave us with doubts or worries about the truth of the proposition in question. We sometimes demand more from knowledge than it requires. Bach argues that contextualism does not really come to grips with scepticism or with the sceptical version of invariantism, according to which knowledge requires the highest degree of evidence, justification, and conviction.

Scepticism denies that we have knowledge by ordinary standards, and sceptical invariantism does not use WAMs to explain why we casually make the knowledge attributions that we do. In any case, these sceptical views are implausible on independent grounds. Much more plausible is a moderate, non-sceptical version of invariantism, which also can interpret contextualist data without resorting to WAMs. What vary in contexts where special concerns arise, whether sceptical or practical, are not the truth conditions of knowledge attributions but the knowledge attributions people are prepared to make. It is not the standards for the truth of knowledge attributions that go up but the attributor's threshold of confidence regarding the relevant proposition.

One element of Bach's defence of moderate invariantism is the answer to a question raised by Lewis's version of contextualism: how can a possibility that is ignored be properly ignored? He suggests that this is so to the extent that the cognitive processes whereby beliefs are formed and sustained are sensitive to realistic counterpossibilities (so-called relevant alternatives). The occurrence of the thought of a possibility contrary to a tempting proposition gives one prima-facie reason to take it seriously. And the fact that such a possibility does not come to mind is (defeasible) evidence for its irrelevance. But this fact shows evidently that one cannot explicitly consider it, since to consider it would bring that counterpossibility to mind.

Contextualism is a mild form of relativism about the truth of sentences. There is a standard form of contextualist strategy for explaining the appeal of sceptical arguments. While vagueness is not merely a case of context-dependence, it does appear to be highly conducive to context-dependence. In his essay, however, Timothy Williamson shows that context-dependence in representation causes its own problems in the retention and transmission of information, especially when language users are unaware of the context-dependence. Moreover, a contextualist treatment of certain problems of practical reason is implausible, because the agent's context seems to have primacy over the speaker's context in practical respects. This primacy is incompatible with contextualism about the relevant terms. The case of practical reasoning suggests a way in which vagueness need not induce context-dependence. Williamson draws an analogy between the case of practical reasoning and the case of epistemic appraisal to show how the vagueness of epistemological vocabulary need not make it context-dependent. The analogy is then argued to be more than an analogy, because there is a practical aspect to epistemic appraisal with respect to the formation and retention of beliefs. Therefore, Williamson concludes, something is wrong with epistemological contextualism; he suggests an alternative explanation.

Contextualists speak of the semantic value of knowledge ascriptions as somehow shifting with context. But what is it that shifts? What is the parameter that shifts with the context? What epistemic gear do the wheels of context turn? In his essay, Jonathan Schaffer considers three possible answers. What shifts might be: (T) the threshold of justification (Cohen), (S) the standard of epistemic position (DeRose), or (A) the set of epistemic alternatives (Lewis). He assesses these three answers in the light of four desiderata. The parameter of shift must be: (D₁) linguistically plausible, (D₂) predictively adequate, (D₃) in accord with contextualist resolutions of scepticism, and (D₄) connected to our practices of inquiry. He argues that

(A) fares best, by these desiderata. Both (T) and (S) fail all of (D₁)–(D₄) while (A) satisfies them all. Thus he concludes that what shifts is the set of epistemic alternatives. A very simple contextualistic treatment of a sentence containing an epistemic modal, for example, *a might be F*, is that it is true iff for all the contextually salient community knows, *a* is *F*. It is widely agreed that the simple theory will not work in some cases, but the counterexamples produced so far seem to be amenable to a more complicated contextualist theory. In their essay, Andy Egan, John Hawthorne, and Brian Weatherson argue, however, that no contextualist theory can capture the evaluations speakers naturally make of sentences containing epistemic modals. If we want to respect these evaluations, our best option is a *relativist* theory of epistemic modals. On a relativist theory, an utterance of *a might be F* can be true relative to one context of evaluation and false relative to another. They argue that such a theory does better than any rival approach at capturing all the behaviour of epistemic modals.

The contemporary debate on contextualism shows more than *one* view on epistemology and the philosophy of language. There are rather *two* views dealing with the role of context, both called by the same name. One clear relationship is that, if contextualism about language is true, the certain semantic arguments against contextualism in epistemology will be disarmed. So even if the essays by François Recanati and Herman Cappelen and Ernie Lepore are closely related to the semantic underpinnings of contextualism in epistemology, their central claims are within the philosophy of language.

Two traditions in the philosophy of language and semantics take effect in our understanding of language and claim to give us an answer to the question: What is the basic notion of semantic content (properties)? According to literalism, we may legitimately ascribe truth-conditional content to natural language *sentences*, quite independent of what the speaker who utters this sentence means. For the contrasting view, namely contextualism, *speech acts* are the instances of content. Therefore it can be concluded that only in the context of a speech act does a sentence express a determinate content. It follows that the same sentence may express different contents in different contexts. Context-sensitivity must be generalized, as François Recanati claims to show in his essay. First, he looks at the historical development of literalism. The extent of context-sensitivity in natural language was progressively acknowledged. Where does this tendency ultimately lead us? Recanati's answer is to contextualism. Secondly, he describes the steps which can lead from a critique of the dominant literalist position (minimalism) to contextualism. In the last sections he presents several possible arguments in favour of

contextualism, corresponding to three distinct versions of contextualism. According to the first version, individual words have determinate contents but semantic composition requires going beyond that content: to determine the content of *complex* expressions (e.g. sentences) we must creatively enrich or otherwise adjust the meaning of individual words—and that cannot be done without appealing to speaker’s meaning. In the second version it is not just semantic composition which requires adjustment and modulation of word meaning. Individual word meanings themselves could not go directly into the interpretation. They are either too abstract and schematic—in such a way that elaboration or fleshing out is needed to reach a determinate content—or they are too rich and must undergo ‘feature-cancellation’, or some other screening process through which some aspects will be backgrounded and others focused on. According to the third, even more radical version, the content which a word contributes is contextually constructed, but the construction can proceed without the help of conventional, context-independent word meanings. Recanati’s turn is that the contextual sense which an expression assumes on a particular occasion of use can be computed directly on the basis of the contextual senses which that expression had on previous occasions of use.

Herman Cappelen and Ernie Lepore take a different view. The context-sensitive expressions which semantic minimalism recognizes are not only obvious, they also pass certain tests for context-sensitivity. Beyond fixing the semantic value of these obviously context-sensitive expressions, contexts of utterance have no effect on the propositions semantically expressed by the sentences we use. Cappelen and Lepore outline a defence of semantic minimalism and speech act pluralism against two common objections. In particular, it is alleged by some to follow from semantic minimalism that comparative adjectives are context-insensitive, and it also has been objected to minimalism that it postulates contents that are explanatorily idle, that play no role in an account of communication. Cappelen and Lepore defend semantic minimalism against the first objection not, as we might expect, by denying that implication, but by endorsing it. They go on to address the second objection and end with a reversal, that is, they argue that, although on their account meaning is not explanatorily idle, those who deny minimalism wind up with an account that is. In particular, they show that Recanati’s view fails to satisfy his own availability principle.

The essays of Jason Stanley, Paul M. Pietroski, and Peter Pagin connect issues of context-dependence with fundamental issues in the philosophy of language: meaning and truth, and compositionality. In his essay, Jason

Stanley continues his recent project of defending the view that the intuitive truth conditions of an utterance are the result of a compositional semantic process. In particular, he argues that this view is compatible with the context-dependency of what is expressed by an utterance. Furthermore Stanley argues that many of those who proceed otherwise are unfairly saddling the position with strange theoretical commitments. Part of his goal in this essay is to elucidate the target concept of the intuitive truth conditions of an utterance, in order to better elucidate the actual commitments of the view that intuitive truth conditions are due to semantics. Along the way, he discusses the proper attitude to take when some particularly troublesome data on the semantics/pragmatics divide, such as deferred reference, are concerned.

It is widely held that, if meaning is compositional, then the meaning of a declarative sentence *S* determines a truth condition for *S*, at least relative to a contextually determined choice of semantic values for any context-sensitive expressions in *S*. There has been a lot of debate about what conclusions we should draw from this conditional. But following Chomsky and a few others in his essay, Paul M. Pietroski asserts that the conditional is false: the meaning of a sentence *S* is pretty clearly determined by the meanings of the relevant constituents. But sentences may not even have (context-sensitive) truth conditions. And as many critics of the Davidsonian programme have noted, available evidence—across a wide range of examples—suggests that truth conditions are not compositionally determined. Indeed, we should be suspicious of the very idea that in natural language names denote things that satisfy predicates. Once we make this idea clear enough to serve as the basis for a theory of meaning that can also be a theory of understanding, the idea faces huge difficulties. And we should not be impressed by the usual philosophical motivations for a truth-theoretic conception of linguistic meaning because of their descent from considerations of radical interpretation or semantic externalism. The real question is whether we can articulate an alternative conception that is more descriptively adequate while retaining the explanatory virtues of the Davidsonian programme. Leading off from remarks by Chomsky, Pietroski argues in the direction of such a conception. He shows that such a conception is required if we want to have any real account of how lexical meanings and compositional principles interact to yield the compositional meanings which are characteristic for natural language.

How is the concept of compositionality to be extended from context-invariant to context-dependent meaning? And how might the compositionality of natural language conflict with context-dependence? Several new

distinctions are needed, including a distinction between a weaker (e-) and a stronger (ec-) concept of compositionality for context-dependent meaning. In his essay, Peter Pagin investigates the relations between the various notions. A claim by Jerry Fodor that there is a general conflict between context-dependence and compositionality is considered. There is in fact a possible conflict between ec-compositionality and context-dependence, but not of the kind Fodor suggests. It involves so-called unarticulated constituents, in John Perry's sense. Because of this phenomenon, some semantic accounts might have a variation in the meaning of a complex expression between contexts without any corresponding variation in the meaning of its syntactic parts. The conflict can be resolved in several ways. One way is to make the unarticulated context-dependence explicit only in the meta-language, which turns it into an unarticulated constituent account. A recent argument by Jason Stanley against such accounts is discussed. According to him, certain readings of English sentences are unavailable in these theories because they involve a binding of contextual variables. After considering a reply to Stanley by Recanati, Pagin presents an outline of a fully compositional theory—of the unarticulated constituent variety—which does deliver these readings.

Michael Glanzberg's essay addresses a more general question: what is the relation between the philosophical notion of expressing a proposition and the linguistic notion of presupposition? First, he offers an analysis of the philosophical notion of failing of expressing a proposition. Secondly, he gives an account of how failure to express a proposition may be recognized in natural language, by way of some discourse-based diagnostics which are more robust and reliable than simple truth-value judgements. Thirdly, he attempts to document what gives rise to the phenomenon of expression failure. This part of the essay involves some detailed investigation of presupposition. Glanzberg shows that some presupposition failures lead to expression failure, but some do not. He goes on to offer an analysis of elementary presuppositions which explains why presuppositions fall into these two categories. This involves a close examination of some important cases of presupposition: clefts, factive verbs, and presuppositions generated by conventional implicatures such as those of 'too' and 'even'. He ends with some speculation about the relation between conventional implicature and presupposition.

This project was initiated by Protosociology, J. W. Goethe-University, Frankfurt am Main, Germany. We would like to express our thanks to the contributors and to Ernie Lepore and Peter Momtchiloff who encouraged us to undertake the project.

*A Tall Tale
In Defense of Semantic Minimalism and
Speech Act Pluralism*

HERMAN CAPPELEN AND ERNIE LEPORE

In *Insensitive Semantics* (2005), we argue for two theses—Semantic Minimalism and Speech Act Pluralism. In this chapter, we outline our defense against two objections often raised against Semantic Minimalism. We begin with five stage-setting sections. These lead to the first objection, namely, that it might follow from our view that comparative adjectives are context *insensitive*. We defend our view against that objection (not, as you might expect, by denying that implication, but by endorsing it). Having done so, we address a second objection, namely, that Semantic Minimalism makes it difficult to see what role semantic content plays in communicative exchanges. We respond and end with a reversal, that is, we argue that even though the second objection fails against us, it works against those who raise the objection. In particular, we show that our critics, in particular, Carston (2002) and Recanati (2004), end up with a notion of communicated content that fails various tests for psychological reality.

Stage Setting 1: Semantic Minimalism

Three features of Semantic Minimalism are important in the context of this chapter (all elaborated on in *Insensitive Semantics*):

- (a) The most salient feature of Semantic Minimalism is that it recognizes few context sensitive expressions, and hence, acknowledges a very

limited effect of the context of utterance on the semantic content of an utterance. The only context sensitive expressions are the completely obvious ones ('I', 'here', 'now', 'that', etc., essentially those Kaplan lists in 'Demonstratives' (1989: 489)). These also pass certain tests for context sensitivity we spell out below.

- (b) It follows that all semantic context sensitivity is grammatically (i.e. syntactically or morphemically) triggered.
- (c) Beyond fixing the semantic value of these obviously context sensitive expressions, the context of utterance has no effect on the proposition semantically expressed or the semantic truth conditions. In this sense, the semantic content of a sentence S is that proposition that all utterances of S express (when we adjust for or keep stable the semantic values of the obvious context sensitive expressions in S).

Some illustrations: keeping tense fixed,¹ any utterance of (1)

- (1) Rudolf is a reindeer.

is true just in case Rudolf is a reindeer, and expresses the proposition that Rudolf is a reindeer.² Any utterance of (2)

- (2) Rudolf has a red nose.

is true just in case Rudolf has a red nose, and expresses the proposition that Rudolf has a red nose. Any utterance of (3)

- (3) Rudolf is happy.

is true just in case Rudolf is happy, and expresses the proposition that Rudolf is happy. Any utterance of (4)

- (4) Rudolf has had breakfast.

is true just in case Rudolf has had breakfast, and expresses the proposition that Rudolf has had breakfast. Any utterance of (5)

- (5) Rudolf doesn't know that penguins eat fish.

is true just in case Rudolf doesn't know that penguins eat fish and expresses the proposition that Rudolf doesn't know that penguins eat fish.

¹ As we will throughout this chapter.

² Semantic Minimalism need not take a stand on whether semantic content is a proposition, or truth conditions or what have you. Throughout we try to remain neutral by couching the issues both in terms of truth conditions and in terms of propositions.

If you find it surprising that we are writing a chapter or (worse) a book defending conclusions so obvious, we have a great deal of sympathy. The problem is that a wide range of our contemporary colleagues rejects these views. (It's probably no exaggeration to say that our views about (1)–(5) are now held only by a small minority of philosophers, at least among those who have thought about the surrounding issues.³) In our book, we rebut these influential objections; here we want to elaborate on some implications of the view defended.

Stage Setting 2: Speech Act Pluralism

Here's one way to summarize Speech Act Pluralism:

No one thing is said (or asserted, or claimed or...) by any utterance: rather, indefinitely many propositions are said, asserted, claimed, or stated. What is said (asserted, claimed, etc.) depends on a wide range of factors other than the proposition semantically expressed. It depends on a potential infinitude 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.

It follows from this view that an utterance can assert propositions not even (logically) implied by the proposition semantically expressed. Nothing even prevents an utterance from asserting (saying, claiming, etc.) propositions incompatible with the proposition semantically expressed by that utterance.

From this it further follows that if you want to use intuitions about speech act content to fix semantic content, you must be extremely careful. It can be done, but it's a subtle and an easily corrupted process.⁴

These points are connected to our defense of Semantic Minimalism because one underlying assumption in many anti-minimalist arguments is the idea that semantic content has to be closely connected to speech act content. If Speech Act Pluralism is correct, then no such close connection exists, and so this requirement is revealed to be a philosophical prejudice. (Another way to see the connection is this: If there really were (or had to be) a close connection between speech act content and semantic content, then all the data we think support Speech Act Pluralism would also serve to undermine Semantic Minimalism.)

³ See ch. 2 of *Insensitive Semantics* for extensive discussion of this point.

⁴ For some instructions on how to proceed, see ch. 7 of *Insensitive Semantics*.

At this initial stage it's worth highlighting one more aspect of Speech Act Pluralism that has both wide-ranging implications and sets our view apart from (all?) other contemporary accounts of context sensitivity. We don't think everything speakers say by uttering a sentence in a context *C* is determined by features of *C*. The speaker's intentions, facts about the audience, the place and time of utterance, background knowledge salient in *C*, previous conversations salient in *C*, etc., are not even together sufficient to fix what the speaker said. According to Speech Act Pluralism, a theory of speech act content has to take into account the context of those who say or think about what the speaker said, that is, the context of those who report on what's said by the utterance can, in part, determine what was said by that utterance. (As far as we can tell, we are on our own defending this view; see Cappelen and Lepore, 1997.)

Stage Setting 3: Opponents

We have many opponents; indeed, it often feels as if we have only opponents. What our opponents have in common is a commitment to some form of semantic contextualism. Semantic contextualists, as we mark them, posit more semantic context sensitivity than is generally recognized. Sometimes their motives are opportunistic: for example, they claim they can solve Sorities, Liar, Skeptical, Moral and Fregean puzzles/paradoxes by positing that vague, semantic, knowledge, moral and psychological attributions are semantically context sensitive. Sometimes they posit context sensitivity because they believe themselves to have uncovered more of it than linguists/philosophers have so far recognized. Sometimes they conclude that entire semantic programs collapse under the weight of their discoveries; sometimes they are more modest, concluding only that their contributions are to the general project of semantics for natural language—namely, modest extensions to the already recognized indexicals and demonstratives. No matter how ambitious or modest their motivations, we have come to the same conclusion: they are all wrong; none of the contextualist candidates are semantically context sensitive.

We have argued against contextualism with a variety of dialectical strategies: One of our favorite argumentative strategies is to present direct and simple tests for context sensitivity; and to show that traditionally recognized context sensitive expressions pass these tests with flying colors, while context-

ualist candidates all fail them. In order to present the first objection to Semantic Minimalism, we'll briefly rehearse a couple of these tests (both discussed at greater length in chapter 7 of *Insensitive Semantics*).

Stage Setting 4: Test 1—Context Sensitive Expressions Block Inter-Contextual Disquotational Indirect Reports

Take an utterance u of S in C . Let C' be a context relevantly different from C (i.e. different according to the standards significant according to contextualists about S). If there's a true disquotational indirect report of u in C' , then that's evidence S is context insensitive. So, take an obviously context sensitive expression, for example, the first person pronoun 'I' and its utterance in the sentence 'I went to Ottawa' made by Sarah-Jane. If Rich tries to report what Sarah-Jane said with 'Sarah-Jane said that I went to Ottawa', his report is false because the expression 'I' fails to pick out what it picked out in Sarah-Jane's mouth. The presence of 'I' in the disquotational report figures prominently in an explanation of why the report is false.

It's (almost) a matter of definition that context sensitive expressions tend to block inter-contextual disquotational indirect reports. The reason why is obvious: e is context sensitive only if e shifts semantic value between relevantly different contexts of utterance. It's obvious that all the traditionally recognized context sensitive expressions ('he', 'now', 'that', 'you', etc.) block inter-contextual disquotational indirect reports.

Stage Setting 5: Test 2—Context Sensitive Expressions Block Collective Descriptions

Here's another test applied to verbs first. If a verb phrase v is context sensitive (i.e. if it changes its semantic value from one context of use to another), then on the basis of merely knowing that there are two contexts of utterance in which 'A v -s' and 'B v -s' are true respectively, we *cannot* automatically infer that there is a context in which ' v ' can be used to describe what A and B have both done.

In short, from there being contexts of utterance in which 'A v -s' and 'B v -s' are true it doesn't *follow* that there is a true utterance of 'A and B both v .' This

is because the semantic value of ‘v’ in the previous collective sentence is determined in one context, and we have no guarantee that that semantic value, whatever it is, ‘captures’ (whatever that means) the semantic values of ‘v’ in those contexts of utterance where they were used alone.

On the other hand, *if* for a range of true utterances of the form ‘A v-s’ and ‘B v-s’ we obviously *can* describe what they all have in common by *using* ‘v’ (i.e. by using ‘A and B v’), then that’s evidence in favor of the view that ‘v’ in these different utterances has the same semantic content, and hence is not context sensitive. A parallel point extends to singular terms.

If an (unambiguous) singular term N is context *insensitive* and there’s a range of true utterances of the form ‘N is F’ and ‘N is G’, then we, for example, in *this* context, can truly utter ‘N is F and G’. Similarly, if N is context sensitive, we shouldn’t be able to do this. As an illustration consider the context sensitive ‘yesterday’: Suppose we know of two contexts in which ‘Yesterday John left’ and ‘Yesterday Bill left’ are true respectively (though we don’t know the days of these contexts). It doesn’t follow that there is a context in which ‘Yesterday John and Bill left’ is true.

Again, all traditionally recognized context sensitive expressions pass this test of collectivity.

There are other tests for context sensitivity; one of our favorites we call the Inter-Contextual Disquotational Test (different from Test 1 above). In discussing this test we distinguish between two kinds of context shifting arguments, Real and Impoverished, arguing that only the former identifies context sensitive expressions. (Context shifting arguments involve an appeal to speaker intuitions about distinct utterances of a single unambiguous sentence shifting in truth-value, or in proposition expressed, or in what’s said.)

The Inter-Contextual Disquotation Test was our first and we feel a sentimental attachment to it. However, audiences tend to find it a bit confusing, so we’ll leave it alone for now and direct those interested to our published work (Cappelen and Lepore, 2003). Instead, we’ll take our two tests involving indirect reporting and collectivity and turn to what most contextualists take to be a fundamental flaw in our position.

First Objection: ‘Tall’ Is Self-Evidently Context Sensitive

We have argued that the contextualist candidates fail the various tests for context sensitivity. This applies to ‘know’, ‘good’, ‘red’, quantifier words, and

so on. A standard reply is that there must be something wrong with our reasoning since words self-evidently context sensitive also seem to fail our tests: for example, comparative adjectives like ‘tall’.

So, for example, look at our first test. Suppose A utters in a context C ‘Rudolf is tall’. Suppose that in C the contextually salient comparison class consists of giraffes. According to contextualists, the proposition semantically expressed by A’s utterance is *that Rudolf is tall for a giraffe*. This result is rendered possible because ‘is tall’ is alleged to be context sensitive. But look at our tests: we take it as obvious that anyone reporting A’s utterance can accurately utter ‘A said that Rudolf is tall’ and this is so regardless of the context the reporter happens to find herself in, that is, even if the context of the report and the context of the reported utterance are relevantly different, that is, even if giraffes are not particularly salient in the context of the report. The reporter might not know that Rudolf is a giraffe; she might be unsure what kind of animal Rudolf is; or suspect he is a reindeer. The point is this: If the context of the first utterance and the context of the second utterance are relevantly dissimilar, then this report ought to be impossible—if ‘tall’ really is context sensitive.

Now turn to the second test. Take distinct utterances of ‘Mount Everest is tall’ and ‘Kobe Bryant is tall’ and ‘The Empire State Building is tall’. Suppose in the first context, mountains are salient, in the second NBA players are, and in the third skyscrapers are. Suppose you are collecting these utterances into a context in which mountains and basketball players and skyscrapers are *not* (particularly) salient. Then any utterance of ‘Mount Everest, Kobe Bryant, and the Empire State Building are all tall’ (or ‘Mount Everest is tall, and Kobe Bryant and the Empire State Building are too’—an appeal to a fourth test involving VP deletion; cf. *Insensitive Semantics*, ch 7) should be false, on the assumption that ‘tall’ is context sensitive. It’s our intuition, however, that there *are* contexts in which such utterances can be true; it’s hard to see how that could be so if ‘tall’ isn’t taking as its semantic value something the original utterances have in common.

Contextualists of every flavor have mocked, ridiculed, snickered, flat out laughed, and even worse, completely ignored our views because of these results. The current attitude seems to be that any argument that leads to the view that ‘tall’ (or any other comparative adjective) is not semantically context sensitive must be seriously flawed.

There are at least three responses to our arguments:

- (a) Our tests for semantic context sensitivity are no good.

- (b) Comparative adjectives do pass our tests, but for one reason or another, we can't hear their uses as passing these tests.
- (c) Or, one might say: That's right. These words fail the tests and they are context *insensitive*—contrary to what we all once thought.

We have considered and replied to the first two options elsewhere (cf. Cappelen and Lepore, 2003). Here we would like to try something bolder: we'd like to run with option (c). This requires investigating what others have thought of as the absolute absurdity of Semantic Minimalism, that comparative adjectives are (semantically) context insensitive.

To this end, we'll tease out our critics' argument; and try to establish it has nothing to do with semantics but rather reflects a metaphysical concern—one we do not think semanticists have to address. We present the objection in three stages, only the third of which will require an extended answer (though it is important to see how that stage is different from the first two).

First Stage Objection to (c): Dismissive Incredulous Stare

The objection to the view that 'tall' is context insensitive typically starts out with the kind of stare Lewis characterized as incredulous. This stare is typically accompanied by a dismissive utterance of something along the lines of:

Are you crazy! Of course, there can be both true and false utterances of (6):

- (6) Osama Bin Laden is tall.

If in one context the topic of discussion is the heights of NBA players, your utterance will be (taken to be) false; and if in another, the topic is the heights of Saudi Arabians, your utterance will be (taken to be) true.

Reply to First Stage

If you followed our brief introduction of Semantic Minimalism and Speech Act Pluralism above, you'll immediately understand that this reply can be pushed only by someone who does not understand our view. It is based on nothing but confusion. It should be obvious that our Speech Act Pluralism can accommodate the same data. In one context, the utterance *says* something true and in another an utterance of the same sentence *says* something false. But intuitions about the speech act content of these distinct utterances are *not*

reliable guides to the semantic content of (6) or even its utterances, and so intuitions about the former need not be a good guide for conclusions about the latter.

Second Stage of Objection to (c): Honest Request for Further Elaboration

Opponents who bypass the deeply confused first stage typically move on to a second stage of confusion. They ask, Well, what is it to be tall *simpliciter*? That is, what is it to satisfy the semantic truth conditions of 'A is tall'? If it is not to be tall for an X, or according to some standard, what then is it?

Reply to Second Stage

Our quick, and we think completely satisfactory, reply is given by (6_{TC}) and (6_P):

- (6_{TC}) 'Osama Bin Laden is tall' is true iff Osama Bin Laden is tall.
- (6_P) 'Osama Bin Laden is tall' semantically expresses the proposition *that Osama Bin Laden is tall*.

Here is our problem: We think this is a conclusive reply, but our opponents insist on further elaboration. They move on to the third stage of the objection.

Third Stage of Objection to (c): Confused Demand for Further Elaboration

Faced with (6_{TC}) and (6_P) our opponents tend to react with something like this:

(6_{TC}) and (6_P) just aren't enough. I can't take this theory seriously unless you tell me more about what the right-hand side of those biconditionals mean (or require, or demand or . . .). You just don't have a semantic theory unless you say more. If you can't tell me what it is to be tall, then you don't have a semantic theory.

Reply to Third Stage

We've presented this third stage of the objection so that it both reflects innumerable conversations we have had about this topic, but (we hope)

also reflects how unreasonable the demand is. We really don't think we, qua semanticists, are required to respond to this challenge. To demand that semanticists tell you what it is to *be tall* is to start down a most slippery slope. For example, why not also require that semanticists tell us what it takes, or is, to *be tall for a man*? Is that something semanticists are supposed to explore? Or take the word 'change'. Are semanticists required to reveal what the property of *change* is in order to do their job? Or what it is to be funny in order to deal with the semantics of the word 'funny'?

Though we take the answers to these various questions, qua semanticists, to be quite obvious, we also realize that sticking to our position is almost impossible (certainly unrewarding) since all the people that we like to talk to about these issues seem to lose interest if we don't elaborate. So, partly for selfish reasons (we don't want people to ignore us), partly out of the goodness of our hearts (we seek philosophical harmony), we'll engage in a little bit of metaphysics. We do this, however, filled with resentment and, ultimately, with the goal of getting our opponents to realize how absurd it is to require that we respond to their challenge.

Here goes. Think about dancing: Some people dance by stepping, some crawl around the floor (like Martha Graham), some have music, some don't have music, some jump in the air, some wave their arms, some hold on to other people, some are alone, some slide on ice, some fly in the air, etc. What do all these activities have in common in virtue of which they are all dancing? This is certainly not our area of expertise but suppose the dance metaphysicians will inform us that to dance is to move in some way *W*, where *W* is what all those different events of dancing have in common. There can be different accounts of *W*, and as far as we can tell both Semantic Minimalism and Semantic Contextualism are compatible with each and every one of them.

Or, think about eating. Some people eat sandwiches, some soup, some apples, some eat in Norway, some in the East Village of New York City, some eat with a spoon, some with their fingers. More generally, there are many things to eat, many places to do it, and many ways to eat. Any event of eating is of a specific thing, in some way, in some location. What is this property of eating? Well, isn't the simplest answer something along the lines of: to engage in the kind of activity that all these different events have in common, that is, what eating soup, apples, sandwiches, with finger, spoons, in Norway or New York, etc., have in common. Again, we're not specialists, but whatever they all have in common, that's what the activity of eating is. Notice, Semantic Minimalism and Semantic Contextualism are compatible with any answer

to these questions. Neither the former nor the latter need take a stand on what eating is.

Now think about funny things: There are funny people, funny jokes, funny paintings, funny movements, etc. People who are funny can be so by moving around funny, by saying funny things, by writing funny, etc. The expression ‘funny’ presumably has as its semantic value whatever all these things have in common. Here are some conjectures about this property. It might be dispositional: for an object to be funny is for it to trigger a certain reaction in an audience. Whether an act is funny might depend on the context in which it is performed (e.g. the interests, expectations, etc., of the salient audience). Any such account of the semantic value of ‘funny’ is compatible with Semantic Minimalism and with Semantic Contextualism.⁵

Finally, turn to the property of being tall. We suppose that to figure out what tallness is, you proceed much as in these earlier cases. Engage in a little bit of tallness metaphysics. Consider, for example, the Empire State Building, Mount Everest, and Kobe Bryant. Ask, what, if anything, do they all have in common? Naturally, one answer is that they are all tall. If that’s so, and it is, then it triggers the following metaphysical question. What is it in virtue of which these three objects are all tall? Or, what do they all have in common? Tallness? But what’s that? What does it take for something to instantiate tallness? Because, as in all matters metaphysical, we are rank amateurs, we don’t have much to say, but here are four preliminary options (there are obviously others):

1. For something to instantiate tallness there must be *some* comparison class or other with respect to which it’s tall. If that’s all it takes to instantiate tallness, it’s very easy to do so. We take this to be an exceedingly unpromising account of tallness.

2. It might be that to instantiate tallness it’s insufficient to be tall with respect to *some* comparison class. For each object there might be one such class that’s privileged, say, for natural kinds, the natural kind they belong to, for artifacts the artifact they instantiate. Since objects belong to many kinds, work would have to be done to show one of these is privileged.

3. A third option is that the circumstances the object is in at a time *t* singles out a comparison class that’s the one the object has to be tall with

⁵ In all these cases you could attempt to respond that it is a philosophical prejudice that there is something that all these things have in common. That it is a pun to say they are all dancing, eating, funny. We do not address that response here, but see ch. 11 of *Insensitive Semantics* for further discussion.

respect to in order to be tall at *t*. Again, work would have to be done to figure out how this comparison class is picked out.

4. The property of being tall corresponds to being taller than the average height for all objects that have height. Since we have no idea how many objects have heights we have no idea exactly what has this property.⁶

If you agree that there's a property of tallness—how could you not?—but have a better account of what it is to instantiate it, that's fine with us. Try it out on us. Which one is correct? We are not sure even how to determine an answer to this question. However, the only serious objection we can contemplate is to deny that there's any such thing as the property of being tall. Such cynicism would be to endorse Metaphysical Nihilism about tallness: that is, it would be to endorse the view that there's nothing A and B have in common if A is tall for a G and B is tall for an F. That view is, as far as we can tell, a rather bizarre view to hold because no one, as far as we know, denies there is any such a thing as *being tall with respect to some comparison class*. No one can deny there's such a thing as *being tall with respect to a privileged comparison class* or a *contextually salient comparison class*; or having the property of *being taller than the average height of all objects that have heights*. If this is so, then everyone agrees with us that at least for these four accounts of what the property of tallness is, each picks out something that exists. Whether it's the 'right' account is another topic.

In sum, our response to the first objection is this. If you think there is such a thing as tallness, then let that be the semantic value of 'tall' in 'Osama bin Laden is tall' and in answer to the question as to what it takes for that sentence to be true we say that it is whatever it takes for Osama bin Laden to have that property. To keep this answer in perspective remember:

- (a) We don't accept that it is a necessary condition on an acceptable semantic theory for English that it tells us what tallness is (even though we have given you some modest pointers for how to proceed).
- (b) According to Speech Act Pluralism, the semantic content of 'Osama bin Laden is tall' is not all of what the speaker who utters that sentence says; more generally, it does not fully determine the content of speech acts performed by people who utter that sentence.

⁶ A more elaborate discussion of these options would, in some ways, mirror contemporary debates about knowledge attributions. Both Stanley and Hawthorne propose theories according to which knowledge is some kind of interest relative property, but where this does not necessarily make 'know' a context sensitive expression (see Stanley, n.d.; Hawthorne, 2003). We imagine analogous arguments being made in connection with comparative adjectives. For some suggestions along these lines, see Graff (2002), and a reply by Stanley (2003).

First (and Only) Digression: Being Tall for an F is No Better than Being Tall

Suppose you're baffled by the idea that there's such a thing as tallness. We'll now try to show that if you are, you should be equally baffled by the idea that there is such a thing as, for example, being tall for a giraffe, or more generally, by the sort of property expressed by *being tall for an F*. This claim is dialectically significant because Semantic Contextualists tend to hold that this alleged problem occurs only for those who are Semantic Minimalists (as applied to comparative adjectives in particular). The fix, according to Semantic Contextualists, is supposed to reside with relativizing comparative adjectives to comparison classes, that is, with a commitment to contextualism for 'tall' and other comparative adjectives.

Adjectives like 'tall' are to be treated as relational with, for example, an unpronounced place for a comparison class that gets indexed in a context of use. So, for example, in effect, the sentence 'A is tall' is equivalent, on this contextualist suggestion, at some level of linguistic analysis, say, at the level of LF, with the representation 'A is tall for an F', where 'F' is an indexical that somehow receives its semantic value in context. For a sentence like (6), in one context of utterance the indexed comparison class (or property, or whatever) might be NBA players; and in another it might be Saudi Arabians.

Recall that the alleged problem for *tallness* is that it's mysterious what it is to be tall *simpliciter*: 'There can be no such thing as tallness *simpliciter*. To claim Kobe Bryant, Mount Everest, and the Empire State Building all have something in common—namely, *tallness*—is a mistake, and any semantics that presupposes there could be such a thing must be mistaken. Since Semantic Minimalism, as characterized, is committed to this possibility, it should be rejected.'

If this objection issues from anyone content with properties like *being tall for an F*, then it is terribly misplaced. Take the property of *being tall for a giraffe* as an example, that is, we're imagining an opponent who thinks that many things can instantiate the property of *being tall for giraffes*. Before proceeding with our inquiry, consider the following basic giraffe facts. Giraffes have hairy ears. The fleshy part of the ear stops before the hairs on the ears stop. Not every giraffe can stretch his neck all the way up; some are old and arthritic. (With assistance they might be able to stretch their necks further than without help.) Giraffes can stand on their back legs and lift their front legs into the air, and thereby, push themselves further up into the

air. That makes them longer. They have hoofs, and these hoofs wear down with usage.

Holding these simple giraffe facts in mind, consider two giraffes, say, A and B. What would it be for A and B to both instantiate the property of being tall for giraffes? The problem is this: there are many ways to be tall for a giraffe. For starters, there are indefinitely many ways to measure the tallness of giraffes. Consider these few illustrations. A giraffe's height can be measured:

- from bottom of his hoof to the fleshy tip of his ear with a self-stretched neck;
- from the bottom of a hoof to the tip of his snout with a self-stretched neck;
- from the bottom of a hoof to the hairy tip of an ear with a self-stretched neck;
- from the bottom of a hoof to the tip of a snout when standing on his back legs with his front legs lifted into the air;
- all of the above, with an artificially stretched neck, that is, by a machine or something else that can stretch the neck out further than the giraffe can by herself. (Remember, some giraffes are arthritic, and have very stiff necks.)

Then, of course, there's the question of which comparison class or property or whatever we are to compare any given giraffe to. Here are but a few options:

- all living giraffes;
- a stereotypical giraffe;
- french giraffes;
- all giraffes that have ever lived, are alive, and will ever live;
- all possible giraffes;
- all giraffes in the vicinity of a certain giraffe.

Then, of course, there's the question of the (optimum) conditions under which to measure a particular giraffe (holding the method of measurement and the comparison class fixed). Here are but a few of indefinitely many options:

- right after a bath (giraffes shrink a bit after having taken a bath);
- right after a long walk (their hoofs wear down);
- when dead (again, death shrinks us all);
- when hungry (they tend to stretch their necks further);
- when pregnant (their necks are rendered less flexible).

Let's stop here even though there is much else that has to be settled. But now ask yourself: what is it to *be tall for a giraffe*? what is *giraffe-tallness*? It all depends on *which giraffes* you compare any given giraffe to, *how* you measure it, the conditions of the giraffe when being measured, and so on. The 'and so on' here is vital. There are no obvious or a priori limits on the different variations on giraffe-tallness.

Just to remind you why this matters. We're imagining a Semantic Contextualist opponent who's completely baffled by the idea that there's such a thing as tallness and that it can be the semantic value of 'tall'. We've just tried to make that seem a little less peculiar by showing that the kind of worry that triggers befuddlement with respect to being tall should also, if legitimate, trigger the same sort of befuddlement with respect to being *tall for a giraffe*. Now, since we expect at least some opponents to be completely non-befuddled about being *tall for a giraffe*, at least before seeing our examples, this discussion might remove or alleviate some of their resistance to tallness.

Of course, we expect many opponents to say: 'Of course, there's no such thing as being tall for a giraffe *simpliciter*. You have to fill it out: you have to add something about the class of giraffes, the condition of the giraffes, the measuring methods, and so on.' To these critics we say: OK, just do it. Let's see how that gets incorporated into semantics and then we'll continue the debate.

Second Objection: Role of Semantic Content in Communication

Remember, according to Speech Act Pluralism, speakers use sentences to make claims, assertions, suggestions, requests, claims, statements, raise hypotheses, inquiries, etc., the contents of which can be (and typically are) radically different from the semantic contents of (the propositions semantically expressed by) these utterances. The speech act content (i.e. what was said, asserted, claimed, asked, etc.) depends on a potentially indefinite range of facts about the speaker, his audience, their shared context, the reporter (i.e. the person recounting what was said), the reporter's audience, and their shared context. These facts have no bearing on the semantic content of the utterance.

Here's a potential worry for this position. What communicators actually *care* about in a discourse exchange *is* the speech act content and *only* the

speech act content. What they care about is what the speaker said, asserted, claimed, stated, suggested, asked, etc. If this *isn't* the semantic content, if the semantic content is, so to speak, always hidden, if it never surfaces, then what purpose does it serve? Isn't it just an idle wheel? What would be lost if our theory just let it go? So, even if there is tallness, and even if the semantic value of 'is tall' somehow involves it, what role can this peculiar property play in communication? Does it have any kind of psychological reality? Let's call this the Psychological Challenge to Semantic Minimalism.

Reply to Second Objection: Semantic Content Does Have a Role to Play in Communication

We think the answer is simple and obvious but we can't overemphasize its importance. We begin by reminding you of some basic facts about communication. Then we respond directly to this psychological challenge. What we are about to say presupposes there being a clear notion of *a shared context*. We doubt there is one, but we'll place our reservations to the side for now. If there are shared contexts, then that will make life even harder for the Semantic Contextualist.

Basic Facts About Speakers and Audiences who Share a Context

Speakers are sometimes wrong (or have incomplete information) about their audience, for example, about:

- what the audience believes and knows;
- what the audience remembers about prior conversations;
- how the audience has interpreted previous conversations;
- how the audience perceives their shared environment; and
- what the audience believes about the speaker.

Audiences are sometimes wrong (or have incomplete information) about speakers, for example, about:

- what the speaker believes and knows;
- what the speaker remembers about previous conversations;
- how the speaker has interpreted previous conversations;
- how the speaker perceives their shared environment; and
- what the speaker believes about the audience.

Audiences and speakers are both often wrong (or have incomplete information) about the context that they find themselves in, for example, about:

- what their perceptual environment is; and
- what the contents of preceding conversations were.

Speakers and audiences know that they can be wrong and have incomplete information about each other in the ways just specified.

Basic Facts about Speakers and Audiences who do Not Share a Context

Sometimes the audience of an utterance doesn't share a context with the speaker. This can happen in any of several ways, the most salient of which being the reproduction of a speech act, as in published articles. Writers often have no idea who their reader is; they know next to nothing about her beliefs; or about her perceptual environment; all they know is that it is not shared. Yet, nonetheless, writers have audiences (no matter how small they might be).

Another typical device through which a speech act can reach an audience in another context is indirect quotation. This is when S says in C to A what another speaker S' said in another context C' to another audience A'. In these cases the sources of confusion are multiplied. The added complications should be obvious; there is not even the illusion of a shared context.

Basic Facts about *Inter*-Contextual Content Sharing

First, people can and often do say the same thing in different contexts. People in different contexts can say that Napoleon was short.

Second, according to Semantic Contextualists, no two contexts (are likely to) share exactly the same content fixing parameters, for example, the intentions are not the same; the background knowledge is not the same; previous conversations are not the same; what's normal is not the same; and so on (cf. e.g. Sperber and Wilson, 1986: 118, 192–3; Carston, 2001: 26–7; Recanati, 2004: 149; Bezuidenhout, 1997: 212–13).

Third, it is possible to say in a context C that people in a range of contexts C_1 – C_n said the same thing, for example, there are true reports, say, in C, of the form 'They all said that Napoleon was short' about different speakers'

utterances in contexts C_1 – C_n . (Similarly, distinct utterances can be collected; true utterances of the form ‘A is tall’, ‘B is tall’, and ‘C is tall’ said in contexts C_1 , C_2 and C_3 can be collected in a single context C_4 with an utterance of ‘A, B, and C are tall’.)

Note that if someone denies these three points we don’t want to talk to her or about her (because she doesn’t think she can say what we say, so she can’t deny what we say, and (according to her) we can’t say what she said, and so we can’t say that we disagree with what she said).

The Cognitive Role of Minimal Semantic Content

What, then, is the cognitive role of minimal semantic content? The answer should be (almost) self-evident by now:

1. Speakers know that their audience can be (and often are) mistaken (or have incomplete information) about the communication-relevant facts about the context of utterance. The proposition semantically expressed is that content the speaker can expect the audience to grasp (and expect the audience to expect the speaker to expect them to grasp) even if they have mistaken or incomplete communication-relevant information.

2. Audiences know that the speaker can be (and often is) mistaken (or has incomplete information) about the communication-relevant facts about the context of utterance. The proposition semantically expressed is that content the audience can expect the speaker to grasp (and expect the audience to grasp) even if she has such mistaken or incomplete information.

3. The proposition semantically expressed is that content which can be grasped and expressed by someone who isn’t even a participant in the context of utterance.

4. The proposition semantically expressed is that content which speakers and audiences know can be transmitted through indirect quotation or reproduction (in the form of tapes, video recordings, etc.) to, or collected by, those who find themselves in contexts radically different from the original context of utterance.

In short: the proposition semantically expressed is our minimal defense against confusion/misunderstanding/indifference, and it is that which guarantees communication across contexts of utterance. It’s what allows us to collect, report, and reproduce others’ utterances.

Possible Counter-Reply

We expect this sort of reply: ‘Hold it: You’re saying that the minimal semantic content is a “shared fallback content” and that this content serves to guard against confusion and misunderstandings. But given what you’ve told us about minimal propositions, how could they serve that purpose? Consider, for example, an utterance of (6). Suppose a speaker utters it to communicate that Osama Bin Laden is tall for a Saudi Arabian (or something like that). That’s what the speaker is trying to say. How would it help an audience to know that the minimal proposition, that is, *that Osama bin Laden is tall*, was expressed? It might not be what the speaker wanted to assert. What help could it be to know that this proposition was expressed?’

Our response is simple: it is a starting point. The audience knows that the speaker is talking about Osama Bin Laden and attributes tallness to him, and not, for example, to Sprite cans, Sweden, Britney Spears, or pig ears. There’s lots to talk about in the universe. The proposition semantically expressed pares it down considerably. Knowledge that this proposition was semantically expressed provides the audience with the *best possible* access to the speaker’s mind, given the restricted knowledge she has of that speaker. In general, audiences know what to look for in such situations; they know what kind of information would help narrow down more closely what the speaker wanted to communicate.⁷

To sum up our reply, consider the following charge from Recanati against Semantic Minimalism and our reply. Recanati writes of minimal propositions:

Let the semanticist use it if he or she wants to, provided he or she agrees that . . . the minimal proposition has no psychological reality. It does not correspond to any stage in the process of understanding the utterance, and need not be entertained or represented at any point in that process. (Recanati, 2004: 89)

If there’s a difference between having a cognitive function and corresponding to a stage in processing/having psychological reality, we don’t know what that difference consists in. If (1)–(4) above are insufficient to ‘correspond to a stage in the process of understanding the utterance and need not be entertained or represented at any point in that process’, then we don’t know what is.

⁷ There are many theories about how speakers go from semantic content to speech act content and we do not mean to, nor do we need to, endorse any one of those here.

In some sense, we're taking a stab in the dark here since we're not at all sure what Semantic Contextualists have in mind by the psychological requirement. What we have said is sufficient to render the propositions semantically expressed psychologically real, but we're genuinely confused since we have no idea how Semantic Contextualists satisfy their own requirement.

Concluding Point: The Second Objection Reversed (or Why Recanati's Account of What-Is-Said Doesn't Satisfy his own Availability Principle)

Suppose we focus, as Semantic Contextualists tend to, on the context of the speaker and her audience. The factors that figure into fixing the what-was-said/explicature include, *inter alia*, (i)–(iv):

- (i) information triggered in the speaker and the audience by prior discourse contents;
- (ii) information conversational partners share about each other;
- (iii) information the conversational partners have acquired through observation of their mutual perceptual environment;
- (iv) information conversational partners have about each other's purposes and abilities (e.g. whether the person is being deceitful or sincere, whether the person tends to verbosity, or is a person of few words).

These in no way exhaust the facts that, according to Semantic Contextualists, are content determinants, but what we have to say about (i)–(iv) generalizes. The problem is this. Suppose (i)–(iv) are factors that fix the explicature (i.e. the proposition expressed) of an utterance *u* of some sentence *S*. Now (i)–(iv) involve the mental states of several people (i.e. the speaker and her audience). None of the participants knows all the relevant facts about all the other participants: Herman doesn't know all the information triggered in Ernie by their many previous discussions; Ernie doesn't know what information Herman has about him. (He undoubtedly knows things about him that he doesn't even know he knows.) He doesn't always know what he will pay attention to in their sometimes shared perceptual environment; and so on.

The point here is obvious: if the explicature is fixed by these sorts of facts (what else?), then no one of the participants has direct access to the explicature. It is fixed intra-personally, and so there's no reason to think the resulting content is 'represented' at any stage of that person's processing of

the relevant utterance. There is no reason to think that the resulting proposition is psychologically real.

Recanati discusses a version of this objection and the utter failure of his reply illustrates just how hard it is for Semantic Contextualists to satisfy their own psychological reality requirement. In particular, it illustrates why Recanati can't satisfy his Availability Principle (his version of the Psychological Requirement).

Hence my 'Availability Principle' (Recanati 1993: 248), according to which 'what is said' must be analysed in conformity to the intuitions shared by those who fully understand the utterance—typically the speaker and the hearer, in a normal conversational setting. I take the conversational participants' intuitions concerning what is said to be revealed by their views concerning the utterance's truth-conditions. I assume that whoever fully understands a declarative utterance knows which state of affairs would possibly constitute a truth-maker for that utterance, i.e., knows in what sort of circumstance it would be true. (Recanati, 2004: 20–1).

Recanati's theory, based on his Availability Principle, is supposed to be an alternative to theories according to which the explicature/content/what-is-said is not psychologically accessible. Recanati's idea is that, since his what-is-said corresponds to the speaker's intuitions about what is said, it will figure in the process of understanding (an utterance of) the sentence. He raises this worry:

Have we not equated what is said with their [i.e. the speaker and audience] understanding of what is said?... We have not. We have equated what is said with what a normal interpreter would understand as being said, in the context at hand. A normal interpreter knows which sentence was uttered, knows the meaning of that sentence, knows the relevant contextual facts (who is being pointed to, etc.) Ordinary users of the language are normal interpreters, in most situations. They know the relevant facts and have the relevant abilities. But there are situations... where the actual users make mistakes and are not normal interpreters. In such situations their interpretations do not fix what is said. To determine what is said, we need to look at the interpretation that a normal interpreter would give. This is objective enough, yet remains within the confines of the pragmatic construal. (Recanati, 2004: 27)

But what's *normal* is not something speakers have psychological access to. What's normal need not 'be in the speaker's mind when the sentence is understood'; it certainly needn't figure into any psychological processes that the speaker goes through when understanding (an utterance of) a sentence. This is so for several obvious reasons; here are perhaps the most obvious ones:

- A speaker can be abnormal, but think that she is normal.
- A speaker might know that she is not normal, but not know what normal is.
- A speaker might think that she is not normal, but not be.
- More generally: even for speakers who are normal and know that they are normal, they might not know what counts as a normal understanding of some specific feature of a context that they happen to find themselves in.

A lot of situations have no ‘normal’ set of expectations associated with them. Suppose you meet someone in a cafe on a hot New York City summer day. What ‘normality’ are we looking for? Normal for you when talking to strangers in a cafe in New York City on a hot summer day? There’s no such thing!

In other words, if what’s normal, in part, determines what-is-said, and if what is normal is not represented at any stage in the processing of the utterance, then the resulting what-is-said cannot be so represented. Then, we suppose (though, as we have admitted, we’re not sure we entirely understand the Semantic Contextualists here), Recanati’s what-is-said is not psychologically real.

In sum: the Semantic Minimalist has a response to the Psychological Objection; it is the Semantic Contextualist who surprisingly does not.⁸

⁸ Reply to Stanley on Binding

This is not the place to rehearse our arguments for Semantic Minimalism (for a summary see our 2005), but we would like to indulge in one retrospective digression. In his contribution to this volume, Stanley responds to our criticism against his Binding Argument and we thought it appropriate to comment on his response. Stanley writes:

If the intuitive reading of (5) is (5^{*}), then it would seem that the advocate of the binding argument is committed to postulating a place variable in the logical form of ‘ $2 + 2 = 4$ ’.

[(5) Everywhere I go, $2 + 2 = 4$

(5^{*}) For all places x , if Sally goes to x , then $2 + 2 = 4$ at x .]

(Where we imagine (5) uttered by a confused mathematical anthropologist who travels the world to find out if mathematical statements are universal.)

I do not see that (5^{*}) is the intuitive reading of (5), and I do not see that Cappelen and Lepore even believe that (5^{*}) is a reading of (5). As Cappelen and Lepore (*ibid.*) point out, ‘... it is close to indisputable that arithmetical statements lack hidden indexicals referring to places.’ Presumably, the reason they are so convinced of this is that it is unclear what it even means to speak of an arithmetical statement being true at a place.

What we agree with Stanley about is (1) and (2):

1. If (5^{*}) is the intuitive reading of (5), then Stanley is committed to postulating a place variable in ‘ $2 + 2 = 4$ ’.

REFERENCES

- Bezuidenhout, A. (1997). 'The Communication of De Re Thoughts', *Noûs*, 31/2: 197–225.
- Cappelen, H., and E. Lepore (1997). 'On an Alleged Connection between Indirect Quotation and Semantic Theory', *Mind and Language*, 12: 278–96.
- (2003). 'Context Shifting Arguments', *Philosophical Perspectives*, 17: 25–50.
- (2005). *Insensitive Semantics* (Oxford: Basil Blackwell).
- Carston, R. (2001). 'Explicature and Semantics', in S. Davis and B. Gillon (eds.), *Semantics: A Reader* (Oxford: Oxford University Press).
- (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication* (Oxford: Basil Blackwell).
- Graff, Delia (2002). 'Shifting Sands: An Interest Relative Theory of Vagueness', *Philosophical Topics*, 28/1: 45–81.
- Hawthorne, J. (2003). *Knowledge and Lotteries* (Oxford: Oxford University Press).
- Kaplan, David (1989). 'Demonstratives', in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan* (Oxford: Oxford University Press), 481–563.
- Recanati F. (1993). *Direct Reference: From Language to Thought* (Cambridge: Blackwell).
- (2002). 'Unarticulated Constituents', *Linguistics and Philosophy*, 24: 299–345.
- (2004). *Literal Meaning* (Cambridge: Cambridge University Press).
- Sperber, D., and D. Wilson (1986). *Relevance* (Oxford: Basil Blackwell).
- Stanley, J. (2003). 'Context, Interest-Relativity, and the Sorites' *Analysis*, 63/4: 269–80.
- (n.d.). 'Context, Interest-Relativity, and Knowledge', unpublished MS.

2. There is no hidden argument place for location in arithmetical statements. (Note that Stanley himself uses the factive 'point out'.)

Where disagree with him is over (3) and (4):

3. That (5^*) is the intuitive reading of the described utterance of (5).
4. That it is unclear (and we think it is unclear) what it means to speak of an arithmetical statement being true at a place.

About (4): We have no difficulty whatsoever making sense of the idea of an arithmetical statement being true at a place. We can all make sense of the claim that $2 + 2 = 4$ is a necessary truth, i.e. that it is true *in every possible world*. If you make sense of that, surely you can make sense of it being true *at a place* or a time (for a time traveler). Test: ask yourself if you agree that our imagined mathematical anthropologist is wrong in assuming that there are places where $2 + 2 = 4$ is false, *because* $2 + 2 = 4$ everywhere.

About (3): The reason we think (5^*) is a natural reading of (5) is that every one (read unrestrictedly) we have asked whether (5^*) is a natural reading of (5) have agreed that it is. Also: the speaker of (5) might use (5^*) to say what she said, if asked to elaborate (or if asked: 'what do you mean?') Exercise: If you think (5) isn't the best way to say what she's wants to say, try finding a more natural way to say it.

If (5^*) is the correct reading of (5) and if it is not generated by quantifying over an argument place in all arithmetical statements, we owe an account of how this reading is generated. This note is not the place to present that account.