

**LYING, MISLEADING, AND LANGUAGE**

by

Matthew Knachel

B.G.S, University of Dayton, 1999

M.A., University of Wisconsin-Milwaukee, 2003

Submitted to the Graduate Faculty of the  
Kenneth P. Dietrich School of Arts and Sciences in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy

University of Pittsburgh

2014

UNIVERSITY OF PITTSBURGH  
KENNETH P. DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Matthew Knachel

It was defended on

April 3<sup>rd</sup>, 2014

and approved by

Michael Liston, Professor, Department of Philosophy, UW-Milwaukee

Mark Wilson, Professor, Department of Philosophy

Dissertation Advisor: Kieran Setiya, Professor, Department of Philosophy

Dissertation Advisor: James Shaw, Assistant Professor, Department of Philosophy

Copyright © by Matthew Knachel

2014

## LYING, MISLEADING, AND LANGUAGE

Matthew Knachel, PhD

University of Pittsburgh, 2014

My focus is the everyday distinction between lies and other deceptive speech acts—acts of misleading—which involve saying something truthful with the intention of causing one’s listener to have false beliefs.

The distinction resists straightforward pre-theoretical formulation, but it is closely tied to the concept of *saying*: if what someone said, strictly speaking, is true, then she did not lie—though she may have misled. Ever since Paul Grice distinguished what a speaker says from what she otherwise communicates (e.g., via implicatures), theorists have attempted to provide a rigorous circumscription of ‘what is said’. I test these theories according to how well they handle our intuitions about the lying/misleading distinction, and find that no extant account can adequately underwrite that distinction—in the process discovering that the boundary in question is even more difficult to draw than initially expected.

Positions in the debate over ‘what is said’ can be characterized as more or less minimal, according to how much pragmatic—as opposed to semantic—content is taken to be a part of what a speaker says. There are different accounts, though, of how to characterize the division of labor between Semantics and Pragmatics, so I turn my attention to those, applying the lessons of my previous investigations. I focus mainly on a prominent recent proposal from Herman Cappelen and Ernest Lepore (“Semantic Minimalism”), according to which semantic content is particularly austere. I argue that their account is deeply flawed. I then briefly examine some alternative views from the literature, expressing guarded skepticism about those.

Finally, I turn to non-linguistic philosophical questions that arise in connection with my focal concern: How should lying be defined? What is the moral (and legal) significance of the lying/misleading distinction? I argue that the literature on the definition of lying would be enriched by more carefully considering competing accounts of assertion, and that a Robert Brandom-style normative account seems the best option. I conclude that it's difficult to make the case for a moral difference between lying and misleading, but that a legal distinction can be made with the help of my earlier conclusions.

## TABLE OF CONTENTS

<b>PREFACE.....</b>	<b>ix</b>
<b>1.0 INTRODUCTION.....</b>	<b>1</b>
<b>1.1 WHAT IS SAID .....</b>	<b>3</b>
<b>1.2 THE SEMANTICS/PRAGMATICS DIVIDE .....</b>	<b>10</b>
<b>1.3 THE DEFINITION OF LYING, ASSERTION, AND THE LAW .....</b>	<b>22</b>
<b>1.4 OUTLINE.....</b>	<b>26</b>
<b>2.0 IN SEARCH OF ‘WHAT IS SAID’ .....</b>	<b>28</b>
<b>2.1 METHODOLOGY .....</b>	<b>30</b>
<b>2.1.1 Alternative methods.....</b>	<b>32</b>
<b>2.1.2 Concerns .....</b>	<b>35</b>
<b>2.2 STARTING AT THE BEGINNING: GRICE.....</b>	<b>40</b>
<b>2.2.1 Implicature: the basics .....</b>	<b>46</b>
<b>2.2.2 Logical Connectives .....</b>	<b>48</b>
<b>2.2.3 Other limitations of the Gricean picture .....</b>	<b>57</b>
<b>2.3 GENERALIZED CONVERSATIONAL IMPLICATURES .....</b>	<b>60</b>
<b>2.3.1 Categories of GCIs.....</b>	<b>62</b>
<b>2.3.2 Q-Implicatures .....</b>	<b>65</b>
<b>2.3.3 I-Implicatures.....</b>	<b>72</b>
<b>2.3.4 M-Implicatures.....</b>	<b>74</b>
<b>2.3.5 Final reflections on GCIs.....</b>	<b>76</b>

2.4	RELEVANCE THEORY .....	78
2.4.1	The basics.....	78
2.4.2	Appraisal.....	80
2.4	LESSONS FOR THE DEBATE OVER ‘WHAT IS SAID’ .....	86
2.5.1	Refining Grice’s minimal notion .....	87
2.5.2	Truth-Conditional Pragmatics .....	89
2.5.3	The role of a minimal ‘what is said’ .....	91
3.0	WHITHER SEMANTICS? .....	107
3.1	CAPPELEN AND LEPORE’S SEMANTIC MINIMALISM.....	109
3.1.1	Minimal propositions: metaphysical problems .....	112
3.1.2	Minimal propositions: psychological problems.....	123
3.1.3	Speech Act Pluralism.....	129
3.1.4	SPAP undermines other arguments .....	134
3.1.5	Summing up: how to use intuitions .....	144
3.2	ALTERNATIVES TO CAPPELEN AND LEPORE .....	147
3.2.1	Borg’s semantic minimalism.....	148
3.2.2	Indexicalism.....	154
3.2.3	Default Semantics.....	162
4.0	LYING, ASSERTING, MORALITY, AND THE LAW .....	168
4.1	THE DEFINITION OF LYING .....	169
4.2	ASSERTION .....	180
4.3	MORAL QUESTIONS.....	200
4.4	LEGAL MATTERS.....	213
	BIBLIOGRAPHY .....	225

*For my parents, Howie and Maryrita Knachel*



## PREFACE

I could not have completed this project without the help and support of many, so I beg indulgence while I use this space to acknowledge my debts and record my gratitude.

Extra special thanks are due to my mother-in-law, Barb Lorenz, who was extraordinarily generous with her childcare services—accepting only Miller Lite and the occasional home-cooked meal as remuneration. I would not have had enough time to write without Nana’s help.

Speaking of time, I’d also like to thank my decrepit 1986 Toro model 521 snow-blower for getting jammed up with wet snow and chopping off the top digits of two of my fingers. This injury convinced the Faculty of Arts and Sciences to extend my deadline an extra two semesters; I’m not sure I could have finished without the extra time.

For helpful discussion of the material from Chapter 2, I thank Rick Scaia, Dan Williams, Poop Secaur, Tim Upton, Johnny Sullivan, R.W. Haskins, Eric Romer, Hank Guelcher, Al Frecker, Jason Mennino, and Chris Lyons.

For their encouragement, guidance, and helpful comments on drafts, I thank the members of my committee: Kieran Setiya, James Shaw, Mark Wilson, and Michael Liston. Kieran deserves special recognition for encouragement and guidance: he convinced me it was possible to finish, facilitated the formation of my committee, and got me through the usual (and in my case, some unusual) administrative bottlenecks (with help from above, from the omniscient and ever-vigilant Kathy Rivet). James deserves special recognition for helpful comments: his notes were incredibly detailed and insightful. This would be a much worse piece of work if not for his

feedback, and that of the others. Of course it should go without saying, but any errors or other defects that remain are at least partially their fault; they read the whole thing, and there's four of them, for Pete's sake.

I have half as many daughters as committee members—but they're twice as fun and *way* better looking. I thank Rose and Alice for being a constant source of motivation, relaxation, consternation, and made-up illustrative examples (see below). I also thank them for their patience: I've been working on this thing, on and off, for Rose's entire life—and she's the older one (9 years). Suffice it to say, there has been a lot of time when dad couldn't play because "he's working on his book."

Finally, my deepest debt, as always, is to my wife, Maggie. If not for her inspiration and forbearance, I could not have done this.

## 1.0 INTRODUCTION

In 2005 I watched Harry Frankfurt appear on *The Daily Show with Jon Stewart* to discuss his (surprisingly best-selling) “book” *On Bullshit*.<sup>1</sup> He and Stewart had an amusing back-and-forth, in the course of which Stewart asked whether political “spin” is a species of bullshit and Frankfurt confirmed (somewhat tentatively) that he thought it was. I found this answer puzzling, since the crucial characteristic of the bullshitter, for Frankfurt, is his lack of concern for the truth: unlike the liar, or the honest man, he has no regard for truth-values—“He does not care whether the things he says describe reality correctly. He just picks them out, or makes them up, to suit his purpose.”<sup>2</sup> The politician, on the other hand, is profoundly concerned with truth, in the following sense: he doesn’t want to get caught saying something false, opening himself up to accusations of being a liar. He must therefore take extraordinary care in crafting his statements, especially, as is often the case, when he aims to cause his listeners to have false beliefs. This activity—speaking truly but manipulatively, often in such a way as to engender false beliefs—requires tremendous skill. Frankfurt himself acknowledges this, noting that in the realms of politics, public relations, and advertising “there are exquisitely sophisticated craftsmen who—with the help of advanced and demanding techniques of market research, of public opinion polling, of

---

<sup>1</sup> Frankfurt 2005. Scare quotes on ‘book’ because its content is just a short essay from decades before, made suitable for binding by manipulating page-size and margins. Here is the URL for the televised segment:

<http://www.thedailyshow.com/watch/mon-march-14-2005/harry-g--frankfurt>

<sup>2</sup> *Ibid.* p. 56

psychological testing, and so forth—dedicate themselves tirelessly to getting every word and image they produce exactly right.”<sup>3,4</sup> This is an apt description; he’s just wrong, I think, to characterize these people as bullshitters, since part of what makes their job so difficult is that they have to respect the truth (advertising and public relations firms are just as wary as the politician of being accused of lying).

The distinction to which these craftspeople must be highly sensitive is a familiar one: it’s the difference between lying and (merely) misleading.<sup>5</sup> It is quite possible to cause listeners to have false beliefs by saying things that are, strictly speaking, true. For example, grocery stores often advertise that, say, cans of soup are on sale by affixing a little placard to the shelf reading something like, “10 for \$10.” This sign reports, truthfully, that for \$10, one can purchase ten cans of soup. What the sign doesn’t mention, but which is also true, is that one can purchase a single can of soup for \$1, or two for \$2, and so on. The deal is: soup’s on sale for a buck a can. The reason they advertise the deal the way they do—rather than a more straightforward “\$1 each”—is that they hope that some customers will make an inference to the false belief that in order to get the special price, they must buy ten cans; the store’s hoping that some people will buy more soup than they otherwise would have. Hence, the poor little old lady laboring to push

---

<sup>3</sup> *Ibid.* p. 23

<sup>4</sup> One of the greatest contemporary craftsmen of this sort, Frank Luntz, has outlined his techniques in detail. See Luntz 2007.

<sup>5</sup> To be more precise, I should say ‘attempting to mislead’, since ‘mislead’ is a success-term, and so would describe the perlocutionary act of (successfully) causing false beliefs; lies, on the other hand, may fail in their deceptive goals (if such goals they have; the case of the “bald-faced” lie may be an exception), and so to lie is a mere illocutionary act. That said, I intend to bypass such fussiness, and use ‘mislead’ as shorthand for the illocutionary variant preceded by ‘attempt to’. (The distinction between illocutionary and perlocutionary speech acts comes from Austin 1962.)

A further note on ‘mislead’: misleading, unlike lying (at least normally; again, “bald-faced” lies may be an exception) is something one could do unintentionally—by accidentally mis-speaking, for example, in a way that causes one’s hearer to have a false belief. In exploring the distinction between lying and misleading, I intend to focus on intentional acts only, so I will use ‘mislead’ in this more restricted sense. It might be suggested that *deceiving* is always intentional, so ‘deceive’ might be a better choice than ‘mislead’; but I take it that lying and misleading (in my sense) are different *species* of deception, so I will stick with ‘mislead’, suitably qualified. (Thanks to Kieran Setiya raising questions about these distinctions.)

the cart laden with multiples-of-ten sized stacks of soup cans, boxed potatoes, and frozen dinners—most of which she doesn't need.

The focus of the present work is the boundary between such misleading speech acts on the one hand, and outright lies on the other. Contemplation of this boundary—where it lies, how to draw it, whether it's sharp or fuzzy, what its moral and legal significance should be, and so on—quickly involves one in the contemplation of a variety of philosophical questions. Viewing the issues from this vantage, I maintain, can give us fresh insights into vexed questions. The purpose of this introduction is to give the reader an overview of the philosophical literatures concerning some of the issues that arise when we focus on the lying/misleading distinction.

## 1.1 WHAT IS SAID

In order to distinguish between lying and (merely) misleading, we rely on some apparently commonsensical considerations. One is the distinction between what one says (strictly speaking), and what one merely implies. We might isolate the former by looking at what was actually uttered. In the case of the misleading sign at the grocery store, we have '10 for \$10'.<sup>6</sup> To decide whether this is true or false, we might inquire after another commonsense notion—what the utterance literally means; this will give us its truth-conditions. Clearly, the claim in question is true just in case \$10 will get you ten cans of soup, and so what was said is strictly speaking or literally true.<sup>7</sup> Any inference to the further claim that one must buy ten cans (or some multiple of

---

<sup>6</sup> Unnecessarily pedantic remark: it wasn't uttered, but written. Yes. Speech acts needn't be spoken; written communication can be analyzed the same way as spoken.

<sup>7</sup> There is an ambiguity here. Is 'strictly speaking' modifying 'what is said', or 'true'? I've apparently used it both ways. Does it make a difference? It could: as we'll see below, there may be stricter or looser standards according to

ten cans) in order to get a special deal goes beyond the literal meaning of what was said; such claims are merely implied. And whether or not they are implied is arguable: the grocery store might defend itself from accusations of deceptiveness, arguing that no such implication was intended; they use the larger numbers to highlight just how good a deal people are getting, to make it more vivid (you could stock up on a whole month's worth of soup for a relative pittance!). This feature of what is implied in communication—that it's debatable—isn't an aspect of what is said; that's an objective fact.

So suggests common sense, at least arguably. As a pre-theoretical first pass at the phenomenon, this account has some initial intuitive plausibility. But as soon as we begin to examine things a bit more closely, we quickly find ourselves in deeper linguistic and philosophical waters. One problem with our account as stated has to do with our proposed identification of what is said by the sign with the sentence on it: '10 for \$10' is not a grammatical sentence. And so (adding just one layer of terminological/philosophical sophistication) it cannot express a proposition, and there can be no question of truth-value. Of course, people communicate their thoughts using mere phrases and non-sentences all the time; we might say that such expressions are *elliptical* for complete sentences. But what exactly we mean by 'elliptical', and how we account for the phenomenon, is a vexed topic in linguistics and philosophy.<sup>8</sup> If what is said in an act of speech must be a sentence, then the sign in our example does not provide one; which sentence should be singled out as *the* thing that it says—or even whether there is only one such thing—is an open question. In trying to answer it, we might start

---

which we identify what is said, as when we report the utterances of others with varying degrees of fidelity to the syntax and semantics of the original; also, there can be stricter or looser standards according to which we evaluate the truth of claims, as when we allow degrees of truth for claims involving vague predicates, for example. I want to set aside degrees of truth, however. There is a sense in which demands for strict truth and strict utterance-content are the same, since the latter can simply be understood as truth-conditions; strict truth-conditions give you strict truth (or falsity). The ambiguity, therefore, makes no difference.

<sup>8</sup> See, for example, Stainton 2004 and 2005.

by noting that one could accurately report what the sign says as something like ‘For \$10, one can purchase ten cans of soup.’ The practice of reporting the sayings of others is quite common; such indirect reports needn’t quote the original exactly in order to be considered acceptable. But would *any* such acceptable report capture what was said by the original? For example, one might report what the sign said by giving its true upshot: ‘Soup costs a dollar per can.’ That may seem OK, but it’s also true that seven cans of soup cost \$7. And yet it seems wrong to characterize what the sign says with a sentence like ‘For \$7 one can purchase seven cans of soup.’ Nor are such problems unique to cases such as ours, in which the original act was committed using less than a full, declarative sentence. Elisabeth Camp has a nice example, in which the original utterance, made by a certain George to an acquaintance named Alex, is “I’ve run into Jim coming out of the Bluebird Diner the last three Monday nights.”<sup>9</sup> Alex, we’re told, knows that the Democratic Party holds its weekly meetings on Mondays at that diner. He therefore reports George’s original utterance as follows: “George says Jim’s been going to the Democratic party meetings at the Bluebird the last few Mondays.”<sup>10</sup> This report seems innocuous enough, until we imagine a chain of similarly altered reports eventually coming back to Jim, who’s confronted by a right-winger about his left-leaning views. Jim, who prefers to keep his opinions to himself, is angry at George, who made the original report of Jim’s comings and goings on Mondays, for saying things about his politics. But, if we suppose that George had no knowledge of Democratic meetings at the Bluebird, he can rightly object that he never said anything about Jim’s politics. So, though the practice of reporting what was said by others often involves the production of sentences that differ from those originally uttered (if, indeed, the original utterance even involved a sentence), we cannot necessarily rely on such reports to give us a correct

---

<sup>9</sup> Camp 2007, p. 206

<sup>10</sup> *Ibid.*

characterization of *what was said* originally. It is evidently not so easy to say what is said by any utterance.

Returning to our original focus, however, may help clear things up. We're interested in the boundary between lies and acts of misleading. As we noted, one might defend oneself against the accusation of a lie by appealing to the truth of what one said, *strictly speaking*. What we're after, then, is a more constrained notion, one which rules out the kinds of wanton speech-reporting that got us into trouble above. Still, we need some criterion by which we might precisely delimit the content of what is said. Paul Grice was the first theorist to provide a systematic account of this and related phenomena.<sup>11</sup> He distinguished, in acts of communication, between what is said, on the one hand, and what is "implicated", on the other.<sup>12</sup> Implicata (i.e., things implicated) go beyond the content of what is said in an utterance; they are aspects of what is communicated by an utterance that are meant by the speaker without being said. We'll explore Grice's ideas in more detail in Chapter 2, but for now I want to highlight what he says about 'what is said.' He's a bit vague on the subject, but states that what he has in mind is something "closely related to the conventional meaning of the words (the sentence)... uttered."<sup>13</sup> Further, he says that its content is determined by "the particular meanings of the elements of [the sentence uttered], their order, and their syntactical character."<sup>14</sup> Kent Bach dubs this the "*Syntactic Correlation* constraint."<sup>15</sup> What we have here is a kind of codification of the thoughts behind our commonsensical first pass above. What is said by an utterance is closely tied to, if not identified with, the actual sentence uttered and its conventional (or literal) meaning. This kind of constraint

---

<sup>11</sup> See Grice 1989.

<sup>12</sup> 'Implicate' was a neologism for Grice: "I wish to introduce, as terms of art, the verb *implicate* and the related nouns *implicature* (cf. *implying*) and *implicatum* (cf. *what is implied*)." *Ibid.*, p. 24.

<sup>13</sup> *Ibid.*, p. 25

<sup>14</sup> *Ibid.*, p. 87

<sup>15</sup> Bach 2001, p. 15



rules out all but the most faithful indirect reports of someone's speech act, thereby avoiding the troubles above. However, it gives rise to its own set of difficulties.

First, as we've already seen, we may not have, in our original utterance, a complete declarative sentence—and so syntactic constraint leaves us without truth-conditions. But even if we set those difficult cases aside, the identification of what is said even by a fully sentential utterance isn't as straightforward as it may seem. Even utterances of full sentences may leave us unable to determine truth-value—not because of any epistemological shortcomings we may have as hearers, but because the sentence, while grammatically complete, seems to fall short of expressing a full proposition. So, for example, if someone says "Rose isn't tall enough," the sentence alone doesn't provide us with enough information to make an evaluation.<sup>16</sup> We need to know what it is she's not tall enough *for*. Not tall enough to ride the rollercoaster, to reach the top shelf, to play for the Lakers? What? In order to evaluate such utterances, we need input from the context in which it is uttered. If we're at the county fair, in line for the Tilt-A-Whirl, we might judge the statement false; if we're evaluating her prospects for the upcoming NBA draft, we might judge it true. This kind of (apparent) contextual dependence of truth-conditions is a pervasive linguistic phenomenon, and it can manifest in a variety of ways. Some sentences, like 'Rose isn't tall enough', are grammatically complete, but their dependence on context to deliver truth-evaluable propositions is relatively obvious. Other sentences may also display such contextual dependence, but less-obviously. Suppose Allen has exactly three children. Now consider an utterance of 'Allen has two children.' False, right? Maybe. In most contexts, the utterance would mis-characterize the facts. But suppose the context is this: one of Allen's friends has just learned about a special tax-credit available to parents with two or more children; he's

---

<sup>16</sup> Rose is my older daughter, who's tall for her age, but not *that* tall. We'll see her in a lot of examples. Probably her little sister, Alice, too.

thinking about whom he should inform about this money-saving opportunity, and utters the sentence in question. Now the utterance is true; in this context, *at least two*, rather than *exactly two*, is the operative sense of ‘two’ in the sentence. Examples like this can be multiplied. The apparent ubiquity of such contextual dependence has led some to conclude that Grice’s strategy for using syntactically constrained literal meanings of uttered sentences to identify what is said is hopeless: in a suitable context, almost any sentence can be used to say a wide variety of things.<sup>17</sup>

In his 1985, Charles Travis takes several apparently straightforward sentences and shows that, despite first appearances, their truth-values are not so simple to determine when we imagine various contexts. Considering ‘The kettle is black’, he says, “Suppose the kettle is normal aluminum, but soot covered; normal aluminum but painted; cast iron, but glowing from heat; cast iron but enameled white on the inside; on the outside; cast iron with a lot of brown grease stains on the outside; etc.”<sup>18</sup> John Searle subjects ‘The cat is on the mat’ to similar treatment.<sup>19</sup> Faced with these considerations, one might begin to despair of using literal meanings of uttered sentences as a means of identifying what is said. Paul Ziff puts the concern pungently: “The factors that serve to determine what is said... appear to constitute a hopelessly unmanageable motley. ... One may have a firm grasp of the phonology, morphology, syntax and semantics of a language, thus a thorough knowledge of the language, and yet not understand what is said in that language.”<sup>20</sup>

This kind of nihilism about literal meaning, one might worry, could provide the liar with a Freddy Riedenschneider-style defense: who’s to say whether what I said was, strictly speaking,

---

<sup>17</sup> See Searle 1978, Travis 1985 and 1997, Ziff 1972.

<sup>18</sup> Travis 1985, p. 197

<sup>19</sup> Searle, *op. cit.*

<sup>20</sup> Ziff, *op. cit.* p. 38

true; “[T]here is no ‘what [is said]’—not in any sense that we can grasp, with our puny minds.”<sup>21</sup> But of course, just as Riedenschneider’s absurd defense rested on a misunderstanding/misrepresentation of (Werner) Heisenberg’s Uncertainty Principle, such a skeptical conclusion about ‘what is said’ based on Searle’s and Travis’s examples would rest on a similar mistake. For their point is not that there is no such thing as what is said by an utterance; rather, their aim is to undermine a particular conception of that notion, viz., a syntactically constrained, literal—one might say *minimal*—conception.

Now we can start to see a bit more clearly the contours of the philosophical landscape I aim to explore. We will examine various closely related debates within this space, using the distinction between lying and misleading as our point of entry. The notion of what is said (strictly speaking) by an utterance, as we’ve seen, is crucial to drawing that distinction, and so one of our main foci will be a debate over how to delimit that crucial notion. There is a range of views in the literature. At one end, we have relatively minimal characterizations of ‘what is said’, according to which it is very closely related to the literal meaning of the sentence uttered. This view must be somewhat moderated in light of the sorts of phenomena we’ve briefly mentioned—where the context of the utterance obtrudes to affect its content—but, as we’ll see, it’s still possible to elucidate a coherent position on which there is minimal difference between sentence-meaning and what is said by an utterance of that sentence. The further we move from this basic minimalist position on the spectrum of views about ‘what is said’, the more we allow contextual

---

<sup>21</sup> Coen and Coen 2001. In the film, Riedenschneider, a high-cost defense attorney in a murder trial, attempts to create doubt about his client’s guilt in the minds of jurors by claiming that there’s no such thing as ‘what happened’, since “[t]hey got this guy, in Germany. Fritz Something-or-other. Or is it? Maybe it’s Werner. Anyway, he’s got this theory: you wanna test something, you know, scientifically - how the planets go round the sun, what sunspots are made of, why the water comes out of the tap - well, you gotta look at it. But sometimes you look at it, your looking changes it. Ya can’t know the reality of what happened, or what would’ve happened if you hadn’t-a stuck in your own goddamn schnozz. So there is no ‘what happened’. Not in any sense that we can grasp, with our puny minds. Because our minds... our minds get in the way. Looking at something changes it. They call it the ‘Uncertainty Principle’. Sure, it sounds screwy, but even Einstein says the guy’s on to something.”

factors to contribute—and these factors can be quite unruly: they may involve, for example, nonce interpretations of newly minted metaphorical usages, that, while clear enough in context, cannot be determined by the application of linguistic rules and consultation of established word-meanings. If we go far enough along this path, we end up with Searle and Travis, where every utterance must be evaluated within its own context, where its content is completely up for grabs; but there are more moderate, intermediate positions, which allow to a greater or lesser degree for the intrusion of contextual factors into the determination of what is said. Our central question will be this: where along this spectrum can we find a conception of ‘what is said’ suitable to underwrite a clear distinction between lying and misleading?

## 1.2 THE SEMANTICS/PRAGMATICS DIVIDE

One way we might characterize the debate over ‘what is said’ is this: the more conservative minimalist positions see *semantic* factors as playing the primary role in its determination, while the more liberal views allow *pragmatic* considerations to play a central role. There is a closely related debate in the literature over how to demarcate the domains of Semantics and Pragmatics—indeed, over how to understand what we mean by those terms. The distinction between Semantics and Pragmatics as separate domains of linguistic inquiry goes back a ways. Rudolf Carnap influentially abstracted the domain of the former—“expressions and their designata”—from that of the latter—“investigation[s] [in which] explicit reference is made to the... user of a language.”<sup>22</sup> The basic idea is that Semantics covers conventional meanings of

---

<sup>22</sup> Carnap 1942, p. 9. This was part of a tripartite distinction, with Syntax the third discipline, abstracted from Semantics, dealing with “only the relations between the expressions....”

expressions abstracted from actual usage, while Pragmatics deals with questions about what speakers mean by using various expressions (with their given conventional meanings). This formulation is too simple, however, in that it ignores the fact that the truth-conditional contribution of many expressions—thought to be a purely semantic topic—depends upon their use. So-called indexical expressions can refer to various things, depending on context; words like ‘I’, ‘here’, ‘now’, ‘this’, and so on display this sort of flexibility. This would seem to cede to Pragmatics some role in what was thought to be the domain of Semantics—the assignment of truth-conditional content. David Kaplan rescues some systematic role for Semantics, though, with his distinction between character and content. According to this account, the conventional meaning of indexical expressions is not a full-fledged content (a referent, e.g.), but a function—called a “character”—*from contexts to contents*. The word ‘I’, for example, has for its character a function that maps the context of its use to its user, the speaker.<sup>23</sup>

With this insight in hand, it is possible to draw a boundary between Semantics and Pragmatics that departs only minimally from Carnap’s original formulation. The domain of Semantics is the literal, truth-conditional contents of sentences as determined by the rules of language. Accommodation can be made for indexicality in this picture: the literal content of indexical expressions is determined in context by their conventional meanings, which are just functions from contexts to contents. This is a highly constrained contribution from context; it contrasts with the unconstrained appeals to contextual factors typical of Pragmatics, which concerns itself with the determination of what speakers mean when they make utterances. Speaker-meaning typically outstrips literal meaning (as when, for example, implicatures are involved), so in making such determinations, one must consider a speaker’s intentions, the

---

<sup>23</sup> See Kaplan 1989.

identification of which might involve bringing to bear various bits of background knowledge about features of the context of discourse that go well beyond those used to nail down the values of indexicals (e.g., speaker, time, location). Those indices are part of what Kent Bach calls “narrow context”, as opposed to the “wide context” consulted in determining speaker-meaning. The former *provides* semantic values; the latter allows us to *ascertain* speakers’ intentions.<sup>24</sup> In this picture, Semantics provides an input for Pragmatics: literal meanings of expressions uttered are part of the information one must consider in determining speaker-meaning; they are an interpretational starting point, which aspects of wide context may or may not modify or augment in various ways as one zeroes in on the proper understanding of an utterance. Semantic and pragmatic processes are complementary aspects of utterance interpretation.

Something like this is what Francois Recanati calls the “Standard Picture” of the Semantics/Pragmatics divide.<sup>25</sup> It has been challenged. Many theorists point out that, even on this understanding of the division of labor between Semantics and Pragmatics, many supposedly semantic phenomena seem to involve pragmatic processes. One issue here is the question whether Semantics can always provide truth-conditional content. On the Standard Picture, it is supposed to; Semantics delivers literal meanings, and what are the meanings of (declarative) sentences but the truth-conditions of the propositions they express? So says a venerable tradition in the philosophy of language. But there are many sentences, it seems, whose semantic content falls short of truth-evaluable propositions. Take, for example, a sentence we’ve already seen: ‘Rose isn’t tall enough’. As we pointed out above, it’s not clear how we’re supposed to evaluate an utterance of this sentence for truth or falsity unless we know what it is she’s not tall enough *for*. And this information can only be provided by context. Now, the contextual information

---

<sup>24</sup> Bach 2004, p. 477

<sup>25</sup> Recanati 2004(a), p. 452

needed to make a determination in this case is not an aspect of “narrow context” as characterized above; it doesn’t seem to involve assigning values to indexical expressions. But if that’s the case, then what’s needed to provide truth-conditions for this utterance is the consultation of “wide context”—i.e., Pragmatics. Stephen Levinson calls this phenomenon “pragmatic intrusion”, to capture the idea that in these cases Pragmatics seems to be crossing the boundary and intruding into supposedly semantic territory—the determination of truth-conditions.<sup>26</sup> This sort of breach is qualitatively different from the (what Levinson calls “pre-semantic”<sup>27</sup>) contribution of Pragmatics to the determination of semantic values for indexicals: it is not—at least not obviously—driven by the syntax of utterances. Indexicals are a syntactical element of sentences whose meanings, as it were, *compel* us to consult context; Jason Stanley and Jeffrey King characterize this kind of pragmatic intrusion as a “weak pragmatic effect.”<sup>28</sup> “Strong pragmatic effects,” on the other hand, “[are] one[s] in which context affects what is communicated, but not by affecting the referential contents of any lexical item in a sentence.”<sup>29</sup> Now, one could argue, in the case of our example of ‘Rose isn’t tall enough’, that there is a lexical element in that sentence that demands contribution from context—viz., the word ‘enough’—so that in this case we may say that any propositional completion of the original (‘She isn’t tall enough to play for the Lakers’) is properly a semantic content on the Standard Picture. This is debatable. But there are myriad examples of the (supposed) phenomenon of semantic underdetermination—when the deliverances of Semantics are not propositional, and hence not truth-evaluable—and for many of them it’s not clear at all how one can massage things so that the effect of context can be characterized as “weak.” I’ll modify an example from Stephen Neale: I may use the sentence

---

<sup>26</sup> Levinson 2000

<sup>27</sup> *Ibid.*

<sup>28</sup> King and Stanley 2005, p. 118

<sup>29</sup> *Ibid.*, p. 119

‘The chair broke within a week’ to express the proposition that the chair *I purchased from Smith* broke within a week *of my purchasing it*.<sup>30</sup> Now, while it might be claimed that a specification of reference-time to fix the period of ‘within a week’ is semantically required, it’s much less plausible to claim that the semantics of ‘chair’ demands a specification of the person from whom it was purchased. The intrusion of strong pragmatic effects is much harder to deny in such cases.

Another problem for the Standard Picture is presented by cases in which there is no underdetermination—Semantics delivers a full proposition—but the purely semantic truth-conditions are, intuitively, the *wrong* truth-conditions. Suppose I utter the sentence ‘Alice finished her green beans and got a treat’ in a typical family-dinner context.<sup>31</sup> This sentence is a conjunction: its literal semantic meaning gives us truth-conditions according to which temporal order doesn’t matter; so the sentence would be vindicated by a scenario in which Alice first sneaked a chocolate chip cookie and then ate her green beans. But that’s not right. The original utterance tacitly included temporal order, as was clear from the context in which it was uttered: Alice finished her vegetables *and then* got a treat. This temporal reading of ‘and’ is typical. Consider ‘Rose put on her boots and went outside’ or ‘Alice took a bath, read a book, and went to bed’. What’s striking about these cases is that, while the temporal reading arguably goes beyond the strict semantic content of the sentences, it arises automatically, by default. It doesn’t seem to require the specification of some special context of utterance; on the contrary, one would have to invent a special context to *cancel* it and revert to the minimal conjunctive reading. These sorts of utterances arguably carry what Grice called ‘Generalized Conversational Implicatures’: “Sometimes one can say that the use of a certain form of words in an utterance would normally (in the absence of special circumstances) carry such-and-such an implicature or type of

---

<sup>30</sup> Neale 2007, p. 81

<sup>31</sup> Recall: Alice is my two-year-old. At dinner, the rule is: no vegetables, no treats.



implicature.”<sup>32</sup> A more generic term might be ‘default meanings’.<sup>33</sup> These are “salient meaning[s] intended by the speaker, or presumed by the addressee to have been intended, and recovered (a) without the help of inference from the speaker's intentions or (b) without conscious inferential process altogether.”<sup>34</sup> The question is how to categorize these, and this is the challenge to the Standard Picture. They’re truth-conditional, and so there’s a temptation to classify them as semantic phenomena; on the other hand, they’re separable from (arguably) more basic sets of truth-conditions—and though they may be settled on non-inferentially by hearers as the salient meanings of utterances, they nevertheless rely on (often prodigious amounts of) background information (the typical order in which children eat different foods, or have their bedtimes structured), making it tempting to classify them as pragmatic.

Faced with this apparent blurring of the line between Semantics and Pragmatics, theorists have developed a wide variety of responses. One can try to maintain the Standard Picture, or abandon it—or seek some middle ground. Among those who try to maintain an unmoderated view, with a strict separation between Semantics and Pragmatics, are King and Stanley.<sup>35</sup> For them, truth-conditional content is entirely a semantic matter, allowing only for pragmatic effects of the “weak” variety. The strategy is to subsume all apparently strong pragmatic effects under a suitably expanded account of indexicality; the logical form of sentences is said to contain myriad empty slots demanding to be filled by context, and these slots needn’t correspond to lexical items in the syntax.<sup>36</sup> This approach acknowledges the challenges posed by examples of apparent semantic underdetermination and contextually dependent default meanings and seeks to meet

---

<sup>32</sup> Grice 1989, p. 37. Levinson 2000 contains an elaborate theory of Generalized Conversational Implicatures..

<sup>33</sup> See Jaszczolt 2010(a). Though this is somewhat tendentious, in that the use of the word ‘meaning’ might be thought to beg the question whether such phenomena ought to be considered semantic or pragmatic.

<sup>34</sup> *Ibid.*, p. 1

<sup>35</sup> See also Stanley 2000 and 2002, Stanley and Szabó 2000.

<sup>36</sup> Neale (2007, p. 78, fn. 3) calls these “aphonic indexicals,” since they don’t have any phonological content.

them head on, by accommodating those cases within a purely semantic framework. Another option is to bite the bullet: admit that only Pragmatics can give us the more expansive truth-conditions intuitively applicable to such examples, but maintain that, appearances notwithstanding, Semantics does provide unique, fully propositional and truth-conditional content in such cases—indeed, in all cases of grammatically complete declarative sentences. This approach has come to be called Semantic Minimalism. Herman Cappelen and Ernie Lepore, for example, maintain that ‘Rose isn’t tall enough’, on its own, expresses a full, truth-evaluable proposition—never mind what she isn’t tall enough *for*.<sup>37</sup> Emma Borg has a somewhat different version of Minimalism, which is more radical than that of Cappelen and Lepore, driven by the conviction that semantic contents must be “formally tractable.”<sup>38</sup> For her, any allowance for speaker-intentions, as features of context, determining semantic contents is illegitimate; they are too unruly to meet the formal tractability condition, and must be ruled out. This means that overtly indexical expressions like pronouns, included by Cappelen and Lepore, and most other theorists, among the well-behaved context-invoking lexical items, but which sometimes require reference to the intentions of the speaker to resolve (what does the speaker have in mind when he uses ‘that’?), cannot be so casually dismissed as semantically unproblematic. Hers is an austere Minimalism; her minimal “reflexive propositions” look not at all like semantic contents as normally conceived.<sup>39</sup> One can also maintain a strict division of labor between Semantics and Pragmatics, but acknowledge the facts of semantic underdetermination. Bach, for example, allows that Semantics sometimes generates something less than a truth-evaluable proposition; in

---

<sup>37</sup> See Cappelen and Lepore 2005.

<sup>38</sup> Borg 2007, p. 355. See also Borg 2004. Borg’s motivation is to develop an account of semantic processing that fits together with a modularity theory of mind (*a la* Fodor 1983, e.g.), where a language faculty operating on formal elements delivers semantic values.

<sup>39</sup> On reflexive propositions, see also Perry 2001, Korta and Perry 2011

such cases, we have only a “propositional radical.”<sup>40</sup> A process of “completion” is required in such cases to generate a full proposition—as when we complete ‘Rose isn’t tall enough’ with ‘to play for the Lakers’.<sup>41</sup> In cases where Semantics does yield a complete proposition, but it differs from the one naturally recovered, Bach says that a process of “expansion” or “conceptual strengthening” takes place.<sup>42</sup> In ‘Alice finished her green beans and {then} got a treat’, the material in braces is the result of expansion in this sense. For Bach, the expanded/completed contents are implicit in the utterance of the original sentences, and so he refers to them as “conversational implicatures” (as opposed to Grice’s conversational implicatures; note the slightly different spelling).<sup>43</sup>

For Recanati, the positions canvassed so far all fall under the rubric of ‘Literalism’, since they all agree that “we may legitimately ascribe truth-conditional content (the property of... expressing a thought or a proposition) to natural-language sentences.”<sup>44</sup> They only disagree the extent to which contextual factors contribute to sentence-meanings. Recanati contrasts this Literalism with “Contextualism,” the view that “it is only speech acts, utterances in context, that have content in a basic, underived sense.”<sup>45</sup> The debate between these camps can be seen a descendant of the mid-20<sup>th</sup>-century debates between “Ideal Language Philosophers” (in the Frege, Russell, Carnap, Tarski tradition) and “Ordinary Language Philosophers” (later Wittgenstein, Austin, Strawson).<sup>46</sup> Like their forebears, Contextualists regard linguistic meaning as a malleable, use-based phenomenon: the meaning of an utterance depends on the situation in

---

<sup>40</sup> Bach 1994, p. 3 [My references are based on pagination of a copy of the paper available on Bach’s website: <http://userwww.sfsu.edu/kbach/>]

<sup>41</sup> *Ibid.*, p. 126

<sup>42</sup> *Ibid.*, p. 134

<sup>43</sup> *Ibid.*, p. 126

<sup>44</sup> Recanati 2004(b), p. 83

<sup>45</sup> *Ibid.*

<sup>46</sup> *Ibid.*, p. 1

which it's uttered, on the background assumptions of interlocutors; conventional meanings of words and language rules play a role, but their prescriptions are only established and maintained by previous and ongoing usage. The effects of context on the truth-conditions of utterances can't be limited to those mandated by conventional meanings (what Recanati calls "bottom-up" pragmatic processes, like the determination of reference for indexicals); optional ("top-down") processes play an indispensable role in determining the propositions we communicate.<sup>47</sup> So, unlike Bach, Recanati is not willing to relegate sentence-expansion to a merely complementary linguistic role, a secondary pragmatic modification to a primary, semantic meaning; the expanded proposition that Alice ate her green beans and *then* got a treat is the primary meaning of the utterance, and Pragmatics provides it. The literal, merely conjunctive meaning, since it plays no role in the interpretive process, is entirely dispensable.<sup>48</sup>

Contextualism comes in more or less radical flavors. One can maintain that the pragmatic processes involved in, say, expansion—while they're indispensable in ordinary contexts of utterance to generate the proposition meant—are nevertheless *optional*. There *could* be a context in which the minimal, literal meaning was picked out (in which, for example, one only meant that two conjoined events occurred, in whatever order). That the pragmatic contribution is optional is indicated by the fact that it is what Grice called "cancellable."<sup>49</sup> One can, without contradiction, retract or cancel the optional proposition, as in "Alice ate her green beans and had dessert—but not in that order." At the more radical end, we get views like Searle's and Travis's; these question the very existence of literal meanings, traditionally conceived, for word- or sentence-types. Consider how we learn language. Words are initially associated with situations,

---

<sup>47</sup> *Ibid.*, 18

<sup>48</sup> *Ibid.*, pp. 64 - 65

<sup>49</sup> Grice 1989, p. 44

and we refine our understanding of them by applying or mis-applying them in new situations, being corrected as necessary along the way. But no amount of training of this sort can ever offer us *complete* guidance: there may always be novel situations in which we're unsure how to proceed—not because of some epistemic failing on our part, but because there's no such thing as a fully determinate “meaning” that could provide answers in every possible situation. This is the point of fanciful examples involving cats and mats floating in orbit, kettles made of aluminum but stained black, and so on. This situation makes utterance interpretation, the determination of speakers' meanings, an essentially, unavoidably pragmatic, contextual affair. For Contextualists, then, Semantics is at best a bit player, at worst completely irrelevant.

Unless, that is, we take arguably the most radical step of all, throwing open the gates to welcome the pragmatic barbarian hordes, as it were, and re-conceiving Semantics and semantic processing, not as a syntactically constrained, minimally contextual affair, as is traditional, but rather a multi-faceted, richly contextual process, with the contribution from syntax and lexicon but one factor among many in determining semantic content. Katarzyna Jaszczołt's Default Semantics takes this approach. She happily accepts the label of ‘Radical Contextualist’<sup>50</sup>, and aims to provide an account of utterance-meaning as the sum (she uses the symbol  $\Sigma$ ) of information from several processing resources, all of which are conceived as on a par with one another: not only syntax/lexicon, but also traditionally pragmatic sources such as background knowledge about the world, society, and culture; the discourse situation; properties of the human cognitive system. The result of this summation (called the “merger representation”) is the semantic content of an utterance, the speaker's “primary” (as opposed to “secondary”)

---

<sup>50</sup> “Default Semantics... sits comfortably in the contextualist camp and in its radical flank, but also goes significantly beyond some of its assumptions.” Jaszczołt 2010(b), p. 2. [My references are based on pagination of a copy of the paper available on Jaszczołt's website: <http://people.ds.cam.ac.uk/kmj21/>]

meaning—that which it was her intention primarily to convey, and which may or may not be explicit in the utterance, may contain implicatures and other indirect speech acts.<sup>51</sup> The summation is supposed to be compositional and truth-conditional, and so Jaszczolt feels comfortable referring to it as semantic content. But one concern with this approach is that to subsume so much of what formerly counted as pragmatic under the semantic umbrella is to obliterate the distinction between Semantics and Pragmatics entirely—or at least to render what many consider a substantive debate a mere question of terminological preferences. Another concern is whether Jaszczolt’s alternative distinction between primary and secondary meaning is compatible with other commonly accepted distinctions, such as those between what is said and what is implicated, explicit and implicit content, etc. And of course we will have questions about whether the framework can save the phenomena relating to the lying/misleading distinction.

Returning to our main focus—the distinction between lying and misleading—we’re now in a position to give a more terminologically sophisticated statement of the commonsense approach with which we began. One merely misleads, and does not lie, if what one says, strictly speaking, is true. We can isolate this strict content of what is said by appealing to the distinction between Semantics and Pragmatics: Semantics gives us what is said, strictly speaking, while Pragmatics provides any content communicated that goes beyond this (*implicata*, e.g.). But now we can see just how naïve this approach is. Not only is it unclear whether purely semantic factors can account for our intuitive notion of ‘what is said’, it’s unclear precisely what counts as purely semantic, as opposed to pragmatic, content. We have a wide variety of views on these questions, and they will give us a wide variety of different accounts of the distinction between lying and misleading. Our project, then, is to canvass these positions, draw out their implications for our

---

<sup>51</sup> See, e.g., Jaszczolt 2010(b).

distinctive focus, and ask whether those implications comport with our intuitions about how to differentiate between lies and other deceptive acts. To the extent that they don't, this is a strike against the view in question. It needn't be a fatal blow, however: a mismatch between the implications of a sufficiently plausible position in the philosophy of language, on the one hand, and our intuitions about lying and misleading, on the other, may provide evidence that it's the latter, rather than the former, that needs to be re-thought.

The initial animating concern for this project was a recognition that a commonly used distinction—between lying and misleading—depends upon concepts the precise analysis of which has been the subject of long-standing philosophical disputes. The hope was that in an examination of the literature, I might find an account well-suited to provide a theoretically robust underpinning for the commonsense distinction. This hope, as we shall see below, has gone mostly unrealized. However, there is value in conducting such a search. First, it never a bad thing to subject one's intuitions about everyday matters to critical scrutiny under the light of the more sophisticated terminology provided by the relevant theoretical debates; it may be that our intuitions are sufficiently imprecise that only vague, ambiguous conclusions can follow from them; if that is the case, then such an examination can help us clarify our thoughts on important subjects. Second, bringing to bear one's intuitions about everyday phenomena can shed new light on theoretical debates. Only very rarely in the literatures I have heretofore discussed has the distinction between lying and misleading been mentioned, much less a central focus.<sup>52</sup> Making it the focus will at the very least help throw into sharper relief the differences between the various relevant positions in the philosophy of language and linguistics, giving us fresh insight into

---

<sup>52</sup> Meibauer 2005 and 2011, Stokke 2013(a) and 2013(b), and Saul (2013) are exceptions.

what's at stake in those disputes; it will also help us identify any serious weaknesses in theorists' positions.

### 1.3 THE DEFINITION OF LYING, ASSERTION, AND THE LAW

It does not go without saying that my focal distinction, between lying and misleading, should even be made. There is a large philosophical literature on lying, of course, and many writers define it in such a way that there is no difference between lying and misleading. Many others make the preservation of that distinction an explicit desideratum. Either way, it's clear that there's a relationship between the positions one takes on the issues already canvassed and one's views about the appropriate way to define lying: for instance, an expansive conception of Semantics, such as that favored by Jaszczolt, is compatible with a broader definition of lying, since the subsumption of traditionally pragmatic phenomena like implicatures under the category of primary meaning apparently makes classic cases of mere misleading into lies; it would be odd to deny that in cases where a speaker's primary meaning is false she is lying. An examination of the literature on the definition of lying would seem to be in order, particularly with an eye toward ways in which it might be informed by the explorations of the issues already canvassed (and vice versa).

There is very little agreement among philosophers about how precisely to define lying, but a consensus forms at least around one necessary condition: lying involves *saying* something.<sup>53</sup> A famous example from Kant illustrates the restriction: he imagines someone

---

<sup>53</sup> We can include various forms of non-verbal saying, e.g. writing, sign language, and so on.



packing his bags in front of someone else, in order to get the observer to believe that he (the packer) is going on a trip.<sup>54</sup> While this is an act of deception, it is not a lie; lies involve statements, and the packer never speaks.

But beyond the general agreement that lies are statements of some kind, there is widespread disagreement about the conditions that must obtain for a particular statement to qualify as a lie. Some maintain that lies are false statements<sup>55</sup>; in this, they're in agreement with the Oxford English Dictionary, which defines a lie as "a false statement made with the intent to deceive." Others deny that lies must be false statements: some to accommodate the intuition that telling someone something one *believes* to be false, but which is, in fact true, is still lying<sup>56</sup>; some whose definitions include among lies even true statements that merely imply (or implicate) falsehoods.<sup>57</sup> Recently, some philosophers have attempted to articulate definitions that do away with the OED's second condition—the intent to deceive.<sup>58</sup> They point to so-called bald-faced lies—falsehoods averred despite the speaker's knowledge that he has no chance of deceiving anyone—and adjust their conditions on lying to accommodate them. Discussion of these and other questions in their vicinity (Is it enough that the liar believes his statement *probably* is false? Or that he lacks belief in its truth? Must the liar have an audience? Is it possible to lie to eavesdroppers? Must the liar intend to deceive the addressee about the proposition he expresses? Must he also intend to deceive the addressee about his own belief in that proposition? And so on.) has generated a large literature trying to codify our intuitions about lying in the form of necessary and sufficient conditions. It has even been suggested by some that the search for

---

<sup>54</sup> Kant 1797.

<sup>55</sup> E.g., Carson 2006.

<sup>56</sup> See Siegler 1966. Many authors share Siegler's intuition about such cases: Simpson 1992, Mannison 1969, Chisholm and Feehan 1977, Williams 2002, Mahon 2008(a), Fallis 2009.

<sup>57</sup> See Bok 1978, Meibauer 2005 and 2011.

<sup>58</sup> See especially Carson 2006 and Sorenson 2007.

necessary and sufficient conditions is fruitless, because the meaning of ‘lie’ is a “cognitive prototype” to which different acts can correspond in varying degrees.<sup>59</sup>

A perhaps more interesting set of issues—from the point of view of the present work—emerges if we focus on the requirement about which there is most widespread agreement, viz., that lying necessarily involves saying something. If we set aside concerns about whether what is said must be false, believed false, and so on, we can ask a more fundamental question, one with which we’re already familiar: what counts as what is said for any given utterance? What contents are we evaluating for falsity (or believed-falsity, or whatever), and how are they determined by the actual words uttered? And further, to what sense of ‘saying’ must an act conform in order to count as a lie? An adequate account of lying must answer these questions, but too often in the literature they are given inadequate attention or ignored entirely. For example, in his otherwise comprehensive overview of the various positions in the literature on defining lying, 2008(b), James Mahon opts to present each alternative definition using the neutral and unexplained ‘make a statement’. Among those who have recognized that something more than this is required, the consensus seems to be that lying is a species of assertion.<sup>60</sup> The accounts of assertions offered by these authors vary in their ambitions to specificity. There is, of course, a whole literature on the speech act of assertion.<sup>61</sup> Only recently has an attempt been made to bring it to bear on the question of how to define lies—by Andreas Stokke in his 2013(a), which seeks to define lying using the account of assertion first elucidated by Robert Stalnaker. More work remains to be

---

<sup>59</sup> Coleman and Kay 1981. They back their claims with experimental data about people’s intuitions.

<sup>60</sup> See, e.g., Chisholm and Feehan 1977, Simpson 1992, Williams 2002, Fallis 2009, Stokke 2013(a).

<sup>61</sup> Seminal works include: Bach and Harnish 1979; Brandom 1983, and 1994; Dummett 1973; Lewis 1979; Searle 1969; Stalnaker 1978; Strawson 1964; Williamson 1996. That list is partial; the literature is vast. Pagin 2008 gives an exhaustive overview, MacFarlane 2011 a nice concise one.

done. It is an open question whether other prominent accounts of assertion are suitable for the task.

Parallel to controversies among philosophers about how best to understand and define the notion of lying are controversies among legal scholars about how best to understand and define cognate legal notions, especially perjury. One could succinctly define perjury as lying under oath, and so some of the considerations that we can bring to bear on the issues so far discussed ought to be relevant in the legal context. The relevant U.S. statute defines perjury as occurring when a witness “states or subscribes any material matter which he does not believe to be true.”<sup>62</sup> This condition is quite similar to those encountered in the literature on defining lies, and makes room for the kind of distinction that’s been our focus all along—between lying or perjuring oneself on the one hand, and merely misleading (in a non-perjurious way) on the other. The most prominent ruling on perjury in the United States is the Supreme Court’s 1973 decision in *Bronston v. United States*. In Bronston’s original trial, the following exchange occurred: “Q. Do you have any bank accounts in Swiss Banks, Mr. Bronston? A. No, sir. Q. Have you ever? A. The company had an account there for about six months, in Zurich.”<sup>63</sup> Bronston’s final answer was true; what he didn’t say was that he had also had a personal account in Zurich. Bronston was convicted of perjury based on that testimony, but the Supreme Court overturned that ruling, maintaining that the literal truth of Bronston’s answer precludes a perjury conviction. There has been vigorous debate among legal scholars about the ruling in *Bronston*, with some maintaining that the Court got it right, and others arguing that the Court’s decision is only the most prominent example of the pernicious influence of “Literalism” in legal thought.<sup>64</sup> Of course, perjury is not

---

<sup>62</sup> 18 USC. §1621.

<sup>63</sup> *Bronston v. United States*, 409 US 352 - Supreme Court 1973

<sup>64</sup> See Solan 2002, Green 2000 and 2001 for the former view, Tiersma 1990 and Simon 2002 for the latter.

the only deceptive act covered in the law—we have libel, slander, fraud, and a host of other crimes—and there is much variation in the language with which these various acts are proscribed: some statutes (like that for perjury) seem to call for a narrower reading, so that what’s forbidden is more like lying; others cast a wider net, making it illegal in certain contexts to engage in what we’ve been calling mere misleading. It is naturally to be hoped that our detailed theoretical investigations into the linguistic and philosophical underpinnings of the notions of lying and misleading, and the distinction between the two, might allow us to make helpful contributions to debates among legal scholars about the merits and demerits of alternative approaches both to drafting and interpreting laws that prohibit deceptive acts.

## **1.4 OUTLINE**

We will take up the issues just briefly introduced in order. Chapter 2 will focus on questions about how to determine ‘what is said’. We will approach the issue from the perspective of a concern to account for the distinction between lying and misleading, and ask what a concept of ‘saying’ would have to look like to underwrite that distinction. We will search through the literature on ‘what is said’ for a suitable theoretical account. We will find none, but the search will prove instructive.

The focus of Chapter 3 will be the question of how properly to draw the distinction between Semantics and Pragmatics. We will attempt to bring to bear the lessons of Chapter 2 on the relevant issues. We will focus at length on one particular proposal in the literature—

Cappelen and Lepore's version of semantic minimalism—and then briefly discuss alternative approaches due to Borg, Stanley and Jaszczolt.

Finally, in Chapter 4, we will take up miscellaneous issues that arise when we consider how best to define lying. Besides various objections and counterexamples, and the corresponding alterations to the definition that they suggest, we will examine in detail how various accounts of assertion fit with these definitional efforts; whether there is any moral significance to the lying/misleading distinction, and how different answers to that question affect our definition of lying; and the legal issues, involving perjury and other proscribed deceptive actions, that are the natural adjunct of a concern with the distinction between lying and merely misleading.

## 2.0 IN SEARCH OF ‘WHAT IS SAID’

The purpose of this chapter is to make a contribution to an on-going debate in linguistics and philosophy of language over the concept of ‘what is said’. Roughly, the debate is between those who favor a more austere conception, on the one hand, and those who favor a more robust one, on the other. This debate often involves reflection on our intuitions about specific cases, and my contribution will feature much of this. There is a logically prior debate, though, about what kinds of intuitions are appropriate to rely on in circumscribing the targeted concept; section 1 addresses this question. My focus on the lying/misleading distinction suggests a middle way between—or a synthesis of—two positions. After discussing these methodological preliminaries, we will proceed to apply our methods to specific cases. The purpose of this activity is to conduct a search: we’re looking through the theoretical literature on the subject for an account of ‘what is said’ that’s suitable to underwrite the distinction between lying and misleading. The theorists whose views we’re discussing, with a few exceptions, did not intend their work to be put to this use, so a caveat is in order: failure of an account to comport with our intuitions about lying and misleading should not be construed as a direct criticism of that view; it may be that there are multiple conceptions of ‘what is said’, and those that fail our tests may help explain other intuitions and linguistic phenomena, or have other theoretical uses. One of our results will be that this is indeed the case, and the work we do will help throw into sharper relief the differences between competing accounts.

Our search through the literature for a theoretical account of ‘what is said’ suitable to explain the difference between lying and misleading will be a failure; no extant view will do. But our failure, and the reasons for it, will allow us to gain some fresh insight. We will start our search from the beginning, as it were, with the pioneering work of Grice; in section 2, we will see that Grice’s austere ‘what is said’ is too narrow a notion to provide a basis on which to draw the distinction between lying and misleading. The cases we examine will suggest a natural way of extending Grice’s notion—including so-called Generalized Conversational Implicatures (GCIs) as part of what is said—and in section 3 we will test this approach, relying on the elaborate theory of GCIs advocated by Stephen Levinson. This more expansive ‘what is said’ will also fail to capture our intuitions about deceptive acts, but in the course of analyzing a variety of cases, one of our most important lessons will become clear: the distinction between lying and misleading is an extremely murky one, with intuitions exhibiting a variability that makes it difficult to see how *any* theoretical account of ‘saying’ could possibly capture it. Not yet completely deterred, though, in section 4 we will continue our search for such an account by looking briefly at a major player in the debate about ‘what is said’: Relevance Theory. After concluding that, despite some promising aspects, it will not serve our purposes either, we will conclude in section 5 pulling back a bit to look at the big picture, extracting whatever lessons we can about the debate between advocates and opponents of austerity about ‘what is said’. Our exemplars of the former school will be neo-Griceans like Kent Bach and Larry Horn; Francois Recanati will represent the latter view. The main lesson will be that the two sides of the debate are talking about two different notions of ‘what is said’—the more austere one merely locutionary, the more robust one illocutionary—each of which has a basis in our linguistic practice and a role to play in our linguistic lives, but which we must be careful not to conflate.

The notion of ‘what is said’ that we need to underwrite the distinction between lying and misleading, while not fully circumscribed by any extant theory, is evidently an illocutionary notion—but a merely locutionary ‘what is said’ sometimes plays a role in explaining our intuitions in this area. This suggests a puzzle: how to explain the complex and variable relationship between these two conceptions of ‘saying’. We have no solution to offer—only the observation, derived from our examination of a wide variety of cases, that no simple solution will do. We conclude with a thorough critique of an attempt, by Horn, to offer such a simple solution.

## 2.1 METHODOLOGY

There are disputes in the literature about specific cases, and types of cases, over whether or not some propositional content is part of ‘what is said’. Our approach to adjudicating these types of disputes will draw on intuitions that we have about circumstances in which the issue of whether someone has lied, or merely misled, is salient. We will invent scenarios in which it is, and apply the following test:

- Let  $u$  be an utterance, let  $t$  be a true proposition (arguably) expressed by  $u$ , and let  $f$  be a false proposition (arguably) expressed by  $u$ .
- Consider a defense against the accusation that  $u$  was a lie, on the grounds that it’s  $t$ , not  $f$ , that was said. If we would reject such a defense, then  $f$ , not  $t$  captures ‘what is



said’<sup>65</sup>; on the other hand, if we would accept such a defense, then *t*, not *f* captures ‘what is said’.<sup>66</sup>

So, for example, imagine the following scenario: you’re paying a mid-day visit to a neighbor who’s an enthusiastic but terrible cook; he offers you a bowl of what he’s calling “gumbo,” but which to the best of your olfactory and visual reckoning resembles nothing more closely than burnt pieces of vulcanized rubber floating in dirty dishwater. You respond to the offer politely, but deceptively, thus: “No, thanks. I’ve had lunch.” In fact, you haven’t eaten anything that day, but you’re not about to break your fast with this so-called gumbo; you’ll grab a sandwich or something as soon as you get home. Now suppose that it comes to pass that, a bit later, back at your house, the bell rings and it’s that same neighbor, now at your door; he’s come bearing a Tupperware container full of his “gumbo,” thinking that maybe you’d like to have some for dinner, when you’re not so full from lunch. But as you answer the bell, you open the door with one hand while holding in the other a hero sandwich the size of a football. You’re still chewing your most recent bite, and there’s Italian dressing dribbling down your chin. Your neighbor, his feelings hurt, exclaims, “I thought you said you already had lunch! You lied to me.” Setting down your grinder and wiping your chin, you defend yourself with this speech: “No, I didn’t lie. What I said to you was, strictly speaking, true. The sentence that I uttered—‘I’ve had lunch’—expresses a proposition with the following truth-conditions: it is true if and only if there is at least one occasion in the past on which I’ve eaten that particular meal. Now,

---

<sup>65</sup> Or at least: *f* is *part* of what is said. We may reject such a defense on the grounds that the speaker says more than one thing, at least one of which is false

<sup>66</sup> The test makes a quite minimal, and I think plausible, presumption: that a necessary condition for lying is saying something false. There’s a vast literature on how to define lying, and some deny even this minimal claim. More on such definitions in Chapter 4.

since my life before today has not been entirely lunch-less—in fact, I’ve eaten lunch many, many times—what I said was true. Hence, I did not lie.”

We can apply our test to this scenario. The utterance  $u$  is “I’ve eaten lunch”; the true proposition  $t$  that it (arguably) expresses is that your life hasn’t been entirely lunch-less before the time of utterance; the false proposition  $f$  that it (arguably) expresses is that you ate lunch before the time of utterance but on the same day as the utterance. Your defense against the accusation of lying, on the grounds that it was  $t$ , not  $f$ , that you expressed, would be rejected. It’s laughable—and not merely because it’s couched in ridiculous-sounding technical terminology (that just makes its absurdity more vivid). Since we would reject such a defense,  $f$  captures ‘what is said’: you told your neighbor that you’d already had lunch *that day*. You lied.<sup>67</sup>

### 2.1.1 Alternative methods

As we’ve noted already, this way of proceeding is novel: with few exceptions, no-one has approached these issues from a perspective focused on the lying/misleading distinction. Whether or not this approach is also fruitful, of course, can best be demonstrated by applying our test to concrete cases. It is nevertheless possible to say some things on its behalf before we proceed to those cases.

In the literature about ‘what is said’, there has been debate specifically over the question of which kinds of intuitions we ought to rely upon to help us nail down this notion. On one side,

---

<sup>67</sup> This claim can be disputed from a different direction: one could argue that, since your motivations in deceiving were pure—you wanted to spare your neighbor’s feelings, refusing his offering without implying that he’s a bad cook—that this is not a genuine lie. Some authors define lying in such a way that the act is intrinsically wrong (see Grotius 2005), which makes these kinds of “white lies” impossible. If lies are necessarily told with malicious intent, then any act which lacks that intent is not a genuine lie. I’ll set this concern aside for the moment and assume that white lies are real lies; the proper definition of lying is a topic for Chapter 4.

there are those who claim that we should favor intuitions about the truth or falsity (or appropriateness, or accuracy) of indirect reports of speech acts, reports of the form “So-and-so said that blah.” To the extent that such reports go through, ‘blah’ is (at least part of) what is said.<sup>68</sup> Cappelen and Lepore: “We take our practice of indirect reporting at face value and assume that the speakers have said what we have the reporters saying that they have said.”<sup>69</sup> Bach formalizes the approach, calling it the “IQ Test”: “An element of a sentence contributes to what is said in an utterance of that sentence if and only if there can be an accurate and complete indirect quotation of the utterance (in the same language) which includes that element, or a corresponding element, in the ‘that’-clause that specifies what is said.”<sup>70</sup> Recanati, on the other hand, insists that there is a better way to isolate ‘what is said’, relying not on intuitions about the viability of indirect reports of speech, but rather on intuitions about the truth-conditions of utterances. “One has simply to provide subjects with scenarios describing situations, or, even better, with—possibly animated—pictures of situations, and to ask them to evaluate the target utterance as true or false with respect to the situations in question.”<sup>71</sup> Normal interpreters understand what is said by an utterance—its meaning, the proposition it expresses—just in case they know how the world would have to be to make it true; utterance understanding is grasp of truth-conditions. It is important for Recanati that what is said by an utterance be consciously available to interlocutors; if it plays no conscious role in utterance interpretation, or if it is not part of what a speaker intends to convey in an utterance, then it cannot be ‘what is said’. Conflict arises between his approach and Bach’s because some propositions that pass the IQ test are not

---

<sup>68</sup> Cappelen & Lepore and Bach, who advocate this method, also allow that more than one thing can be said by a single utterance—hence the parenthetical hedge.

<sup>69</sup> Cappelen and Lepore 2005, p. 50

<sup>70</sup> Bach 1999, p. 8 [My references are based on pagination of a copy of the paper available on Bach’s website: <http://userwww.sfsu.edu/kbach/>]

<sup>71</sup> Recanati 2004(b), p. 15

consciously available. For example, I might say “Alice and Rose are sisters.” By now you, dear reader, know that those are my two girls, and so you interpret me as saying something the truth-conditions for which include the requirement that they have the same parents. For Bach, this is too expansive an understanding of what is said; there is no element in the syntax of the sentence to indicate this. Rather, a more minimal, literal proposition corresponds to what is said by my utterance, viz. that each of the girls has at least one sibling. They may have different parents and live in different counties; they may never have met, for all I’ve *said*.

Now, on the face of it, my approach, with its focus on lies (and hence the truth and falsity of what is said), has more affinity with Recanati’s than with Bach’s. And indeed, for the case just mentioned, if we invented a scenario in which the context made it clear that I was communicating that the girls were related, but in fact they weren’t, my test would render the same verdict about what was said as Recanati’s.<sup>72</sup> And while I’ll often agree with his assessments, our results will not always line up exactly. This is because my test relies on *both* intuitions about intuitive truth-values *and* on intuitions about the appropriateness of indirect reports of utterances. It combines both competing approaches. The truth-conditional element of my test is clear: I’m relying on intuitions about when statements are lies, and hence false, or not, and hence true. But since my scenarios involve an imagined *defense* against an accusation of lies, a defense that takes the form of an indirect report on one’s own speech (“No, what I said, strictly speaking, was...”), I incorporate Bach’s approach as well. And so, for example, we will have some cases in which the minimal sense of ‘what is said’ that he defends does indeed play a role in our linguistic lives. My approach will help us to say how his and Recanati’s understandings of that key notion differ.

---

<sup>72</sup> Exercise for the reader.

At this point I'll just issue a promissory note: the dual nature of my test will enlighten. We will see that a strict reliance on one type of intuition or other leads to counterintuitive results, but that their combination will help clarify things. Perversely, the nature of this clarification will be a realization of just how fuzzy the boundaries are between lies and non-lies, Semantics and Pragmatics, what is said and what isn't. Focus entirely on truth-conditions or indirect quotation gives the illusion of clear-cut answers; my approach reveals the reality that things are a damn mess.

### **2.1.2 Concerns**

Despite its promise, we should take note of some concerns that one could have about my methodology. First, one might object that the reliance on intuitions to solve philosophical problems is an outmoded technique that ought to be displaced by a more rigorous, experimental approach. This is the position of so-called Experimental Philosophers. Joshua Knobe and Shaun Nichols characterize the enterprise thus: “[E]xperimental philosophers proceed by conducting experimental investigations of the psychological processes underlying people’s intuitions about central philosophical issues.”<sup>73</sup> I will not be conducting any such studies; let’s just say I don’t have the budget for that sort of thing (I’m still making payments on my armchair). But this is not to suggest that such investigations would be out of place in a project of this kind. In fact, systematic studies of people’s intuitions about lying and about ‘what is said’ have been conducted.<sup>74</sup> Some studies, for example Gibbs and Moise 1997 and Nicolle and Clark 1999,

---

<sup>73</sup> Knobe and Nichols 2007, p. 3

<sup>74</sup> For studies of intuitions about lying, see Coleman and Kay 1981 and Arico and Fallis (forthcoming). These studies will be more relevant to our discussion in Chapter 4.

present subjects with scenarios with utterances-in-context and ask them directly what they think has been said. The former claims to vindicate Recanati's views, since subjects consistently reported the intuition that 'what is said' goes beyond minimally semantic propositions; the latter casts some doubt, failing to reproduce the same results. Either way, Recanati 2004 objects to the approach of both, on the grounds that they're testing the wrong intuitions: as we've seen, he thinks the proper way to delimit 'what is said' is to focus on intuitions of ordinary speakers about the truth-conditions of utterances.<sup>75</sup> Bach also objects to the approach, claiming that what such data "are indicative of is how people apply 'what is said'," not what is, in fact said in any given utterance.<sup>76</sup> People's intuitions about 'what is said', he points out, can be "sensitized" by exposing them to examples of cancellability: for example, Gibbs and Moise reported that people confronted with an utterance of 'Martha gave John her key and he opened the door' had the intuition that part of what was said was that John opened the door with the key in question; but their intuitions might change if they were confronted with 'Martha gave John her key and he opened the door—not with her key, but with another one' and asked whether it was contradictory; Bach predicts they wouldn't.<sup>77</sup> People's untutored intuitions may change with some tutoring. Of course, this is a testable hypothesis. Gibbs and Moise don't test it, and neither does Bach—but he does offer his own testable predictions, pointing the way for future research. The best I can do is the same. I'll be using my own intuitions to arrive at my own conclusions, the generalization of which will also be testable. It would be instructive to conduct studies to see whether I am vindicated; we can speculate along the way about what such studies should look like, how their results might differ from others using different approaches.

---

<sup>75</sup> Recanati 2004(b), p. 15

<sup>76</sup> Bach 2002, p. 6 [My references are based on pagination of a copy of the paper available on Bach's website: <http://userwww.sfsu.edu/kbach/>]

<sup>77</sup> *Ibid.*

As interesting as such investigations might be, however, I don't think we should conclude too hastily that they are the only, or even the best, source of data about the phenomena we're investigating. The language itself is another source. For example, the prevalence in the language of the phrase 'some, but not all', is a linguistic datum with which we must reckon. Such a phrase, though it might be grammatical, would certainly not be so frequently used if it were redundant. When we use the word 'some' by itself, we frequently communicate the sense of 'some, but not all': I might utter 'Some of my students are hard workers' in a typical context and naturally and correctly be taken to mean that not all of them are, for example. The fact that the phrase 'some, but not all' is not redundant, however, indicates that 'not all' is separable from 'some'. Any account of 'what is said' by such utterances must account for this fact: to what degree is this evidence for a minimal construal of such utterances? And the fact with which we're dealing here goes beyond speakers' intuitions; it concerns the prevalence of a certain usage. Another, related example: it is a fact that, among the four corners of the traditional Square of Opposition for Aristotelian logic<sup>78</sup>, only the bottom-right subcontrary (O) lacks its own typical lexical incarnation; for A we have 'all', for E we have 'no(ne)', and for I we have 'some'—but nothing for O (we don't contract 'not all' to get, e.g., 'n'all'). This lack of lexicalization of the O-corner is true not only of English, but across languages. And it generalizes within languages: one can map other linguistic relationships onto an oppositional square, and it is always the bottom-right subcontrary that goes unlexicalized. So we have 'and' (A), 'or' (I), 'nor' (E), but no 'nand' (O); we have 'always' (A), 'sometimes' (I), 'never' (E), but no 'nalways' (O).<sup>79</sup> Larry Horn takes this fact to be significant for debates about which we're concerned. As we just noted, utterances

---

<sup>78</sup> A = Universal Affirmative, 'All S are P'; E = Universal Negative, 'No S are P'; I = Particular Affirmative, 'Some S are P'; O = Universal Negative, 'Some S are not P'

<sup>79</sup> See Horn 1989, §4.5 and Horn 2009 for more examples.

featuring the left-hand (I) subcontrary, e.g. ‘Some philosophers are bald’, typically communicate the content from the right-hand side as well, giving us ‘Some {but not all} philosophers are bald’. Horn’s view is that the bracketed material is an implicature, not part of what is said, since a pragmatic account of this phenomenon is “more general and more explanatory than rival theories.”<sup>80</sup> The details of his argument for this claim needn’t concern us. The point I want to make is simply that speakers’ intuitions aren’t the only sorts of data to which we can appeal in our investigations of the issues at hand.

Another focus of concern about my methodology—especially in light of Bach’s observation about sensitizing intuitions—might be *me*. That is, one might worry that my intuitions, being the result of extensive tutoring in philosophy, linguistics, and logic, won’t be at all representative of normal speakers’ intuitions. This is a legitimate concern, but we shouldn’t make too much of it. Besides being conscious of my own potential theoretically informed biases and taking as much care as I can to avoid them (I have consulted as frequently as possible with the relatively untutored intuitions of friends and family as a safeguard), I can only offer the following: increased tutoring might in some cases provide me with *more* insight into the intuitions of ordinary speakers than I would otherwise have. The kind of tutoring I have in mind is logical—and *I’m* the tutor. Over the course of many years of teaching logic, I have become (painfully, at times) aware of a variety of common linguistic intuitions among untutored speakers. For example, I know there is a strong tendency to consider ‘not all’ to be part of the “meaning” of ‘some’. I’ve seen first-hand the preference for exclusive ‘or’. I have much data on the tendency toward conditional perfection (in which a mere conditional is interpreted as a biconditional); every attempt to explain the conditional reading of ‘unless’ provides fresh

---

<sup>80</sup> Horn 2009, p. 7



evidence. And so on. This kind of tutoring has provided me with a greater sensitivity toward common intuitions, and should be considered an asset.

One final concern we might have about my methodology stems from its focus on lying. I regard this focus as the source of one of its virtues: I think that contemplating scenarios in which there is something at stake—a moral judgment that we care about—can only sharpen our intuitions. But one might worry that the addition of moral considerations may muddy the waters. The intuitions we’re after (about ‘what is said’, literal meanings, etc.) are *linguistic*, not moral; but might we not make an apparently linguistic judgment for (unconsciously) moral reasons? The design of my test seems to open up this possibility. It asks us to imagine cases in which a defense is being made against the accusation of having lied, and to make a judgment about whether such a defense would be acceptable. But because of the morally salient features of the case under discussion, we might be less (or more) inclined to accept a defense. If, for example, the deceptive act has particularly dire consequences for its victim, we might be more inclined to call it a lie—since that judgment carries more negative normative weight—rather than a mere deception, no matter what the linguistic facts may be. I think this possibility of confusing moral and linguistic intuitions is a serious concern. I can offer three thoughts to address it.

First, simply being aware of the possibility, and keeping it in mind when we’re contemplating scenarios, will be a safeguard against its occurrence. Second, if we suspect that moral considerations have obtruded into our linguistic deliberations about a particular case, we can conduct a second test, a kind of control. We keep the linguistic phenomena constant, but change the moral facts. For example, if we’re inclined to reject an imagined defense against lying in a particular case, and we suspect that it’s (only or mainly or partially) for moral reasons (the consequences of the speech act are abhorrent), we can alter the scenario so that the utterance

is the same, but the moral circumstances are less severe (we can always invent a case in which the lie, if it is one, is a “white” one, told with benevolent intentions). If our judgment changes, then we were right to suspect moral intrusion; if not, then our linguistic intuitions are vindicated. And third, to some degree at least, the possible admixture of moral and linguistic intuitions is not so much a bug as a feature. One of my central concerns in this work is to try to find, within the vast philosophical literature on the relevant topics, a theoretical basis on which to draw an *everyday* distinction—between lying and misleading. This distinction is part of everybody’s linguistic practices, and so everybody has intuitions about it. If those intuitions (about ‘what is said’) are affected by the salient moral features of the situations in which it’s at stake, then so be it. Insofar as we’re trying to find a suitable theoretical conception to underwrite the everyday distinction, we must take those intuitions at face value and ask whether they’re captured by the proposal at hand.

## **2.2     STARTING AT THE BEGINNING: GRICE**

Our topic is the difference between lying and misleading. Both are deceptive acts, in that they share the same goal of causing the listener to have false beliefs. But they use different means. When you lie, you just say the false thing you want your audience to believe; when you mislead, you’re more subtle. You avoid saying the falsity, and if you’re really clever, you say nothing but truths. The falsehoods you convey are not part of what you say, but rather part of what you imply or otherwise indirectly communicate. Confronted with your deception, you can at least fall back

on a defense to the effect that you didn't lie. This mitigates the damage somewhat, since lies are typically taken to be more serious moral and legal transgressions.

When you mount a defense against lying, you rely on the notion of 'what is said'; that is, you claim of the thing that you said that it is (strictly speaking, perhaps) true. The question is: how do we delimit the content of what we say, as opposed to what we merely imply? Philosophers have discussed this and related questions for some time, so it's not unreasonable to hope that, somewhere in this discussion, we might find a theoretically sophisticated basis on which we can draw the boundary in question and thereby provide a principled and sharp distinction between lying and merely misleading. Since Grice was the first philosopher to attempt to map of the philosophical terrain in this area, we shall start our search with him.

The seminal work for the philosophical investigation of the distinction between saying and implying is Grice's 1967 William James lectures at Harvard (collected, revised, and reprinted in Grice 1989). It was here that he introduced the idea of "implicatures", dividing the totality of what we communicate between these and what we say. Grice's favored notion of *saying* will be the subject of this section. Our question is whether this notion can underwrite a principled distinction between lying and merely misleading.

Grice is, by his own admission<sup>81</sup>, somewhat vague about how to understand 'saying' in his favored sense. But we can understand it well enough to evaluate its suitability for underwriting our distinction. One thing we know, we've had occasion to remark on already: what someone says in an utterance is supposed to be closely related to the syntax of what is uttered. There is some question in the literature about whether or not the content of what is said is to be

---

<sup>81</sup> Grice 1989, p. 86

equated with truth-conditional content, or whether it is something distinct (and more austere).<sup>82</sup> However such questions are resolved, though, the Syntactic Correlation constraint makes Grice's notion unsuitable for underwriting our distinction, as we'll see. For our purposes, the most interesting aspect of Grice's conception of 'what is said' is that it is, by definition, distinct from any contents otherwise communicated—i.e., implicated. We can apply our test to show that any such understanding of 'what is said' is too narrow to form a basis for distinguishing lies from non-lies.

By requiring that the content of what is said by an utterance be closely related to the linguistic elements of that utterance, Grice constrains that notion to an unacceptable (for us) degree. As we've already seen, some utterances of fully grammatical sentences fail to provide truth conditions unless we allow elements of context—pragmatic factors—to contribute the missing pieces. Grice's 'what is said' is almost entirely a semantic affair. *Almost*, because he does allow minimal contributions from Pragmatics: "[F]or a full identification of what the speaker had said [on a given occasion], one would need to know (a) the identity of [the speaker], (b) the time of utterance, and (c) the meaning, on the occasion of utterance, of [any ambiguous words or phrases]."<sup>83</sup> So, for Grice, context can resolve ambiguities and provide the values for any indexical expressions. This is *some* contextual input, but often not enough to give us something truth-evaluable (see above, e.g., 'Rose isn't tall enough'). This fact alone makes Grice's notion of 'what is said' too limited for our purposes: the mark of an act of (mere) misleading is that what is said in that act is (strictly speaking) true; lies involve saying things that

---

<sup>82</sup> See Levinson 2000, p. 170 – 171 and Bach 1994, §5. One question is whether or not syntactically distinct, but truth-conditionally equivalent utterances count as sayings of the same thing. For example, if I utter 'John loves Mary', do I say the same thing as when I utter 'Mary is loved by John'? Levinson and Bach both think not: Grice's notion is at the level of semantic representation (logical form, structured propositions), and so it would distinguish these two.

<sup>83</sup> Grice 1989, p. 25

are false. If what is said by an utterance may not be truth-evaluable, we cannot decide between the two.

One might suppose we could repair the damage here by allowing for some additional contextual factors to contribute in cases of semantic underdetermination and still get a recognizably Gricean, minimal conception of what is said—and that it would suit our purposes. But even if we could deal with the problem of underdetermination, we would still have other problems. For in many cases in which the minimally semantic ‘what is said’ is a truth-evaluable proposition, it’s the wrong one for our purposes. Instead, the version of ‘what is said’ that we need to distinguish lies from non-lies in such cases turns out to include content that is communicated, by Grice’s lights, by implicatures.

One more note before we proceed: a scrupulous reader of Grice might object to our imputation to him of a kind of semantic minimalism about what is said—the view that what we say, as opposed to what we implicate, is determined by the meanings of the terms uttered with only a highly constrained contribution from context (which can disambiguate and determine reference for a short list of indexicals). There is, after all, little textual evidence to support this imputation: we have the quote I cited two paragraphs ago, and the quotes that Bach takes to establish the Syntactic Correlation constraint (what is said is “closely related to the conventional meaning of the words... uttered” and determined by “the particular meanings of the elements of [the sentence uttered], their order, and their syntactical character.”<sup>84</sup>). A determined Grice exegete could point to this paucity of evidence, and then interpret these passages as consistent with a denial of semantic minimalism about what is said. Further, our sympathetic Gricean could point out that there is an alternative view, not equivalent to semantic minimalism, for which we

---

<sup>84</sup> *Ibid.*, p. 25 and p. 87

have ample textual evidence to support the claim that it is genuinely Grice's—viz., the thesis that the way to isolate what is said (as opposed to merely implicated) is to rely on various diagnostic tests for the presence of implicatures, such as cancellability and reforcability (implicatures can be cancelled without contradiction and felicitously reinforced without redundancy).<sup>85</sup> Since this view and minimalism are not equivalent, one might worry that even if I'm successful in demonstrating that the notion of 'what is said' revealed by the minimalist criteria is not suitable to underwrite the distinction between lying and misleading, I have not thereby shown that the (possibly different) notion of 'what is said' revealed by the diagnostics also fails in this respect.

So there are two worries to which we need to respond: (1) Grice may not really be a semantic minimalist about what is said; and (2) by targeting that conception of what is said, we may miss a viable alternative—'what is said' as revealed by the usual Gricean diagnostics. I'll take these two concerns in turn. First, in imputing to Grice a semantic minimalism about what is said, I am not doing anything out of the ordinary. In the literature to which I'm presently contributing, it is taken for granted that this is Grice's view. So, for example, Recanati says, "Even though he construed saying as a variety of non-natural meaning, Grice espoused Minimalism. On his view, disambiguation and saturation [fixing indexicals and anaphora] suffice to give us the literal interpretation of the utterance--what is literally said."<sup>86</sup> This is typical; Grice's minimalism is never questioned by the authors with whom I'm engaging. Now, maybe everybody's wrong, and some great Grice exegete can demonstrate that, but I'm not interested in Grice exegesis. If we're all wrong, then my target is not the real Grice, but some alternative Grice\* who holds the minimalist views everybody attributes to him. Second, on the topic of the diagnostics (cancellability, reforcability, etc.), we can set aside, I think, concerns about their

---

<sup>85</sup> Our sympathetic Grice exegete is not a fictional character: thanks to James Shaw for stressing these points.

<sup>86</sup> Recanati 2004, p. 27

reliability for identifying implicatures (expressed by Grice himself and later by others<sup>87</sup>); the main concern is that, whatever its merits as a means of identifying implicatures and thereby isolating what is said, it may generate results to which my tests are blind. I don't think this should be much of a concern. In what follows, I will focus my attention on a Grice (Grice\*?) who identifies what is said as semantically minimal content, but with few exceptions, this approach, and one based on the diagnostics, will agree on what is said. Consider an example from earlier, in which I said to my neighbor, "I've had lunch." According to a semantic minimalism about what is said, the proposition expressed here is that there exists a time, before the time of the utterance, at which I, the speaker, had lunch; it's true just in case my life hasn't been entirely lunch-less. Any content beyond that is implied or implicated. The Gricean diagnostics render the same verdict: I can non-redundantly reinforce the interpretation that I have had lunch *today*, simply by adding that word to the utterance ("I've eaten lunch today"); and I can cancel that aspect of meaning without contradiction, for example, thus: "I've eaten lunch—not today, mind you, but I've lunched before." In what follows, almost without exception, the same pattern will repeat: I will try to show that a minimalist conception of 'what is said' is unable to support our intuitions about the distinction between lying and misleading, by identifying content that goes beyond the minimal, but which our test nevertheless suggests ought to, at least arguably, be included as part of what is said; this content will also be cancellable, reinforcing, etc., so the same considerations will apply to both Gricean approaches to isolating what is said.

---

<sup>87</sup> See Grice 1989, p. 43: "I [would not go] so far as to suggest that it is possible, in terms of some or all of these features, to devise a decisive test to settle the question whether a conversational implicature is present or not...[i]ndeed I very much doubt whether the [diagnostics] can be made to provide any such knock-down test...." See also Sadock 1978, p. 293: "The test [cancellability] does not distinguish cases of ambiguity from cases of univocality plus possible conversational implicature. One of the senses of a grammatically ambiguous sentence may always be contradicted."

### 2.2.1 Implicature: the basics

Before we get into this, though, we should briefly survey the different kinds of implicatures, their characteristic features, and the ways they are generated by speakers and interpreted by hearers. An implicature is any component of what a speaker communicates that goes beyond what she merely says. Grice draws a distinction between implicatures that are “conversational” and those that are not. The non-conversational implicatures arise from the meaning of the words used in an utterance, and so Grice calls them “conventional” implicatures. Grice’s original example is an utterance of ‘He is an Englishman; he is, therefore, brave’. For Grice, the utterer *says* that the person in question is an Englishman and that he is brave, but he has not said that the one implies the other; that is implicated.<sup>88</sup> The implicature arises because of the conventional meaning of the word ‘therefore’; hence, it’s a conventional implicature.

On the other side of the distinction, we have conversational implicatures. These arise not because of the meanings of the words used, but because of certain principles governing conversation generally. These are unwritten rules or maxims to which competent speakers can be assumed to be adhering in the course of conversation. Since communication is a cooperative endeavor, Grice identifies as the principal maxim a Cooperative Principle: “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.”<sup>89</sup> Under this master rule fall various sub-maxims, falling into (Kant-inspired) categories: Quantity, Quality, Relation, and Manner. So, for example, under the heading ‘Quality’, we have a principle according to which one ought not to say anything false. The details of maxims needn’t detain us. The important point

---

<sup>88</sup> Grice 1989, p. 25

<sup>89</sup> *Ibid.*, p. 26



is that the presumption that a speaker is observing these rules, along with the data provided by what she says, allows the hearer to infer implicata. For Grice, someone saying that *p* conversationally implicates that *q* if (1) the speaker can be presumed to be following the maxims; (2) *q* must be supposed in order to maintain this presumption; and (3) the speaker thinks the hearer realizes this (condition (2)). Grice's first example has one interlocutor saying "I'm out of petrol," and the other responding, "There's a garage around the corner." The respondent implicates that (he thinks) the garage is open and has petrol to sell; these propositions follow from the presumption that he's following the maxim: "Be relevant." Other implicatures can arise when maxims are violated, as when one implicates one's lack of knowledge by being less than fully informative; also, when maxims are flouted, as when a letter-writer for a student implicates that he has a low opinion of the student's abilities by praising irrelevant attributes ("so-and-so has excellent penmanship and is a snappy dresser").<sup>90</sup>

Grice makes one more distinction: he identifies two types of conversational implicatures, generalized and particularized. The former arise normally or automatically, in the absence of special circumstances, because of the form of words used. We've already seen some examples of Generalized Conversational Implicatures (GCIs): 'some' normally gives rise to an implicature of 'not all'; certain instances of 'and' normally give rise to a temporal implicature of 'and then'. Grice's first example is: "X is meeting a woman this evening," which implicates that the woman is not X's wife.<sup>91</sup> There are many others. Particularized Conversational Implicatures (PCIs), on the other hand, don't arise normally or automatically: they are communicated only because of the peculiar circumstances of the utterance that carries them.

---

<sup>90</sup> *Ibid.*, pp. 32 – 33

<sup>91</sup> *Ibid.*, p. 37.

### 2.2.2 Logical connectives

We are searching for a notion of ‘what is said’ suitable for underwriting the distinction between lying and misleading. I have asserted that Grice’s notion won’t do, both because of cases of semantic underdetermination, and because in some cases the notion we require must include content that, for Grice, is only inferable from implicatures—and hence can’t be part of what is said, by definition. It’s time to back up this assertion. Grice’s principal focus in his William James lectures was on the logical operators: he wanted to defend the claim that there is no divergence in *meaning* between the truth-functional operators of classical logic ( $\supset$ ,  $\cdot$ ,  $\vee$ , etc.) and their natural-language counterparts (‘if’/‘then’, ‘and’, ‘or’, etc.); any perceived differences can be given a pragmatic explanation in terms of implicatures. So, for example, the semantic contribution of ‘and’ to an English sentence is just the standard truth-functional definition: the conjunction is true iff both conjuncts are true. Additional elements of speaker-meaning that may be conveyed by an utterance of a conjunction (causal connection between conjuncts, temporal order) are implicata; *what is said* by the utterer is just the conjunction, and has the usual truth-conditions. I want to run my test—presenting scenarios involving deceptive utterances that may or may not be lies—to see whether Grice’s minimal ‘what is said’ is right for us.

Imagine the following scenario. Setting: Knachel household, dinnertime. Suppose that I, the *paterfamilias*, am a tyrant obsessed with healthy eating. I’ll do anything to get my kids to eat right. The kids are being their typical, finicky selves. Suppose I say this to Alice: “If you don’t finish your green beans, you won’t get a cupcake for dessert.” I brandish a tray of tempting confections. Alice, enticed by the prospect of dessert, is cowed by the threat, and quickly devours her vegetables. As she chokes down the last bite, she asks, “Can I have a cupcake now?” But I never had any intention of giving her a cupcake (do you know how much sugar is in those

things?); I just wanted to get her to eat her vegetables. “No,” I respond, “No dessert for anybody!” I make a show of dumping the cupcakes into the trash; they were just a prop. Now, young children are relatively unsophisticated in some ways, but they come equipped with an innate and finely tuned sense of fairness. Alice accuses me of lying to her. But I was ready for that accusation. I defend myself thus: “What I said was, strictly speaking, true. There were two components in my utterance: ‘Alice doesn’t finish her green beans’ and ‘Alice doesn’t get a cupcake for dessert’. As it turned out, the first is false and the second true. According to the way classical truth-functional logic defines its operators, these facts make my utterance true. For a sentence of the form ‘if p then q’, when p is false and q is true, the whole statement comes out true according to the definition of if/then; in this case, p = ‘Alice doesn’t finish her green beans’ (false), and q = ‘Alice doesn’t get a cupcake for dessert’ (true), so the whole thing is true.”

Now, setting aside concerns about giving such an explanation to a two-year-old, let’s evaluate my defense.<sup>92</sup> It’s constructed along Gricean lines: for him, any deviation between what one communicates with a conditional sentence and what one says in uttering such a sentence is purely pragmatic; there is no divergence in *meaning* between ‘if’ and ‘ $\supset$ ’, and so what is said in uttering a conditional sentence has the truth-conditions of the material conditional.<sup>93</sup> True, there’s a strong tendency, for certain utterances of conditionals, to communicate something more. In this case, my conditional was understood to have the significance of a *biconditional*: ‘If you don’t finish your green beans, you won’t get a cupcake; but if you do, you will.’ This phenomenon is referred to as “conditional perfection” in the literature. For Grice, the additional conditional (‘If you finish your green beans, then you will get a cupcake’) is an *implicatum*; it is

---

<sup>92</sup> I’m a logic teacher: suppose by the time my kids are two they’re quite familiar with the truth-functional definitions of the basic operators in classical logic. Or suppose it’s not Alice, but my wife—who’s sympathetic to the child’s concerns and resents my tyrannical obsession with health—making the accusation and evaluating my defense.

<sup>93</sup> See Grice 1989, Ch. 4

not part of what is said. So, does my defense succeed? I'm somewhat inclined to say that it doesn't, so that part of what I said was the falsehood, 'If you finish your green beans, then you'll get a cupcake'. But this may be a case about which intuitions would vary; conditionals are tricky, after all. Testable hypothesis: we could make people less disposed to call this a lie if we made a defense along these lines: "What I said was that if you *didn't* eat the green beans, then you wouldn't get a cupcake; I never said anything about what would happen if you *did* eat them. Since you did eat them, my claim about what would happen if you didn't couldn't be false; hence, I didn't lie." This explanation taps into the urge to say of conditionals, the antecedent of which is false, not that they're true, but rather that they lack truth-value or something; they're certainly not false. It's easy to elicit this intuition in an introductory logic class. There are two points to make about this modification. First, it's a departure from the test I've proposed that we use, in that the defense does not involve claiming that what I said was true, only that it wasn't false. But the test can easily be modified to accommodate the abandonment of bivalence; if saying something false is a necessary condition for having lied (as we're assuming it is), then a defense against lying can succeed merely by demonstrating that what was said wasn't false (rather than the stronger claim that it was true). Second, though, even if such a defense succeeds, it does not vindicate Grice's claim that what is said by the utterance of the conditional is nothing more than a proposition with the material conditional's truth-conditions; acceptance of something like truth-value gaps is not consistent with that program. But this is a minor problem: nobody believes anymore that the proper semantics for if/then sentences is captured by the material conditional; it may be a departure from Grice's actual position, but an account of what is said by such sentences that adopts a different semantics would retain the Gricean spirit, so long as it counts perfections of conditionals as implicata, rather than part of what is said.

So what does an example like this show? We're looking for a theoretical basis on which to make the commonsense distinction between lying and merely misleading—a distinction about which competent language users have intuitions. If a Gricean account of 'what is said' is at odds with those intuitions—even by a little bit—then we must have some degree of doubt about its suitability for our purposes. One might respond to this concern on Grice's behalf by claiming that—especially in light of the data indicating that people's intuitions are variable—what we have here is a case in which intuitions can be sensitized (to use Bach's term) and reformed. Grice—or an account in the spirit of Grice—has it right, and folks can be brought around to this view with the proper logical tutoring. That one *can* be brought around is supported by the (anecdotal) data that the change in defense above can alter people's intuitions. However (again, anecdotally), not everyone changes her mind about this case, but perhaps a Gricean could argue that one *ought* to come around. That's a bigger argument, though; it requires at least that we consider more than one example—that we examine Grice's notion from many more angles—before we make an overall judgment about its viability. To this task we now proceed.

Straight 'if/then' conditionals are tricky; there's a tendency to "perfect" them, but given all the other noise surrounding our intuitions about them, it's hard isolate the signal and say with any conviction whether such elaborations of content belong to 'what is said'. Alternative conditional locutions, on the other hand, are not so tricky. I know—again, from experience as a logic instructor—that the tendency to interpret sentences featuring 'unless' as biconditionals is *very* strong. So, consider the following scenario. You're an avid golfer and a jerk who hates spending time with his family. It's Friday, and your wife starts in with the nagging: "What are you doing this weekend? I was thinking it would be nice to take the kids to the museum." You reply, "I'm sorry dear, but I've got a ten o'clock tee time at the club. I'm afraid I won't be

joining you at the museum unless it rains.” In fact, the forecast calls for rain all weekend, so your wife is mollified. But that was your plan. You didn't want to get into a whole big thing with her right now (you're trying to enjoy a pre-dinner scotch and watch Hannity), but you have no intention of going to the museum this weekend; you'll golf if there's no rain, and find something else to do if there is. The next morning, it's pouring down rain, and as your wife is busy cleaning up after breakfast, you're headed out: “I'm going to the tavern for Bloody Marys with the boys, and then we're gonna hit the strip clubs. Don't wait up.” “You lied to me again, you creep!” she shouts, hurling a juice glass at your head. “No I didn't,” you reply quickly, ducking and scurrying toward the door, “since the sentence I uttered can be rendered as the conditional, ‘If it doesn't rain, then I won't go to the museum.’ It goes beyond what I said to perfect that conditional, as you apparently have, and add ‘If it does rain, then I will go to the museum.’ That conditional is indeed false, but it's not part of what I said. Anyway, have fun at the museum!”

Does your defense succeed? I know your wife isn't buying it; I'm not either. Testable hypothesis: presented with this scenario, normal language users would be much less inclined to accept the defense against having lied than they were in the previous example. And what we have here is a version of the same phenomenon—conditional perfection—at least according to the standard treatment of ‘unless’ as a conditional. But as any logic teacher knows, one has a hell of a time convincing students that this is the correct treatment. It's almost irresistible to supplement ‘I won't go to the museum unless it rains’ with something like ‘in which case I will go to the museum’. Untutored language users have a strong urge to treat ‘p unless q’ as ‘p iff  $\sim$ q’, and sometimes all the tutoring in the world can't shake that conviction. Take another simple example: an arch-criminal says, “I'll detonate the bomb unless you pay me one million dollars”; the people who give into his demands will be justifiably aggrieved if he detonates the bomb

anyway. The use of ‘unless’ strengthens the tendency to perfect conditionals. The Gricean treatment of this example fails to capture the intuition that what is said, in this case, includes the added content.

We can construct examples that make the minimalist, truth-functional interpretation of the locutions corresponding to other logical operators also seem dubious. Consider ‘either/or’. For Grice, the meaning of this locution is captured by the inclusive sense codified in the truth table for the wedge ( $\vee$ ): ‘either p or q’ is true if and only if it’s not the case that both p and q are false. The tendency to interpret disjunctions as exclusive is explicable in terms of implicatures: the exclusion of the possibility of both disjuncts being true is (often) implicated by the utterance of the disjunctive sentence; *what is said* in the utterance is just the inclusive sense. But consider this example. You’re buying a house, negotiating directly with the sellers. It’s an older home, and it’s got a few big problems: it needs a new roof and it needs serious foundation work. Both repairs cost big money. They agree to come down on the price a bit, and you say, “Alright. I’m going to make an offer at that price. But it will either contain a provision that you pay for the roof repairs, or a provision that you pay for the foundation repairs.” They think that sounds fine, so you tell them you’ll have your lawyer write out an official offer and send it over the next day. They’re giddy with excitement; as soon as you leave, they immediately proceed to pore over a book of fabric swatches, looking for the perfect toilet-seat cover for the master bedroom in their new McMansion in the exurbs. But this excitement was part of your plan; you wanted them to think that they’d only have to pay for one of the two major repair jobs, so they’d get excited and be more likely to cave when you presented your real offer. It arrives the next day, and to their horror you’ve asked that they pay for *both* the foundation and the roof. They accuse you of having lied. You reply that what you said was, strictly speaking, true, on Gricean grounds: what

is said by the utterance of a disjunction is the inclusive-or, which is true when both disjuncts are true.

Does the defense succeed? I think not. In some circumstances, the exclusive reading of ‘or’ is unavoidable—is part of what is said. One might object that the example above was contrived in such a way to make the exclusive interpretation more salient. It was! But this is no objection. Sure, if you vary the circumstances, an inclusive reading might be more reasonable (I haven’t said something false if I declare my intention either to eat steak or lobster for dinner, then later decide, “What the hell—surf and turf!”), but this is evidence in favor of the view that what we say is sensitive to contextual information, contrary to Grice’s picture. We can construct examples of disjunctive statements for which, because of special features of the context, the exclusive reading is inevitable: “With my tax return, I’ll either buy a lawnmower or a moped” (my return, as a matter of (lamentable) financial fact, is not going to be big enough for both). Heck, sometimes the *laws of nature* preclude an inclusive reading: “The flipped coin will come up either heads or tails.” Do we really want to say, in such cases, that *what I said* allows for the possibility of both, that the exclusion of that possibility is merely implicated?

Another concern one might have about this example—and about any one of our examples—is that what’s driving our intuitions is not linguistic, but moral considerations. We’re inclined to call this a lie because of our moral disapprobation: the character in our story is negotiating in bad faith, and we want to condemn him in the strongest possible terms. We discussed this concern above in §2.1.2, and noted that one way to address it is to control for the moral factor: invent a different scenario with the same linguistic features, but in which moral disapprobation would be absent (or diminished). So one might attempt to do this as follows: suppose I tell Alice that she can have either a cupcake or a cookie, but then I end up giving her



both. The linguistic features are the same (the exclusive understanding of ‘or’ is more salient), but the intuition that I lied goes away. Doesn’t this show that it was in fact moral considerations driving our intuitions in the house-selling example, and that we’re not entitled to draw the linguistic conclusions that we did? It does not. What this example shows is that one needs to be quite careful in constructing scenarios to test our intuitions about the difference between lying and misleading. Appropriate examples must involve deception, since we’re trying to distinguish two species of deception. The speaker has to intend to cause the hearer to have a false belief (and to make the example convincing and vivid, has to have some motivation for doing so). I can’t imagine a context in which I’d want to convince Alice that she can have a cookie or a cupcake, but not both--when in fact she can have both. So there’s no surprise we don’t have the intuition that I lied. These scenarios are hard to invent; they need to present us with a choice between the two types of deceptive act. Here’s an attempt to construct an alternative scenario for ‘or’ in which we control for the moral factor. Suppose your family is playing lawn darts—the old school kind, where the things are actually metal darts that you throw way up into the air, and are dangerous as hell—and the family dog runs into the field of play and gets skewered with a dart. The wound is mortal, and you know it, but you want to comfort your kid with a white lie: “Either we’re going to remove the dart or put Rufus down.” Your kid gets the idea that there’s hope: if you get that dart out of him, Rufus won’t be put down--because ‘or’ is exclusive. But the real plan is to do both: remove the dart (to relieve the dog’s suffering, and to ensure you can continue the game with a full complement of darts (sorry - leave that last bit out; it’ll cause the moral disapprobation we’re trying to avoid)) AND put Rufus down, because there’s no way he can make it. Despite the lack of evil intentions here, I’m still inclined to call this a lie.<sup>94</sup>

---

<sup>94</sup> Thanks to Kieran Setiya for raising the issues addressed here.

Finally, let's consider 'and'. For Grice, what is said in utterances of sentences with this operator is only a conjunction, with the usual truth-conditions. As we've noted, however, such utterances are often interpreted as conveying additional information, such as a temporal ordering of the conjuncts, or a causal connection between the two. To see whether we ought to include this content in what is said by such utterances, we can apply the usual test. Consider this scenario. Someone is asking me for advice about where to eat. He's considering a new Italian restaurant, Papa Giorgio's. I tell him the following: "I ate at Papa Giorgio's and got violently sick." Here are the relevant autobiographical details: I did, in fact, eat at the restaurant in question last Wednesday; the only time in my life that I've ever been violently sick was on my twenty-first birthday, for the usual reasons; I have a longstanding grudge against Papa Giorgio, the owner of the restaurant, and I hope to do everything I can to ruin his chances for success and steer people away from his new business (the only reason I went to the restaurant last week was so I could give him the Evil Eye; I was hungry, though, so I had a nice veal picatta while I was there). Now, suppose Papa G. hears about what I've been telling people, and accuses me of lying. I defend myself along Gricean lines: what I said, strictly speaking, is a conjunction; it's true just in case both of its conjuncts are true, and they are—I did go to the restaurant, and there was an occasion prior to my utterance on which I was violently sick.

Does my defense succeed in this case? I think it would clearly be rejected. What if I had printed the sentence in question in the newspaper, as part of a restaurant review, and Papa had sued me for libel? I think I'd be in serious trouble. The tendency to interpret utterances of sentences featuring 'and' as conveying temporal and/or causal information is called "conjunction buttressing" in the literature.<sup>95</sup> My test indicates that, in at least some circumstances, this

---

<sup>95</sup> See, e.g., Levinson 2000, p. 117

interpretation is almost irresistible—strong evidence that the intuitive notion of ‘what is said’ for which we’re searching ought to make room for such content, and evidence, once again, that Grice’s more restrictive notion will not do for our purposes.

### **2.2.3 Other limitations of the Gricean picture**

I’ve been focusing on the contribution of the logical operators to what is said in an utterance, because those were Grice’s primary focus; he developed his distinction between what is said and what is implicated largely to help him resolve disputes about their proper semantic treatment. I have argued that Grice’s understanding of ‘what is said’ is too narrow to underwrite the distinction between lying and misleading, because it produces unintuitive results for utterances involving those operators. But we needn’t limit our criticism of the Gricean picture to cases involving logical operators. If we apply our test, we will see that his minimalist conception of ‘what is said’ produces unintuitive results in a wide range of other cases.

We’ve already noted the problems raised by the phenomenon of semantic underdetermination: sometimes the sentences we utter fail to express truth-evaluable propositions; if we constrain our notion of ‘what is said’ as Grice suggests, then we will have cases in which what one says in producing an utterance will (arguably) not be truth-evaluable.<sup>96</sup> We pointed to utterances like ‘Rose isn’t tall enough’, which seem to require pragmatic input in order to express a proposition and be truth-evaluable. The notion of ‘what is said’ that we require to underwrite a distinction between lying and misleading must be one according to which what we say has a truth-value, and so it seems in these cases Pragmatics contributes to (our favored

---

<sup>96</sup> I include the parenthetical because some people deny the phenomenon. We will talk about this soon.

sense of) what is said. Bach refers to the process by which we flesh out such underdetermined utterances as *completion*.<sup>97</sup> He notes that there is another, complementary process at work in some other cases, which he calls *expansion*. In these instances, we have an utterance of a sentence that *is* literally truth-evaluable, but whose truth-conditions differ from those communicated. Our examples involving the logical operators are arguably instances of this phenomenon, but there are many more. I will complete our discussion of Grice, then, by briefly examining some of these cases, using my test involving intuitions about lying.

We've already seen one. Recall the example of the neighbor who's a bad cook. We declined his offer of gumbo with the utterance, "No thanks. I've had lunch." A minimal, Gricean, account of what is said in this case would deliver truth-conditions according to which the utterance is vindicated just in case the speaker has had lunch at least once in his life. We saw that a defense against lying on such grounds would fail; what is said in such a case is an expanded proposition: 'I've had lunch *today*.' In his 1994, Bach has other examples of expansion, which we can subject to similar analyses. We're asked to consider a mother consoling a child who's cut his finger; she says, "You're not going to die." Now suppose that later in life, the child is diagnosed with a life-threatening disease, and he accuses his mother of having lied to him that day when he cut his finger. She could easily defend herself: "I didn't say you were immortal; what I said was that you wouldn't die *from that cut*." That is, what she says to the child is an expanded proposition, whose truth-conditions differ from the minimal semantic interpretation.<sup>98</sup>

Another example from Bach: an utterance of 'France is hexagonal'. Suppose the context of the

---

<sup>97</sup> See, e.g., Bach 1994. Bach, we should note, disagrees that this process constitutes a pragmatic intrusion into the determination of what is said; he tries to defend a more minimal, neo-Gricean conception of 'what is said'. We'll return to this below.

<sup>98</sup> Here's an instance in which semantic minimalism and one of the Gricean diagnostics make different predictions. The minimal content is a claim of immortality; but the extra content (from this cut) is *not* cancellable, since it's entailed by the minimal content. Note, however, that it *is* reinforcing: there's no sense of redundancy if the utterance is 'You're not going to die from that cut'.

utterance is this: I'm trying to help a child unfamiliar with European geography pick out France on a map. Now, if the child got out his protractor and—failing to measure the requisite six  $120^\circ$  angles—accused me of lying, I could easily defend myself by pointing out that I didn't say it was *exactly* hexagonal, only that it was *roughly* so.

Or consider typical uses of quantifiers. If I tell Rose, "Everybody's coming to your eighth birthday party," and all the invitees, plus her grandparents, uncles, and mom and dad all show up, she's in no position to accuse me of lying because it's not the case that everybody in the *world* showed up. But on a minimalist understanding, what I said in my utterance has just that truth-condition. If Alice needs a beverage (to go with her green beans (or cupcake)), and I tell her, "There's some milk in the fridge," I'm a liar if it turns out that all there is in the refrigerator, milk-wise, is a tiny puddle of spilt milk that dripped down into the crisper. It's no defense to fall back on a minimal semantics for the existential quantifier and make myself out to be a truth-teller in virtue of the fact that the fridge was not, in fact, entirely devoid of milk. We routinely communicate using quantifiers to refer to domains without explicitly identifying them, and when we do so, we *say* things about those domains.

I think these cases are decisive. Unlike the examples involving the logical operators, for which I think—as I noted—that it would be reasonable to expect (depending on the case) varying degrees of disagreement about the intuitions involved, there seems to me no disputing the conclusions about these examples of expansion. The notion of 'what is said' that we use to distinguish lies from non-lies is one according to which Pragmatics makes contributions to content. Grice's minimal notion does not allow this kind of robust contextual contribution, and so it cannot underwrite our distinction. Whether a minimal, purely semantic notion of 'what is

said’ plays some other role in our linguistic practice and/or theorizing is, for now, an open question. We will return to it later.

### 2.3 GENERALIZED CONVERSATIONAL IMPLICATURES

Continuing our search for a conception of ‘what is said’ suitable for our purposes, we may follow up a lead suggested by the previous section. It seemed clear that completions and expansions, to use Bach’s terminology, were part of what is said in the sense we need. Maybe we should look to Bach for the stable, theoretically sophisticated notion of ‘what is said’. Recall, he refers to the contents generated by completion and expansion as ‘implicitures’. Perhaps if we just add these on to literal meanings, we have what we need.

There are two problems with this proposal. First, Bach himself would strenuously deny that implicitures are part of what is said. He conceives of them as a level of content *between* what is said and what is implicated; it is a distinct level of meaning. And he argues forcefully and at length against the kind of expanded notion of ‘what is said’ that I’ve been advocating.<sup>99</sup> But even if we set aside Bach’s own personal qualms with our appropriating his terminology, it’s not clear that the notion of implicature is as precisely delimited as we might hope for it to be. As Levinson observes, Bach “fails to give us a clear boundary between implicature and implicature.”<sup>100</sup> Levinson doesn’t provide much of an argument to support this claim, but I think we can make some observations to substantiate it. Bach says the following:

Although both implicature and implicature go beyond what is explicit in the utterance, they do so in different ways. An implicature is completely separate from what is said and

---

<sup>99</sup> See Bach 1994, 1997, 2001, 2002, 2005, 2007. For example. We’ll return to his concerns.

<sup>100</sup> Levinson 2000, p. 197

is inferred from it (more precisely, from the saying of it). What is said is one proposition and what is communicated in addition to that is a conceptually independent proposition, a proposition with perhaps no constituents in common with what is said. ...In contrast, implicatures are built up from the explicit content of the utterance by conceptual strengthening..., which yields what would have been made fully explicit if the appropriate lexical material had been included in the utterance.<sup>101</sup>

These are suggestive remarks, and they do a fine job of distinguishing the phenomena Bach has in mind from Grice's typical examples of particularized conversational implicatures. In this exchange

A: Smith doesn't seem to have a girlfriend these days.

B: He has been paying a lot of visits to New York lately.<sup>102</sup>

the implicatum—that maybe Smith has a girlfriend in New York—is intuitively “separate from” and “conceptually independent” of what B says. But separateness and conceptual independence are rather vague notions; can they be relied upon to sustain a sharp distinction between implicatures and implicatures? Bach's gloss of “conceptual strengthening” as making explicit content that would have been, but for the absence of “appropriate lexical material,” might be thought to help make things more precise. But consider another of Grice's classic examples, the utterance of ‘There is a garage round the corner’ to the man out of petrol. We could state the implicatum simply by adding appropriate lexical material: ‘There is a {n open} garage {with petrol for sale} round the corner.’ That seems to be “built up from explicit content” of the original utterance. Is it an implicature, then? The same problem arises, perhaps more acutely, if we try to distinguish implicatures from *generalized* conversational implicatures. Consider an example we've already dwelt on: utterances of sentences like ‘Some logicians are bald’ typically implicate propositions like ‘Some {but not all} logicians are bald’. Again, the extra content seems to be built up from the original by the addition of lexical material (in braces), and so it's

---

<sup>101</sup> Bach 1994, p. 11

<sup>102</sup> Grice 1989, p. 32

hard decide whether to call this an implicature or an implicature. Given the close connection between GCIs and lexicalization—they arise typically with certain forms of words—we would expect Bach’s suggested method of demarcation to leave us with this kind of indeterminacy quite regularly. Are the expanded contents in statements of the form ‘A or B {but not both}’ and ‘A and {then} B’, for example, implicatures, implicatures, or both? Unlike the category of *conventional* implicatures, which Bach thinks is empty<sup>103</sup>, GCIs aren’t on the chopping block—and there’s no indication that he wants to assimilate them into the implicature category.

But if Bach’s notion of implicature is somewhat wanting in terms of clear criteria for demarcation, the concept of Generalized Conversational Implicature is not. In Levinson 2000 (*Presumptive Meanings: The Theory of Generalized Conversational Implicatures*), we have an exhaustive defense and taxonomy of GCIs. If we’re looking for a theoretically sophisticated account on which to base our preferred understanding of ‘what is said’, perhaps this is it.

### 2.3.1 Categories of GCIs

Grice, recall, distinguished implicatures in this category from their particularized cousins on the basis of their *default* character: they arise, associated with particular forms of words, normally, and are only cancelled in the presence of “special circumstances.”<sup>104</sup> This feature of GCIs makes them difficult to distinguish from the conventional meanings of the words and phrases that typically give rise to them. They are therefore prime candidates for assimilation into an intuitive understanding of what is said by an utterance. We’ve already seen—in cases involving the logical connectives and in Bach’s examples of implicatures—that this sort of supposedly

---

<sup>103</sup> See Bach 1994 §6, Bach 1999 (“The Myth of the Conventional Implicature”)

<sup>104</sup> Grice 1989, p. 37



implicatural content arguably sits more comfortably on the same side of the meaning-divide as what is said, at least in the sense of that notion that we are seeking to elucidate.

While Grice is characteristically sketchy in his attempts to circumscribe the domain of GCIs, Levinson is the opposite: painstaking and thorough. He bases his taxonomy on some of Grice's suggestive remarks, extending and elaborating them to present a comprehensive system. He divides GCIs into three types, based on the conversational principles that give rise to them. He reduces the number of maxims to three (Grice listed more), and refers to them as *heuristics*; the thinking here is that these rules of thumb allow us to streamline the communication process, speeding up interpretation by providing an inferential—but automatic, often sub-conscious—bridge between relatively impoverished linguistic inputs and richer, expanded interpretive outputs. GCIs help increase the efficiency of communication by licensing inferences to default, presumptive meanings for common expressions; this is why they're only cancelled in abnormal circumstances.

Levinson calls his three heuristics the Q, I, and M Principles, based on/inspired by Grice's first maxim of Quantity ("Make your contribution as informative as required"<sup>105</sup>), second maxim of Quantity ("Do not make your contribution *more* informative than is required"<sup>106</sup>), and maxims of Manner ("Be perspicuous"; "avoid obscurity of expression"; "avoid prolixity"<sup>107</sup>) respectively. Let's run through Levinson's principles in order:

*The Q Principle.* The pithiest statement of this heuristic is "What isn't said, isn't."<sup>108</sup> More specifically, it's a maxim according to which the speaker is supposed not to "provide a statement informationally weaker than [her] knowledge of the world allows", to make the

---

<sup>105</sup> Levinson 2000, p. 35

<sup>106</sup> *Ibid.*, p. 37

<sup>107</sup> *Ibid.*, p. 38

<sup>108</sup> *Ibid.*, p. 31

strongest possible statement “consistent with the facts.”<sup>109</sup> Hearers can then make inferences on the basis of the assumption that the speaker is abiding by this rule. This results in implicatures of the so-called “scalar” variety, for example, which are based on scales ranging from stronger to weaker expressions (e.g., <all, most, many, some>), with the use of a weaker expression implicating that the stronger one(s) could not have been used (some +> not all).<sup>110</sup>

*The I Principle.* The pithy version: “What is simply described is stereotypically exemplified.”<sup>111</sup> More specifically, according to this maxim, the speaker ought to “produce the minimal linguistic information sufficient to achieve [her] communicational ends” and the hearer ought to “[a]mplify [or enrich] the informational content of the speaker’s utterance, by finding the most *specific* interpretation” consistent with the speaker’s presumed point.<sup>112</sup> Under the rubric of this heuristic, Levinson gathers a motley assortment of phenomena, some of which we’ve already seen: conditional perfection (if +> only if), conjunction buttressing (and +> in that order), and many others.

*The M Principle.* “What’s said in an abnormal way, isn’t normal; or Marked message indicates marked situation.”<sup>113</sup> The demands this heuristic imposes on speakers and hearers are, I think, pretty clear: speakers should choose unusual modes of expression if and only if the situation is unusual, and hearers should make inferences about the normality of the indicated situation accordingly. There are as many implicatures arising from this principle as there are marked expressions, so they’re somewhat difficult to categorize further, but examples include an utterance of “Sally was knitting. Occasionally, *the woman* looked out the window.” implicating

---

<sup>109</sup> *Ibid.*, p. 76

<sup>110</sup> I’ll borrow Levinson’s convention and use the symbol ‘+>’ to indicate that the expression on the left implicates that on the right. (I’m being a bit loose here: there should be utterances on the left and propositions on the right; you get the idea, though.)

<sup>111</sup> *Ibid.*, p. 32

<sup>112</sup> *Ibid.*, p. 114

<sup>113</sup> *Ibid.*, p. 33

that ‘the woman’ refers to someone other than Sally; the choice of ‘tome’ rather than ‘book’ in “He was reading a tome” implicating that the volume in question is unusually large; the choice of ‘caused ... to die’ rather than ‘killed’ in “The Spaniards caused the Aztecs to die” implicating that the way in which the Spaniards killed them was unusual and indirect (disease?).<sup>114</sup>

So these are the categories of GCIs. What remains for us to see is whether we ought to include these species of conveyed meaning under the category ‘what is said’, as we’re understanding that notion. That is, we need to ask whether content supposedly implicated according to these heuristics should actually be included as part of the truth-conditional content of speakers’ utterances when the issue of whether they lied or merely misled is on the table. Let’s examine Levinson’s three categories one at a time.

### 2.3.2 Q-Implicatures

We’ve already seen some example of scalar implicatures, and made the case for one of them that at least in some circumstances it makes sense to include the supposed implicatum as part of what is said in an utterance. The understanding of sentences featuring ‘or’ as excluding the possibility of both disjuncts arises, on Levinson’s model, because of the existence of a scale <and, or> and an inference from the use of the weaker locution to the negation of the stronger, in accordance with the Q heuristic: what isn’t said, isn’t; therefore, not both. As the example of the dishonest real estate negotiator (§1.3.1.2) showed, sometimes this interpretational move is so strongly suggested by the context that it makes sense, for our purposes, to include the exclusion of the possibility of both disjuncts as part of what the speaker said. This is at least partially because the

---

<sup>114</sup> *Ibid.*, pp. 138 - 141

very nature of deceptive situations is such that it *matters* to the hearer whether or not the possibility of both is open, and so the interpretive move to exclusive-or is strongly motivated. This is promising. Perhaps we can conclude that, in general, Q-implicatures are prime candidates for inclusion in the notion of ‘what is said’ that we seek. We’ll have to examine more of them to see.

Let’s consider the other scalar implicature that we’ve already mentioned—that from utterances featuring ‘some’ to propositions whose truth-conditions include the negation of universal affirmation—not all. The locutions ‘some’ and ‘all’ fall on a scale, and the former is weaker, so an utterance using it implicates the negation of the latter, by the Q principle. Following our usual procedure, we need to imagine a scenario in which the speaker is being deceptive, so that the inference to the GCI makes a difference, matters to the hearer. Suppose I bake up a batch of my Grandma’s famous cinnamon rolls, my favorite dessert. I and the girls scarf down a couple straight out of the oven, and then save the rest of the batch for tomorrow. That next day, I come home from work and find that the rolls have all been eaten! I set about interrogating the family, and ask Rose about the missing treats. “I ate some of the cinnamon rolls,” she admits. But suppose the fact of the matter is that she ate *all* of them. After interrogating my wife and little Alice, neither of whom snuck so much as a bite, I confront Rose with her deception: “You lied.” She defends herself in the usual way: “No dad, I didn’t lie. What I said was, strictly speaking, true: in fact I did eat at least one of the rolls, so the truth-conditions of ‘I ate some of them’ are satisfied.”<sup>115</sup> Does this defense succeed?

I’m not sure what to say about this case; my guess is that opinions will differ. I think we might expect some degree of variation in intuitions about different instances of the some +> not

---

<sup>115</sup> I should have known kindergarten was too early to introduce quantifier logic.

all implicature, depending on features peculiar to the scenario. Consider another example. Suppose we're living in a dystopian future USA in which Secret Police agents periodically troll for dissidents by calling random citizens in for questioning. I'm so called, and respond to questions about my neighbors with an utterance of "Some of my neighbors are law-abiding, patriotic citizens." In fact, I know that *all* of them fit that description, but I'm nursing a grudge against the lot of them: I overheard them making disparaging remarks about the quality of my noodle salad at last summer's block party. Subsequent to my testimony, everybody's rounded up in the middle of the night and subjected to all manner of harassment and "enhanced" interrogation techniques. When they hear about what I said, both my neighbors *and* the Secret Police accuse me of lying, and I reply that since at least one of my neighbors was as described, what I said was strictly speaking true.

When I present people with this case, they're less-inclined to reject my defense than they are to reject Rose's above. Why? I suspect that one reason is that the facts about which I'm testifying in the latter case—the political views and activities of all my neighbors—are facts about which it's quite possible for me to be ignorant. In Rose's case, the fact in question—whether or not she ate all the cinnamon rolls—is *not* a fact about which it's possible for her to be ignorant. It's relatively rare, in normal speech, to use the word 'some' and not implicate the particular negative, but some cases in which it is appropriate are those in which one is ignorant of all the facts. When I'm teaching introductory logic, and illustrating this minimal use of 'some', I ask my students to imagine visiting home on spring break and telling their mom about the classes they're taking; when she hears about logic, she asks (for some reason), "I've always wondered, what are logic teachers like?" I tell the students that they might respond, given what they know about me, "Some logic teachers are jerks." And for all they know, I point out, *all* of

us are jerks; maybe it's just something about logic that attracts unpleasant people. This leads to another observation: people's intuitions in these cases probably vary partially as a function of the extent of their exposure to logic.

Let's look at some other Q-implicatures. There are various scales of strong-to-weak expressions. Consider the deontic modals: being *forbidden*, *permitted*, or *obligated* to do thus-and-so. Obligation is stronger than permission, so we have a GCI from the latter to the former. Let's cook up a scenario: An old French acquaintance of yours from the Vietnam War sent his son to America for college; the son liked it here so much, he decided to stay; he's just been naturalized—a newly minted citizen. Still, he doesn't really know much about how things work here in the USA. His dad told him to seek you out as an advisor. Big mistake: because of their perfidy in the build-up to the Iraq War, you've developed a strong antipathy to all things French; plus, he didn't know it, but you always hated this particular Frenchman's guts. Anyway, the son comes to you for advice. Suppose he's 22 years old. Hoping to lead the kid astray, you say to him, "Oh, you should know that (male) citizens between the ages of 18 and 25 are permitted to register with the Selective Service Agency, in case there's a military draft." The little punk considers himself a lover, not a fighter, and opts to forgo that little perk of citizenship. But later, when he's at the DMV and they refuse to give him a license because he's not registered (as is, of course, *required* by law)<sup>116</sup>, he comes back to you, angry, and accuses you of lying. "I didn't lie," you reply, "since what I said was, strictly speaking, true: the law does not forbid your registering for the draft; it is permitted." This strikes me as an extraordinarily lame response. I think it's reasonable to include a lack of obligation as part of the content of your utterance.

---

<sup>116</sup> This is how it works in some states.

Testable hypothesis: people would be more inclined to reject defenses against lying in permitted +> not obligated cases than in some +> not all cases.

Number-terms fall on a scale, of course. The tendency to understand utterances of a particular numeral as conveying that exactly-that-many are involved can arguably be treated as a GCI.<sup>117</sup> Use of ‘three’ literally means *at least three*, but implicates *at most three*. I think this inference is so well-entrenched in usage that the supposed implicatum ought to be included as part of the content of what is said. Here’s a scenario using our usual test. In China, according to the “Family Planning Policy,” it’s illegal for certain couples to have more than one child. Suppose my wife and I were a Chinese couple subject to this restriction, and the Party’s Family Planning Compliance Inspector came to our door (I don’t know how this actually works in China; bear with me). I report to this official: “We have one child.” He checks a box on his clipboard and moves on to the next house. Later, though, he sees me taking my two girls out for ice cream. “You said you only had one kid, you liar,” he says. Instead of trying to pass off Alice as the neighbor’s girl or something, I respond, “What I said was true: we have at least one child.” This defense is simply not going to cut it.

Other phenomena that Levinson classifies as GCIs strike me as harder to judge; my intuitions are murky, and I would expect wide disagreement about them. Take the scale <hot, warm>. Suppose I own a hot tub (I don’t), and my wife, who’s about to have a soak, asks me how the water is. I say, “It’s warm.” Background: the heating mechanism for the thing has been on the fritz, so the water’s been unheated for some time. I had just been tinkering with it, and the results of my efforts have been an over-correction: now it’s pumping out energy like crazy,

---

<sup>117</sup> Arguably. Levinson himself is somewhat reluctant (see §2.2.2.2.2 in Levinson 2000). There’s a case to be made that the relevant phenomena should be treated as semantic ambiguity or non-specificity (see Atlas 2005 and Horn 2009).

heating the water to near-boiling temperatures. My wife's been bugging me lately about fixing it, and I wanted to get back at her. And did I! To the tune of second-degree burns over 60% of her body; she'll think twice before she nags *me* again. She accuses me of lying, but I maintain that I merely misled her; the water was, in fact, warm. Part of me wants to say that my defense succeeds—the falsity seems cancellable: “It’s warm; in fact, it’s boiling!”—but part of me wants to agree that I lied. If I surveyed people, they might be inclined to call me a liar, but I worry that would mostly reflect moral disapprobation (there’s no denying that my actions are despicable). If we switched the hot tub to a bowl of soup, and the injuries were just a little scalding on the roof of her mouth, I think the intuition would be to characterize my statement as a lie might diminish; it becomes, arguably, more of a misleading understatement. This is a hard case.

An easier one: Levinson also mentions the scale <know, believe>. If I say that I believe that *p*, I implicate that I don’t know that *p*. I can find in my imagination no scenario in which it’s tempting to include that implicatum as part of what is said. If I tell someone I believe something, but in fact I know it, it just seems bizarre for them to accuse me of lying.

Implications arising from the use of different tenses can be analyzed as GCIs of the Q variety. There are cases for which I would include such supposed implicata as part of what is said. Let’s go back to our dystopian future USA, with the Secret Police and all that. Suppose that things have shifted so far to the right politically that being a member of the Democratic Party is as bad as being a Communist was in the ’50s. I’m brought in for questioning, and I’m sweating it, because I’m a Democrat (I attend secret meetings, distribute underground newsletters, the whole nine yards), and I have been for decades. But I know that the authorities go easier on you if you’ve mended your ways and abrogated your Democratic affiliations. So I tell my interrogators, “I used to be a Democrat.” If they catch me the next day handing out copies of my



pamphlet, “Jimmy Carter Wasn’t So Bad,” and accuse me of lying, can I successfully defend myself by claiming that what I said was, strictly speaking, true, since there was a time before my utterance at which I was a Democrat? I’m inclined to say that my defense fails, and that part of what I said was that I’m not a Democrat now. I’m somewhat less strongly inclined to reject the defense, though, if I alter the original utterance: “I was a Democrat twenty years ago.” And my conviction weakens further reflecting on the simple “I was a Democrat.” But it strengthens again somewhat, considering “I had been a Democrat.” So we have, for the same basic phenomenon—Q-implicatures arising from the use of tense—a range of intuitions about the relationship between what is said and what is implicated.

What have we learned? Far from providing a firm theoretical grounding for a clear and sharp distinction between lying and (mere) misleading, Q-implicatures instead seem to reveal a range of cases falling on a spectrum between those two extremes. Perhaps there is no sharp distinction to be made between lying and misleading; or if there is, the basis of that distinction must involve something other than GCIs. It’s looking impossible to classify these default inferences uniformly; their behavior in lying/misleading scenarios is far too erratic. A quick examination of I- and M-implicatures will confirm this observation.

### **2.3.3 I-Implicatures**

We’ve already seen two of the phenomena Levinson classifies as I-implicatures: conditional perfection and conjunction buttressing (“If you don’t finish your green beans, you won’t get a cupcake” and “I ate at Papa Giorgio’s and got violently sick”). Those two cases alone provide evidence of the varying degrees to which it seems appropriate to include implicata in ‘what is

said’: the intuition to call my conjunction a lie was much stronger. Other I-implicatures confirm this variability.

Recall, I-implicatures, as Levinson puts it, “go in just the reverse direction to that in which Q-implicatures tend. ... [They] are inferences from *the lack of further specification to the lack of need for it*, whereas *Q-implicatures are inferences from the lack of informational richness to the speaker’s inability to provide it*.”<sup>118</sup> The I heuristic instructs the speaker to say as little as necessary, relying on the hearer to fill in the gaps according to how things usually, stereotypically go. So one inference of this type Levinson calls “Inference to stereotype,” and his example is an utterance of “John said ‘Hello’ to the secretary and then he smiled” communicating that John said ‘Hello’ to the *female* secretary and *John* smiled.<sup>119</sup> Secretaries are stereotypically female, and this information settles the reference of ‘he’. Now, even Grice allows that reference-fixing is part of the determination of what is said (though the fact that implicatural inference is involved in the process creates an awkward consequence for his account—according to which ‘what is said’ is supposed to be the *input*, not the *output* of implicature-processing—the Levinson calls “Grice’s Circle”), and so cases in which inference to stereotype is active at that level are unproblematic for our purposes, since there’s no tendency to begin with to separate the implicature from what is said. It’s much more difficult in other instances to see how one could genuinely *lie* by relying on such inferences, though. Kindergarten teachers are overwhelmingly female, but it’s hard to see how an utterance employing the expression ‘the kindergarten teacher’ could be rendered *false* by the teacher’s being male.

But for some other I-inferences, there are cases in which the inferred content ought to be included as part of what is said. Take “negative strengthening,” where an utterance of, e.g., “I

---

<sup>118</sup> Levinson 2000, p. 116

<sup>119</sup> *Ibid.*, p. 117

don't like Alice" implicates that I positively dislike her.<sup>120, 121</sup> Imagine I'm trying to ingratiate myself with some (potential) new friends; they're all big sports fans, but I know next to nothing about sports. They're going on about Lance Armstrong and what a heel he was. I've never even heard of the guy, but everybody really seems to hate him, so I chime in: "I don't like Lance Armstrong." I lied to fit in, one might say. "No," I might respond, "I didn't. What I said was, strictly speaking, true. Since I've never met him and don't even know who he is, I don't have feelings one way or the other about Lance Armstrong: I neither like him nor dislike him. It's not the case that I like him, so what I said was true." This strikes me as pretty weak. Maybe even weaker: suppose they're all Lance Armstrong apologists; despite his admissions to Oprah, they still think he's innocent. I opine, "I don't think he did it, either," not knowing who he is—or even what *it* is. Can I defend the truth of my statement based on this level of ignorance?

An even more striking case: the "mirror maxim," according to which, e.g., "Harry and Sue bought a piano" implicates that they bought it together, rather than each of them buying one.<sup>122</sup> Suppose you're homosexual, but you're lucky enough to live in one of the states where gay marriage is legal. You're unlucky, though, in that your boss is a closed-minded jerk; if he ever found out about your sexual orientation, you'd lose your job. It's time for the annual company Christmas party, a big fancy black-tie affair that everyone attends--usually with their significant other. You're determined to throw your boss off, so (with your husband's somewhat reluctant blessing<sup>123</sup>) you get your sympathetic next-door neighbor's wife, Tina, to accompany you as a beard. At the party, you tell anyone who'll listen (especially the boss), "Tina and I are

---

<sup>120</sup> *Ibid.* Not *that* Alice; some other Alice. It's Levinson's example. He'd *love* my Alice.

<sup>121</sup> Some species of negative strengthening are called "NEG-raising," a term coined in Horn 1989.

<sup>122</sup> *Ibid.*

<sup>123</sup> I forgot to mention: you're a man.

married.” You lied! It’s no defense to point out that, since you have a spouse and Tina has a spouse, what you said is, strictly speaking, true.

#### 2.3.4 M-Implicatures

These sorts of inferences arise in the presence or absence of marked expressions—unusual ways of putting things. Roughly, unusual expressions indicate unusual circumstances; normal expressions indicate normal circumstances. Unlike their I-heuristic cousins, these implicatures do not often seem suitable for assimilation into what is said. Let’s consider two representative examples: one in which someone tries to conceal the unusual nature of a situation by using an ordinary mode of expression, and one in which someone tries to conceal the straightforwardness of a situation by using an inappropriately marked mode of expression.

First, consider Dirty Harry. His Lieutenant’s finally had it: he’s going to start garnishing Callahan’s wages to pay for all the property he destroys. It’s hard enough to get by on a meager Inspector’s salary already, so Harry can’t allow his pay to be cut. But he’s not about to change his investigative *modus operandi*; it may be “unconventional”, but he gets *results*, dammit. So he decides to deceive his Lieutenant. As part of the narrative in his latest report, Inspector Callahan has written this sentence: “I opened the door.” What actually happened was: Dirty Harry took out his .44 Magnum and fired, putting a hole the size of a basketball where the lock and doorknob used to be, then kicked the door completely off its hinges. Did he lie in the report? I’m inclined to call this a classic case of misleading with the truth. He may have done it unconventionally, but he did open the door. If it *is* a lie, it’s at least not *clearly* a lie. But suppose we alter the scenario slightly: instead of writing “I opened the door,” let’s consider the sentence “I unlocked the door and opened it.” Now my intuitions aren’t so clear. This sentence seems worse, if not an outright

lie, at least *closer* to being one. Why? Is it that the idea of unlocking a door is more closely associated with a typical, normal procedure than the idea of opening one? Is it the conjunction of *two* dishonest characterizations rather than just one? I can't say for sure. The only lesson I want to extract from this is that our intuitions about such cases seem mutable, variable—sensitive to small changes in the circumstances. This is an interesting datum.

Next, let's take up the opposite kind of case, in which the utterance is marked, but the situation normal. Consider the nameless outlaw narrator of Bob Marley's "I Shot the Sherriff." Suppose that he's less ingenuous than he is in the song. He wants to appear forthcoming, but he wants his listener to believe, falsely, that he's not guilty of murder. So instead of confessing as he does in the song, he says, "I and my gun were part of the causal chain that ended in the death of the Sherriff. But I and my gun were *not* part of the causal chain that ended in the death of the Deputy." This is a marked way of putting it. It calls to mind elaborate scenarios, like Donald Davidson's stampede of wild pigs.<sup>124</sup> Now suppose (contrary to fact probably; maybe it's his wife) the listener has every reason to believe that the man is being a cooperative, maxim following communicator. According to the M-principle, then, this marked way of confessing implicates that the situation was not a simple, ordinary shooting. But suppose it was: as a matter of fact, the guy just walked up behind the Sherriff in the local saloon, pulled out his pistol, and shot him in the back. (He really didn't have anything to do with the Deputy, though.) Again, I'm disinclined to classify this statement as a lie. It's hard to justify including the lack-of-normalcy as part of what is said. If there's a degree to which things are lies or not—a continuum of

---

<sup>124</sup> Davidson 2001, essay 4. In case the reader's unfamiliar: Smith is trying to shoot Jones (or maybe it's the other way round; I forget), misses badly, but succeeds in setting off a stampede of a nearby herd of wild pigs, who trample Jones to death.

possibilities between the extremes of lies and acts of (mere) misleading—I'd put this one even closer to the 'misleading' side than the Dirty Harry case.

### **2.3.5 Final reflections on GCIs**

So what have we learned? Generalized Conversational Implicatures, despite our initial tentative hopes, do not provide a unified set of phenomena—at least with respect to the distinction that is our focus, between lying and misleading. We have been in search of a theoretically rigorous basis on which to draw this commonsense distinction, which depends crucially on the notion of what is said (strictly speaking) by a speaker in making an utterance: if what he said, strictly speaking, is true, he didn't lie. Some reflection led us to conclude that a notion of 'what is said' suitable for our purposes could not be minimal; it must accommodate the observation that features of the context of utterance play an important role in determining what we say—a role that goes beyond merely determining the references of indexical expressions. We noted that GCIs, as default pragmatic inferences that are closely tied to specific linguistic forms, and are therefore close kin with the conventional meanings of those forms, might be prime candidates for assimilation into the more expansive conception of 'what is said' that we were seeking. Perhaps, our thinking went, if we simply add the GCIs to the semantic contribution of expressions, we'd have a stable foundation for our target distinction—and a theoretically rigorous one, at that, since Levinson 2000 provides an elaborate taxonomy of GCIs with a thoroughly argued, evidence-based rationale.

But, alas, our hopes were dashed. When we got down to concrete cases, and applied our test designed to elicit intuitions about 'what is said' based on the contemplation of scenarios near the border between lying and misleading, we discovered not unity, but variety. Some GCIs were

clearly strong candidates for inclusion as part of the content of what was said by speakers; others were clearly not suitable for such inclusion. And perhaps most disturbingly, some species of GCI exhibited flexibility, ambiguity, on this central question. We weren't sure, in some cases, *what* our intuitions were about exactly what was said; it seemed there were in-between cases, acts not easily categorized as either lies or (mere) deceptions.

Faced with this failure to find a theoretical basis in GCIs for a sharp distinction between lying and misleading, we have two choices: abandon the search, or press on. There is some reason for thinking that the proper choice is the former. Our discovery of variegated intuitions about what is said in various circumstances could be taken as evidence that there is no precisely definable notion suitable for our purposes. Our commonsense distinction between lying and misleading may be inherently blurry and ambiguous, not clear and distinct. We may want to embrace the indeterminacy, and try to search for an explanation of why we might expect it. In fact I think this is the proper course. But before we proceed along it, I want to consider the alternative one last time. For perhaps the reason our intuitions were confused was not that it's impossible to draw a sharp boundary between lying and misleading, but because we were trying to draw the boundary using the wrong tools. GCIs cut across the distinction, but perhaps some other linguistic theory can provide us with the tools to sensitize our intuitions in such a way that a sharp boundary can be drawn. In that spirit of open-minded inquiry, I want to examine Relevance Theory.

## 2.4 RELEVANCE THEORY

Relevance Theory (RT) is an attempt to build on the basic Gricean model of communication, incorporating subsequent insights that the extent of pragmatic contributions to the interpretive process extend beyond the strict limits Grice originally set.<sup>125</sup> Semantics provides a logical form (LF) for utterances, but this is in general not enough to give us the explicit content of the utterance, which is called its “explicature.” Since explicatures may have content that goes beyond Grice’s minimal ‘what is said’, we have in RT another potential candidate to provide a sound theoretical basis for the more expansive understanding of that notion that we seek—one that can underwrite the distinction between lying and misleading. We shall briefly consider whether Relevance Theory can give us what we want.

### 2.4.1 The basics

Relevance Theory is ambitious. It aims to provide a comprehensive account of communication that is based on plausible and empirically testable assumptions about cognitive processing mechanisms. The central assumption is the Cognitive Principle of Relevance: “Human cognition tends to be geared to the maximization of Relevance.”<sup>126</sup> Relevance is a property that cognitive inputs—principally, for our concerns, utterances, but also perceptions, etc.—can have to a greater or lesser degree depending on their cognitive effects and on the effort required to process them. The relevance of utterances will vary directly with the degree to which they have “positive

---

<sup>125</sup> Seminal texts include Sperber and Wilson 1986, Carston 2002.

<sup>126</sup> Sperber and Wilson 2004, p. 610



cognitive effects” like “answering a question [one] has in mind, improving [one’s] knowledge on a certain topic, settling a doubt,” and so on; relevance will vary inversely with “the effort of perception, memory, and inferences required [to process them].”<sup>127</sup>

We saw that Levinson, in his theory of GCIs, streamlined Grice’s framework of conversational maxims, identifying only 3 principles (the Q-, I-, and M-heuristics) involved in the inference of implicata. Relevance Theory is even more austere, with only a single principle governing communication, the Communicative Principle of Relevance: “Every [utterance] communicates the presumption of its own optimal relevance.”<sup>128</sup> From the presumption that speakers are making relevance-maximizing conversational contributions, hearers infer the contents of their utterances. The important difference between Grice and Levinson, on the one hand, and relevance theorists on the other, is that in RT, the maxim is not limited to aiding in the derivation of implicatures—contents that go beyond a speaker’s primary, literal meaning. The Principle of Relevance facilitates understanding at *all* levels, including the primary level, at which we find what speakers communicate explicitly—so-called “explicatures.” This content often goes beyond the more minimal Gricean ‘what is said’, and so can include what Grice and his followers would prefer to call implicatures (especially GCIs), or what Bach would call implicatures.

Sperber and Wilson spell out a multi-step comprehension procedure that hearers are supposed to follow as they recover speakers’ meanings: it’s organized into subtasks, including the decoding of the logical form of utterances (the contribution of semantics), the formation of competing hypotheses for the content of explicatures (which are “development[s] of a logical

---

<sup>127</sup> *Ibid.*, pp. 608 – 609

<sup>128</sup> Sperber and Wilson, p. 158. I limit the principle to utterances; Sperber and Wilson speak in terms of “act[s] of ostensive communication”—a more general notion, the explication of which needn’t detain us.

form”)<sup>129</sup>, the evaluation of these in relation to the Principle of Relevance, the recovery of implicatures that go beyond explicit content, and so on. The tasks can be accomplished in parallel, not linearly, but the process *is* supposed to be deductive. They distinguish between weak and strong implicatures, with only the latter necessary to meet the demands of maximum relevance (among the former are various forms of loose talk and metaphor). And within explicit content, they distinguish “basic explicatures” from “higher-order explicatures,” all of which can be conveyed by a single utterance. So, for example, in an exchange between Peter and Mary, in which Peter asks Mary if she’ll pay him the money she owes by Tuesday, and Mary responds, “I will pay it back by then,” we have for a basic explicature the proposition ‘Mary will pay back the money by Tuesday’, and for higher-order explicatures the propositions ‘Mary is promising to pay back the money by Tuesday’ and ‘Mary believes that she will pay back the money by Tuesday’.<sup>130</sup>

## 2.4.2 Appraisal

On the face of it, there’s much for us to like about Relevance Theory. When we were attempting to mine the theory of GCIs for a suitable candidate for an expanded notion of ‘what is said’, we had to do violence to the intentions of the theorist: GCIs are *implicatures* on that account—content that is, by definition, separate from what is said (in the strict sense).<sup>131</sup> If we consider

---

<sup>129</sup> *Ibid.*, p. 182

<sup>130</sup> Sperber and Wilson 2004, p. 623

<sup>131</sup> Levinson’s view on ‘what is said’ is quite subtle, actually. He thinks that there is pragmatic intrusion into what someone says that goes beyond the disambiguation and indexical fixing that Grice and others allow. In some constructions, GCIs become part of the truth-conditional content, and part of what someone says. For example, if someone utters “She either got married and had a child, or had a child and got married; but I don’t know which,”

adding explicatures to an expanded notion of ‘what is said’, we do much less violence to the intentions of the relevance theorists. They would be reluctant to say something along the lines of “what is said includes explicatures,” but only because ‘what is said’ is so tied up with the Gricean picture that they prefer just to talk of explicatures. Under the auspices of RT, though, we’re free to say of some of our liars above that the explicit contents of their utterances were falsehoods like ‘I ate at Papa Giorgio’s and, as a result, soon thereafter I became violently ill’, ‘The offer will either contain a provision that you pay for the roof repairs, or a provision that you pay for the foundation repairs, but not both’, and ‘We have one, and only one, child’. RT provides no redoubt for the sneaky liar, trying to defend himself by adverting to some mythical, minimal semantic content; for relevance theorists, minimal semantic content is just LF, which is hardly ever propositional (it is underdeterminate, typically, and requires *development* or *enrichment* to achieve propositionality, truth-evaluability).

What relevance theorists don’t provide is an acknowledgment of the kinds of in-between cases that many of our examples above highlighted—cases in which there are conflicting, ambiguous or indeterminate intuitions about the truth-conditional content of utterances. As we’ve seen, there are cases in which contents that RT would classify as explicatures (the ‘not all’ understanding of ‘some’) do not clearly belong on one side or the other of the said/implicit (or explicated/implicated) divide. But we have, in RT, the resources to provide an explanation for this phenomenon. As mentioned, the picture of language processing that relevance theorists favor is one according to which multiple interpretations for a given utterance are entertained by a sub-personal inference-module, operating deductively with the Principle of Relevance as a guiding

---

then what they say must include the implicatural contribution (and then), lest it be the utterance of a tautology (followed by a confusing remark: “I don’t know which.”). But while he allows this sort of intrusion, he does not countenance a notion of ‘what is said’ according to which the GCI becomes part of the content in an utterance of plain old “She got married and had a child,” and of course it is just this sort of inclusion that we were testing.

maxim. The procedure one follows involves testing different hypotheses—candidates for the correct interpretation—in order of their “accessibility” (roughly, the cognitive-computational ease with which they’re accessed) until one finds a suitably relevant one. Now, in this process, there will be degrees of accessibility for the various hypotheses; less-accessible ones will never be considered *during* the interpretive process, but they could be considered *afterwards*—when, for example, questions about the literal truth or falsity of someone’s claims are being entertained. At this point, one may find, on reflection, that while, for example, the ‘not all’ reading of the particular affirmative utterance was the most relevant (and accessible), the ‘some and possibly all’ reading was relatively close, in terms of cost of cognitive computation, to the interpretation that was chosen. In such cases, when an alternative understanding is relatively close to the chosen interpretation accessibility-wise, we might experience the kind of conflicting intuitions that we experience in cases where we weren’t sure whether the speech act was a lie or a mere deception. There could be a spectrum here: from cases in which the strict meaning proffered by the deceiver to defend against an accusation of lying is *much* less accessible (if at all) than the falsity actually communicated, to cases in which it’s quite a close call which of the two interpretations—the minimal truth or the more expansive falsity—is more accessible. This would explain the data above.

This is all quite promising, and I think it may be on the right track. Nevertheless, I’m very reluctant to embrace Relevance Theory as a framework for drawing the distinction between lying and misleading. The features mentioned above are nice, but RT brings along much extra theoretical baggage that I’m reluctant to take on board. This is not the place for a comprehensive review of the merits and demerits of Relevance Theory, so rather than set out detailed arguments, I’ll just briefly note some of my concerns. First of all, RT is committed strongly to a particular

picture of the way the mind works in general, and how it processes language in particular. As Robyn Carston puts it, “The relevance-theoretic approach... aims at a causal mechanistic account... in terms of interacting sub-personal systems,” among which are a “decoding system (the language module)”, a “deductive device” which operates inferentially on interpretive hypotheses using “the logical elimination rules [as] its proprietary database”, and a series of representations for background knowledge conceived as “cross-referenced encyclopaedic entries attached to particular conceptual addresses....”<sup>132</sup> I have personal reservations about both modular and computational theories of mind, so I’m reluctant to hitch my wagon to a theory that so fully embraces versions of each approach. But even if we set those concerns aside, and grant that there is in some sense a division of mental labor in linguistic processing, I’m not sure about the particular job-descriptions that relevance theorists provide. Utterance interpretation that’s as sensitive to contextual factors as the relevance theorists’ model suggests seems a paradigmatic example of an abductive, all-things-considered inference to the best explanation (of what the speaker is trying to convey). Sperber and Wilson’s efforts to avoid the frame problem by simultaneously characterizing this inferential process as non-demonstrative *and* deductive strikes me as a not-entirely plausible attempt to have their cake and eat it, too.<sup>133</sup>

Second, key relevance-theoretic notions seem a bit too vaguely defined to be of much predictive or explanatory use. I was able, above, to give a sketchy account of why some cases might present us with conflicting intuitions about what was said, in terms of the relative accessibility of different interpretive hypotheses. That’s a well and good. The problem is, I can’t see how one might provide anything more robust, an explanation less-sketchy. Relevance is supposed to be a function of positive cognitive effects and processing effort, but there is no way

---

<sup>132</sup> Carston 2002, p. 7

<sup>133</sup> See Sperber and Wilson 1986, Ch. 2.

to quantify these notions, no specification of how to measure them.<sup>134</sup> But without such specifications, I don't see any way to fill in the details of an account of accessibility differences that would be at all satisfactory—that would predict, say, that certain locutions' minimal readings will be more salient than others, and that our intuitions about deceptive utterances featuring them are likely to tend more toward the 'misleading' end of the spectrum.

Also, the key notion of *explicature* seems not to have clearly defined boundaries. Explicatures are supposed to be developments of logical form, but it's hard to see how this criterion creates a clear distinction between explicatures and implicatures. Take Grice's first example of a particularized conversational implicature, the utterance of "There is a garage round the corner" in response to a stranded motorist's "I am out of petrol." As we saw above when we were discussing Bach's quite similar notion of implicature, one can conceive of the communicated content as an enrichment (Bach would say *expansion*, relevance theorists would say *development*) of the original utterance, as in "There is a {n open} garage {with petrol for sale} round the corner.' As Recanati notes, Carston explicitly declares that this is an instance of implicature, but a plausible case can be made that according to RT, it should be considered an explicature.<sup>135</sup> Understanding how to circumscribe the domain of explicature is further complicated by the theory's inclusion of "higher-order" explicatures, as noted above. It is intuitively unappealing to include propositions like 'Mary believes she will pay back the money' as part of the explicit content of an utterance in which she promises to do so. Rather, it seems that such a proposition should count as *entailed* by the utterance.

---

<sup>134</sup> Sperber and Wilson deny that any absolute, quantitative characterizations of these notions is necessary—or even desirable. Rather, "intuitive methods of assessment" of *relative* degrees of cognitive effect, accessibility, and so on are supposed to do the job (2004, p. 610).

<sup>135</sup> Recanati 2004(b), p. 46

Finally, Relevance Theory's scientific ambitions—its status as an empirically verifiable (or falsifiable) theory—seem hard to justify. Sperber and Wilson point to a study that they say provides support for RT. It involves asking people at train stations for the time in various circumstances, some of which prompt the subjects to give precise answers (11:58, e.g.), while others prompt subjects to speak loosely (“It’s noon”). The fact that experimenters were able to elicit these different kinds of responses by altering scenarios shows that people respond to “subtle clues as to what might make [an utterance] relevant for the questioner,” and so the predictions of Relevance Theory are confirmed.<sup>136</sup> My question is this: did we need a full-fledged “cognitive psychological theory”<sup>137</sup> to tell us that people adjust the precision of their utterances according to what they perceive to be the hearer’s needs and expectations? How does this result confirm RT in particular, as opposed to any other account that would make the same utterly banal prediction? This citation by Sperber and Wilson comes in an article published in 2004, nearly thirty years after their initial formulation of RT. If pointing to this study is the best they can do to defend their theory’s claims to scientific rigor, after all this time, that strikes me as a serious problem.

## 2.5 LESSONS FOR THE DEBATE OVER ‘WHAT IS SAID’

Time to take stock. We’ve been searching through the philosophical and linguistic literature for a suitable conception of ‘what is said’ to underwrite our commonsense distinction between lying and misleading. We haven’t found one. But our search has been instructive: in the course of

---

<sup>136</sup> Sperber and Wilson 2004, p. 627

<sup>137</sup> *Ibid.*, p. 625

testing particular proposals, we had a chance to consider a wide variety of cases in which our focal distinction was at issue and to examine our intuitions about it. They turn out to be a bit of a mess. That is, our intuitions do not always give a clear answer to the question: Is this a lie, or an act of (mere) misleading? There are two possible explanations for this fact: either (a) our intuitions are not suitably sensitized by exposure to the proper theoretical basis for the distinction, and if they were, we would reform them in such a way that a clear boundary could be drawn between lying and misleading; or (b) the commonsense distinction cannot be made sharply, is inherently ambiguous, with a range of cases falling on a spectrum between the two extremes. While I'm still open to suggestions on option (a), I want to embrace (b) for now and see what consequences follow.

In particular, I want to explore the consequences for the longstanding debate over the proper understanding of the controversial notion 'what is said'. As we've briefly discussed already, there are, in this debate, basically two opposing camps: those who want to construe the notion minimally, in the style of Grice, and those who regard the minimal conception as untenable, preferring instead a notion that more clearly lines up with normal speakers' intuitions about truth-conditions. The most prominent exponent of the former view is Bach, though there are others; Recanati is perhaps the standard-bearer for the latter view. In the remainder of this chapter, I want to examine this debate in more detail and extract what insights I can from our explorations so far.

### **2.5.1 Refining Grice's minimal notion**

Grice's original characterization of 'what is said' was a minimal one, constrained by syntax and allowing only minimal contribution to content from context (fixing indexicals and resolving



ambiguity). The original formulation gives rise to some intractable problems, though, so in order to maintain a minimal conception of ‘what is said’, theorists in the Neo-Gricean camp have made certain modifications. Bach, as usual, is the principal here.<sup>138</sup>

Grice wanted his distinction between saying and implicating to be exhaustive of speaker-meaning: part of what we mean in an act of communication is what we explicitly *say*, and the rest is what we implicate. A consequence of this is that, for Grice, saying entails meaning; whenever we say something, we mean it. But this is obviously not true, if we understand *saying* in a particular (intuitive) way: when we’re being ironic or metaphorical, we say things that we don’t mean. When I say, sarcastically, “Well that’s just *great*,” after getting a flat tire, I mean the opposite of what I say; and when I say, “You’re a delicate flower” to my wife, I’m not making a botanical observation. Grice was aware of this problem, and so for such cases he distinguished between genuine saying and “making as if to say.”<sup>139</sup> But, as Bach puts it, “[i]t is more natural to describe these as cases of saying one thing and meaning something else instead.”<sup>140</sup>

For Bach, the key to refining Grice’s notion of saying is to be mindful of the difference between locutionary and illocutionary acts. The distinction, going back to Austin 1962, marks the crucial difference between merely saying something and what one *does* by saying it—promising, warning, informing, asserting, etc. A locutionary act produces a content; an illocutionary act assigns a force to that content. In order to get the content of the locutionary act, all one needs is the conventional meanings of the terms, plus disambiguation and resolution of indexical reference, if necessary. These are the minimal semantic ingredients that Grice

---

<sup>138</sup> See especially Bach 1994, §5 and Bach 2001, §1.

<sup>139</sup> Grice 1989, p. 30

<sup>140</sup> Bach 2001, p. 17

explicitly allows to determine what is said, and so it's the locutionary sense of 'saying' that he needs. His mistake, according to Bach, was to use 'say' in an illocutionary sense: "[T]he verb *say*, as Grice uses it, does not mark a level distinct from that marked by such illocutionary verbs as *state*, *tell*, *ask*, etc., but rather functions as a generic illocutionary verb."<sup>141</sup> If we restrict 'saying' to its locutionary sense, we can easily handle cases of irony and metaphor: one says one thing (locutionary act) and means (asserts, states, or whatever—illocutionary act) something else. This more restrictive sense of 'saying' also allows us easily to account for cases in which we say things unintentionally (through misuse of terms or slips of the tongue) and cases in which we say things but don't mean anything at all (as in recitation, e.g.).<sup>142</sup>

An important feature of Bach's minimalist conception of what is said is that it needn't be fully propositional. In utterances like "Rose isn't tall enough," what is said falls short of being a truth-evaluable proposition; there is nevertheless no problem in referring to this incomplete entity (a "propositional radical," for Bach) as the content of what is said, since there is no problem indirectly quoting it: The man said Rose isn't tall enough. It passes the "IQ Test." This potential lack of propositionality, of truth-evaluability, makes Bach's minimalist notion unsuitable as a candidate for 'what is said' in the sense that underwrites the distinction between lying and misleading. This we've already seen. But now that we've reached a point in our investigation of despair about ever finding a theoretical account that can capture all of our intuitions about that distinction, we're free to ask a different question. While Bach's minimal 'what is said' cannot *in general* distinguish lying from misleading, might it nevertheless play a role in *specific instances* in which the distinction is at issue? For example, when it *is* fully propositional, might it in those cases serve as the deceiver's fallback—what he said, strictly

---

<sup>141</sup> Bach 1994, p. 13

<sup>142</sup> Bach 2001, p. 18

speaking? Again, the answer is no—or at least, not always. As we’ve seen time and again, the purely semantic contents of utterances like “I’ve had lunch” and “You’re not going to die” are no safe redoubt for the deceiver, no defense against having lied.

What this suggests is that the sense of ‘what is said’ that we’ve been in search of is not a bare-bones, merely locutionary one, but rather an illocutionary one. Lying and misleading are, after all, illocutionary acts.<sup>143</sup> They are attempts to engender false beliefs by *saying* false (or true) things, in the sense of *stating*, *telling*, or *asserting* them. Our ordinary use of the word ‘say’ is ambiguous between the illocutionary and locutionary senses, but it’s the former we’ve been seeking. This suggests that Bach’s locutionary sense of ‘saying’ plays no role in drawing the distinction we’ve focused on. We shall see.

### 2.5.2 Truth-Conditional Pragmatics

The alternative to Bach’s austere, locutionary ‘what is said’, the content of which is entirely semantically determined, is provided by Recanati’s so-called truth-conditional pragmatics. His project can be seen as a kind of resuscitation of Grice’s original dichotomy between what is said and what is implicated, achieved by dropping the requirement that ‘what is said’ be syntactically constrained. In this picture, what is said by a speaker in making an utterance is determined by a variety of pragmatic processes. This conception of ‘what is said’ is propositional, truth-

---

<sup>143</sup> Again, a caveat: *misleading* is perlocutionary, since it indicates success in causing an effect (false belief) in the hearer. *Attempting to mislead* is the corresponding illocutionary act, but to save syllables I’ve just been using ‘misleading’ to refer to the illocutionary act.

evaluable, and illocutionary—a fine candidate, on the face of it, for underwriting our distinction between lying and misleading.

Recanati identifies ‘what is said’ with the intuitive truth-conditional content of utterances. As we’ve noted repeatedly, this often goes beyond what semantics gives us. He identifies various processes (“primary pragmatic processes”) that enrich bare semantic content to give us truth-evaluable propositions. Some are mandatory, “bottom-up” processes, like *saturation*, which takes us, for example, from an utterance of the sentence ‘She is smaller than John’s sister’ to a proposition specifying the referent of ‘she’ and the relevant relation between John and the sister mentioned.<sup>144</sup> This process is mandatory in the sense that we need to saturate in order to get a truth-evaluable proposition. It’s this process that’s involved in completing familiar examples like “Rose isn’t tall enough.” Optional pragmatic processes aren’t necessary to get a proposition, but their output is the intuitive truth-conditions of the utterance in question. So, *free enrichment* takes us from an utterance of ‘Mary took out her key and opened the door’ to the proposition that she did so with the key mentioned; it’s responsible for the specification of ‘rabbit’ in ‘He eats rabbit’ to *rabbit meat*; it gives us a particular contextually salient restricted class of books in ‘All the books are on the table’ (along with a specific table).<sup>145</sup>

This is all very good so far. The propositional results of pragmatic processes like Recanati’s correspond to what we’ve previously identified as ‘what is said’ by liars who unsuccessfully try to defend themselves by adverting to minimal interpretations of their utterances. Free enrichment gives us ‘I ate at Papa Giorgio’s last week and {shortly thereafter, as a result} got violently sick’ and ‘Tina and I are married {to each other}’. The problem is that relying on intuitive truth-conditions of utterances, ignoring alternative minimal interpretations,

---

<sup>144</sup> Recanati 2004(b), p. 23

<sup>145</sup> *Ibid.*, pp. 23 – 25.

gives us the misleading appearance of certainty about what's a lie and what isn't. Recanati's pragmatic processes get us from 'some' to 'not all', 'warm' to 'not hot', 'permitted' to 'not obligated' in the cases we discussed. But as we saw, it's not clear in those cases whether or not we should include the extra content as part of what was said; we had conflicting intuitions about whether the utterances were false, whether or not to call them lies.

It seems, then, that while the pragmatically enriched, intuitive truth-conditions of utterances play an important role in our linguistic lives, there remains some pull in the direction of minimal interpretations. And this was exactly the state of affairs that our focus on the lying/misleading boundary, and the associated test of intuitions, was meant to examine. A complete reliance of Recanati's intuitive truth-conditions gives the false impression that pragmatic propositions are all that matters; but a complete reliance on Bach's IQ test gives the equally misleading impression that minimally semantic contents play a more important role than they do. The truth is somewhere in between. Our deliberations will have been successful if they allow us to say something about the respective roles in our linguistic lives of both the minimal, merely locutionary 'what is said', and its enriched, illocutionary counterpart.

### **2.5.3 The role of a minimal 'what is said'**

Recanati objects to the very idea of a minimal, locutionary notion of 'what is said', on the grounds of its lack of psychological reality—its lack of “availability”—for interpreters of utterances. According to his “Availability Principle,” ‘what is said’ must be analyzed in terms of normal language-users intuitions about what is said by utterances, and those are their intuitions

about truth-values.<sup>146</sup> These are not determined by minimal semantic contents, and so those are not part of what is said. Moreover, minimal propositions are dispensable in the actual process of utterance interpretation.<sup>147</sup> When we're processing an utterance of 'Al has three kids' in a normal context, for example, we never consider the supposedly more basic *at least three* interpretation; we skip straight to *exactly three*. This is true of all sorts of default meanings. But if they're not part of what is said, and they're not part of the cognitive processing of utterances, then one wonders what the point is of talking about minimal propositions at all.

Bach argues that this reasoning rests on a confusion of psychological facts and semantic facts: "The process of utterance comprehension is obviously a very interesting topic for psychology, but it's hard to see why facts about *hearers'* cognitive processes should be relevant to what a *speaker* says."<sup>148</sup> That the minimal 'what is said' (often) plays no role in the cognitive process of interpretation doesn't show that it's a mere fantasy. He goes on: "All this shows is that hearers can infer what a speaker is conveying without first identifying what the speaker is saying. The semantic notion of what is said pertains to the character of the information available to the hearer in the process of identifying what the speaker is conveying, not to what goes on in this process...."<sup>149</sup> He pointedly uses the term 'available' here, which is Recanati's key notion. The sense in which minimal meanings are available to hearers is not that they *are* consciously present in the process of interpretation, but that they *could be made* so present, after the fact. The key datum here is cancellability. Even if we don't entertain the proposition that Al has at least three kids when interpreting an utterance of 'Al has three kids', we can be made aware of that way of understanding the utterance with the addition of a cancelling phrase: "Al has three kids—and

---

<sup>146</sup> *Ibid.*, p. 14

<sup>147</sup> *Ibid.*, §4.5

<sup>148</sup> Bach 2002, p. 10

<sup>149</sup> *Ibid.*

maybe even more than that.” The implicit qualification of the original to *exactly three* is thereby cancelled, and the minimal content laid bare. The same option is available whenever there’s an implicature, explicature, impliciture, whatever: cancellation reveals a more minimal proposition.

So here we have a clearly defined role for a minimal, locutionary conception of what is said: it explains the phenomenon of cancellability. In addition, as we’ve already seen, it allows us in cases of irony and metaphor to distinguish between what a person says and the (vastly different) thing that she means; and it gives us a sense of ‘saying’ in which one can say something unintentionally (misspeaking, e.g.), or say something without meaning anything at all (performing, e.g.). It is therefore going too far, as Recanati does, to dismiss the notion entirely. It has at least these roles to play in our linguistic lives. The really interesting question is whether and to what extent it plays the *additional* role of providing minimal content to fall back on in cases where the distinction between lying and misleading is at stake. And as we’ve seen, the answer to this question is not at all straightforward. In some cases, the pull of locutionary, literal meaning is sufficient to mark the relevant difference; in others, it clearly isn’t—and in still others, it seems to exert *varying degrees* of pull against pragmatically enriched illocutionary meaning, so that the verdict is not clear.

Instances of irony and metaphor, which make perhaps the strongest case for a merely locutionary ‘what is said’, also provide striking examples of the separation between this minimal sense of ‘saying’ and the more robust, illocutionary one. There is no pull in the minimal direction when we contemplate the use of irony and metaphor to deceive. Think back to my neighbor, the horrible cook. I and my wife finally relent and agree to attend a dinner party at his house. When dinner’s over, and everybody’s disingenuously congratulating the chef, I do the same, but indirectly: “Oh, that was just a *terrible* meal. I can’t remember ever eating something so

revolting.” I make this speech with an ironic tone of voice; he knows how sarcastic I can be, and takes me to mean the opposite of what I said—to be giving him a strong compliment. In fact, what I said was literally true: it was a terrible meal, the most revolting I’d ever had. Now, if the host found out about my true feelings, and called me out for lying, it would be laughable if I attempted to defend myself by pointing to the literal truth of what I said. Cases of metaphor are even clearer. Suppose I’ve succeeded in putting Papa Giorgio’s out of business, but he never discovered that I was the one who undermined him. He still thinks I’m his friend, and he lists me as a reference on his résumé. When his interviewers call me, I tell them, “That Papa: he’s no workhorse.” In fact, he’s one of the hardest working guys you’ll ever meet, but my deception costs him many job opportunities. Suppose it’s revealed to Papa what I said about him, and he’s mad at me for lying to prospective employers. A defense appealing to the literal truth of my claims—he is not, after all, really a horse—would fail. In both these cases, there’s a clear separation between the contents of my locutionary and illocutionary acts—what I’ve merely said and what I’ve *said*, in the sense of stating, asserting, telling. I’ve clearly committed myself to falsities, and the true content of my locutions provides no cover.

So when does minimal content exert a pull? We might have expected it not to in cases in which there’s such a clear separation between what we say and what we mean, like irony and metaphor. Obviously, there are cases on the opposite end of spectrum, in which we speak literally: the content of our illocutionary acts just are the contents of our locutionary acts. These are not interesting from our point of view, though; the act of (merely) misleading can only occur when locutionary and illocutionary contents differ, when one is true and the other false. Our field of inquiry is between these extremes, where what we assert or state is related, but not identical to, what we say.



One particularly murky phenomenon in this area is what Grice called conventional implicature. These extra contents (supposedly beyond what is said) arise because of the conventional meanings of terms used in utterances. So what we say when we utter ‘He is a philosopher, but he’s rich’ has the regular conjunctive truth-conditions; we implicate the contrast between being a philosopher and being rich. Bach denies that there’s any such thing as conventional implicatures; he includes the extra contents as part of what is said, in his sense.<sup>150</sup> He claims that a single sentence can express multiple propositions. Since these phenomena arise due to conventional meanings of terms in utterances, we might expect there to be a strong connection between locutionary and illocutionary contents, so that it would be hard in such cases to make a successful defense against having lied. But results are mixed, depending on the terms in question. Let’s consider a made-up scenario.

My brother-in-law has just married a girl from Poland (whom he met at the Polish bar/bowling alley where he works). Every week, the family convenes at our house for Sunday dinner, and this week he’s bringing his new bride. Now that she’s a citizen (by dint of the marriage), I’m determined to influence her political beliefs so she votes the way I want—for Democrats. She doesn’t know anything about the American political parties, but any political convictions she may acquire will be informed by her strongly Catholic upbringing (the Pope was Polish when she was growing up; that was a big deal). In particular, her Catholicism will strongly dispose her to favor candidates and parties with “pro-life” stances. This, of course, presents a roadblock for my plan to make a Democrat out of her. So I decide to resort to deception. At the Sunday dinner table, I say the following: “Rebecca Kleefisch is a Republican, but she’s against abortion.” Now, it’s true that Kleefisch is a Republican, and it’s true that she’s

---

<sup>150</sup> See Bach 1999.

against abortion, so on a strict truth-functional reading of the sentence, it's true.<sup>151</sup> It's deceptive in that the word 'but' conventionally conveys a contrast between the conjuncts, so that, in this case, there's something surprising about the second conjunct in light of the first. Did I lie, or merely mislead? An accusation of lying here seems to me a bit strained; it's hard to justify including the relevant false proposition—that Republicans are not, generally, “pro-life”—as part of what was said (strictly speaking). But suppose I go on to say the following: “Senator Bob Casey is a Democrat; he is, therefore, against abortion.” Again, true that Bob Casey's a Democrat, and true that he's “pro-life.” What's false is that there's a relationship between the two—that the former implies that latter. Did I *say* it did? Seems to me a stronger case can be made for the word ‘therefore’ than for ‘but’ that the supposed conventional implicature is part of what is said.

So I agree with Bach, and disagree with Grice: these phenomena should not be considered implicatures. Grice is just straightforwardly wrong about ‘therefore’, in my view, when he says:

If I say (smugly), *He is an Englishman; he is, therefore, brave*, I have certainly committed myself, by virtue of the meaning of my words, to its being the case that his being brave is a consequence of (follows from) his being an Englishman. But while I have said that he is an Englishman, and said that he is brave, I do not want to say that I have *said* (in the favored sense) that it follows from his being an Englishman that he is brave, though I have certainly indicated, and so implicated, that this is so. I do not want to say that my utterance of this sentence would be, *strictly speaking*, false should the consequence in question fail to hold.<sup>152</sup>

The reason he doesn't want to say that it's strictly speaking false, I would contend, is that he fails to consider an actual scenario in which that kind of judgment is called for—one in which the

---

<sup>151</sup> Ms. Kleefisch is Wisconsin's Lieutenant Governor.

<sup>152</sup> Grice 1989, pp. 25 - 26

difference between lying and misleading is at stake. This is an example of how our approach can help shed light.

Now, what should we say in general about putative conventional implicatures? They evidently can't be treated uniformly, as the difference between 'but' and 'therefore' in our example illustrates. If we follow Bach, and take secondary propositions like 'Republicans generally favor abortion rights' and 'Democrats are generally against abortion' to be part of what is said, we need an explanation for why one seems so much closer to being part of illocutionary content—to being asserted—than the other. In these cases I think the explanation is relatively straightforward: we're committed to the secondary propositions to different degrees in each case, and the degree to which we're committed is a function of the *meanings* of the words involved. It is part of the semantic contribution of 'therefore' that it commits one more strongly than the word 'but'. If we take assertion as the representative, generic illocutionary act, and following Robert Brandom 1983 and 1994, for example, take it essentially to involve commitment to contents; and if we take it for granted that one can be committed with varying degrees of strength, then we have a fine explanation for a phenomenon that was heretofore puzzling—viz., the fact that our intuitions about whether an act is a lie or (merely) misleading can vary, can come in degrees.

If only it were always so simple. What marks conventional implicatures as distinct from their conversational cousins, in Grice's scheme, is that the former are not cancellable. One cannot retract the implication, by saying, for example, "He is a philosopher; therefore he's rich—though I don't mean to imply that there's any connection between being a philosopher and being rich." Such an utterance would be infelicitous at best, contradictory at worst. This persistence of the extra content helps explain why it's so attractive just to include it as part of what is said—and

is part of the explanation of why this aspect of meaning resists abandonment in defenses against lying. Other contents that we've considered, though—like implicatures, explicatures, and GCIs—for inclusion in what is said, *are* cancellable, and so when they too resist abandonment, we can't appeal to the same sorts of explanations. In particular, the variability in our intuitions about different locutions can't be explained in terms of their meanings. We noted, for example, that among scalar implicatures, the minimal content of particular affirmatives (some... and possibly all) was easier to appeal to in making the case for the literal truth of an utterance than the minimal content of the corresponding deontic modal (permitted... and possibly obligated). It's hard to see how there's something about the *meanings* of 'some' and 'permitted' that explains the stronger tendency of the latter to rule out its superaltern.<sup>153</sup>

So what is the explanation? Why are some cancellable default interpretations so much more persistent than others? It's a matter of context, we might say. We can always imagine a context in which even the most intransigent extra contents are cast aside, making the minimal understanding the most natural. (In Taylor 2001, we get an imaginary scenario in which 'I've had breakfast' would most naturally be understood to mean that the speaker's life has not been entirely breakfast-less.) So too, one might say, can we imagine contexts in which the pull of the minimal interpretation varies from stronger to weaker. Maybe we can, but this does not explain the phenomenon we've identified. Our examples were all rigged to make the pragmatically enriched interpretations maximally salient; that's how you deceive: by making the false interpretation as natural as possible. And still we observed variation in the degree to which minimal propositions could be appealed to.

---

<sup>153</sup> The deontic modals, like the quantifiers, can be arranged in a square of opposition, with obligation in the A corner, and permission in the I corner. See Horn 1989, Levinson 2000.

We've identified a genuine puzzle, one that I've not seen recognized anywhere in the literature. The distinction between lying and misleading is hardly ever discussed, but even when it is, the complexity of the issues it raises is not acknowledged. For example, in Horn 2009, we get a defense of the minimal, neo-Gricean conception of 'what is said' on the basis of its utility in marking the distinction. The paper, a forty-year retrospective on the Gricean framework, is mostly a recapitulation of arguments found elsewhere and earlier; the material on lying and misleading is meant to provide a new line of defense. The abstract lays out the strategy, and is worth quoting in full:

40-plus years ago Paul Grice initiated modern pragmatics by defining a relation of conversational implicature within a general theory of cooperation and rationality. While critics have disputed the formulation and derivation of Gricean principles, the overall framework, with appropriate emendations, remains the most natural and explanatory approach to predicting constraints on lexical incorporation, the behavior of scalar predicates, pragmatic strengthening, and other linguistic phenomena. Despite recent arguments for an enriched conception of propositional content, a range of real and fictional exchanges bearing on the distinction between lying and misleading supports the neo-Gricean view of an austere conception of what is said.<sup>154</sup>

The "recent arguments for an enriched conception of propositional content" come from the usual suspects: Recanati, Relevance Theorists, etc. The kinds of cases Horn mentions early in the paper are familiar: "I haven't had breakfast {today}. ...John and Mary are married {to each other}. ...Robin ate the shrimp and {as a result} got food poisoning."<sup>155</sup> He wants to defend the position that the bracketed material is not part of what is said. But if this is his goal, and the material in the final section on lying and misleading is meant to be the culmination of his efforts, then his choice of examples in that section is strange. He doesn't consider any scenarios in which the familiar quoted examples are the disputed utterances—an odd choice, given how much weight his stated opponents put on them. And as we've seen, these are just the sorts of cases that,

---

<sup>154</sup> Horn 2009, p. 3

<sup>155</sup> *Ibid.*, p. 19

when situated within contexts in which the lying/misleading distinction is at issue, make it most difficult to maintain a minimalism about ‘what is said’. Instead of focusing on the examples we might expect, Horn gives us a fanciful “travelogue,” with an eclectic mix of samples of deceptive communication, drawn from history and fiction. But the cases tend to miss the mark, either because the examples of mere misleading they present can be accommodated within a non-Gricean framework, or because the chosen cases lack the features necessary to test the claims about our intuitions about lying and misleading.

We’ll start with the examples that are simply of the wrong kind. He gives us a snippet of dialogue from a television show, featuring an utterance that will be familiar:

*Kama:* ‘Is he gonna die?’

*Dr. Foreman:* ‘No, *no one’s gonna die.*’

*Kama:* ‘*In the whole world? Ever? That’s so great!*’<sup>156</sup>

Horn doesn’t explain what he thinks this bit of dialogue is supposed to show us, but it’s featured in a section meant to demonstrate that a minimal conception of ‘what is said’ can explain our intuitions about the difference between lying and misleading—and whatever this example gives us, it doesn’t support that claim. The most glaring problem is that the scenario doesn’t feature a deceptive act; as we’ve seen, if we want to gain insight into differences between species of deception, we need examples of deception. Further, the kinds of examples Horn needs are those in which the minimal proposition is intuitively *what is said*; we’ve looked at this kind of case before (the mother consoling the boy with a cut: “You’re not going to die.”), and it elicits the opposite intuition. At best, such an example demonstrates that we can be made aware of a more

---

<sup>156</sup> *Ibid.*, p. 27

minimal sense of ‘saying’, but in this case, at least, this minimal sense does not help us distinguish lying from misleading. It only explains why we take Kama’s remark to be a joke.

Horn gives us another example from a televised medical drama, in which the Meredith character apparently is having an affair with the surgeon mentioned:

*Benjamin: ‘Did you have sex with that brain surgeon?’*

*Benjamin’s sister: ‘Benjamin!’*

*Meredith: ‘It’s OK. Nope, I haven’t. [PAUSE.] Not today, anyway.’<sup>157</sup>*

This example is somewhat better than the first, in that it involves deception—at least until Meredith adds the qualification after the pause. But it still doesn’t suit Horn’s purposes. This case is similar to one that we began with: the lousy-cook neighbor and his appalling gumbo. In that scenario, an utterance of ‘I’ve had lunch’ clearly communicates—and, I argued, *says*—that the speaker has had lunch *that day*. As Taylor 2001 points out, though, sex is different from lunch: most of us don’t have it every day.<sup>158</sup> So in this scenario, the utterance of ‘Nope, I haven’t’ communicates that the speaker *never has*. And if we leave off the qualification after the pause, and suppose that Meredith is trying to conceal her affair, we have an act of deception. It’s a lie, though: everybody—neo-Gricean minimalists, Relevance Theorists, Truth-Conditional Pragmaticists—would agree that *what was said* was false. There is no separation between what is said and what is communicated, so the question of mere misleading doesn’t arise. We could alter the scenario somewhat, so that it would: suppose (contrary to (fictional) fact) that Meredith had never slept with the surgeon, and (again, contrary to (fictional) fact) she wanted the sick kid to believe, falsely, that she had. Now Meredith’s entire utterance, including the remark after the pause, would (arguably) be (merely) misleading: strictly speaking, it’s true that she hadn’t had

---

<sup>157</sup> *Ibid.*

<sup>158</sup> Actually, breakfast is the contrasting meal for Taylor.

sex with the surgeon that day; what's false is the implication that, though she hadn't that day, she had at some other time. This isn't the story we're given, though, so all this shows is that Horn needs different kinds of examples.

And there are other examples. He talks about the old Jesuitical practice of "mental restriction." This method of withholding the truth involves tacitly (mentally) qualifying one's utterances in such a way that they're true—*with* the qualification. So a person accused of killing can say, "I have not killed anyone," while mentally adding "since I got out of prison," and claim to be speaking truthfully, even if he did the killing in question, provided it occurred before his last prison term.<sup>159</sup> Such examples at least promise potentially to be relevant: they clear the bar of actually featuring deceptive acts; moreover, the distinction between lying and misleading is in play, since the Jesuits think this practice allows them to deceive without committing the (worse) sin of lying. Still, though, these examples fail to show that the minimal, merely locutionary sense of 'saying' can explain our intuitions between lying and misleading, again because they are the wrong kinds of cases. We have enough experience by now to know what the right kinds of cases look like: what Horn needs are scenarios in which there is misleading rather than lying, because what is said, strictly speaking, is true—where what is said is his favored neo-Gricean, minimal proposition. The Jesuits don't give us such cases. Here, the minimal proposition said is that the speaker has not (ever) killed anyone; the tacit qualification can't affect what is said. But this is false; the Jesuit lies. Use our test: imagine the murderer trying to defend himself against having lied by pointing to the mental qualification; such a defense is laughable.

The examples we've looked at so far fail because they lack the features necessary to highlight the differences between lying and misleading. Horn's other examples do have the

---

<sup>159</sup> *Ibid.*, p. 25



requisite features, but again they fail to meet his needs. Horn seems confident that the difference between lying and misleading is explicable in terms of “[t]he manipulation of these two dimensions of meaning—what is (austerely) said vs. what is effectively communicated.”<sup>160</sup> His confidence that a minimal ‘what is said’ can in general underwrite the distinction in question is, as we’ve shown, misplaced. Even if we set aside instances in which the minimal content is not even a proposition (but a mere “propositional radical,” as Bach puts it), so that it’s *obviously* unsuitable as a redoubt for the deceiver (what is said can’t be true if it’s not truth-evaluable), there still remain cases, many of which are enumerated in earlier sections of this chapter, in which the minimal proposition fails to provide an intuitive basis for a claim of having merely misled rather than lied. There are, of course, scenarios in which such a notion can successfully be appealed to, and we surveyed some of those; but it’s a fallacy to conclude from a few cases in which minimalism comports with our intuitions about lying and misleading that it can provide a *general* account of the distinction. But the most serious problem with Horn’s actual strategy is that his examples fail to isolate his preferred minimalist account as the best option. He chooses examples in which (1) we would intuitively classify an utterance as merely misleading; (2) any false content communicated would be a mere implicature or implication, not part of what is said; and (3) the austere, neo-Gricean proposition is true. But this is not enough to show that we “*need* to invoke minimal meanings to account for ordinary language intuitions of lying and misleading.”<sup>161</sup> This is because, in the scenarios Horn describes, advocates of a less austere conception of ‘what is said’ can also maintain that this content is true. But if both the minimalists and their opponents can agree that what is said is true, then such cases give us no reason to prefer the former view to the latter, contrary to Horn’s intentions.

---

<sup>160</sup> *Ibid.*, p. 24

<sup>161</sup> *Ibid.*, p. 30. Emphasis mine.

Here's an example: Horn describes the case of Saint Athanasius, pursued on the Nile by agents of the Roman emperor.<sup>162</sup> His pursuers approached him, but didn't recognize him, so they asked, "Is Athanasius close at hand?" Athanasius replied, "He is not far from here." Here we have, I think, a clear case of misleading—deceiving without lying. It would certainly pass my test: if someone accused Athanasius of lying, he could successfully defend himself by arguing that what he said was, strictly speaking, true. Horn and other minimalists would concur. But wouldn't just about anybody? I don't see why Recanati or Relevance Theorists, for example, would be forced to claim that what Athanasius said was false. They advocate for a more robust conception of 'what is said', but they don't deny that there are implicatures—aspects of what is communicated that go beyond even their beefed-up 'what is said'. I think everyone would agree that the falsity in this case—that Athanasius is not right there at that very spot—is merely implicated. It doesn't seem obviously to be a development of the logical form, as required by RT for explicatures; nor does it seem to be the output of Recanati's primary pragmatic processes.

Other examples from Horn share the same defect: they fail to distinguish minimalism from these other approaches as the preferred account of 'what is said'. He tells us the story of Queen Iseult and her lover, Tristan. The Queen is (rightly) accused of adultery, but she concocts a cunning stratagem: she arranges for her lover to appear on the Day of Judgment disguised as a poor pilgrim, then contrives to fall down in the mud near him; he picks her up and places her on dry land, at which point the Trial commences and she swears, "[N]o man born of woman has ever held me in his arms other than my lord King Mark, and that poor pilgrim...."<sup>163</sup> Again, clearly a misleading truth—but I don't think the minimalist would be alone in making that judgment. And Horn's final example is one we've seen, and to which we'll return, from the

---

<sup>162</sup> The case is borrowed from MacIntyre 1994.

<sup>163</sup> Horn 2009, p. 28

Supreme Court: *Bronston v. United States*. The defendant in a case before a lower court, a certain Samuel Bronston, was questioned during the original trial thus:

Q: Do you have any bank accounts in Swiss banks?

A: No, sir.

Q: Have you ever?

A: The company had an account there for about six months, in Zurich.

In fact, Bronston's company did have the account described; but Bronston also had a personal account (since closed). Bronston was convicted of perjury, but the Supremes overturned the ruling, since (in their judgment) the utterance in his last answer wasn't false, but only a misleading truth. Again, I think that minimalists, Relevance Theorists, and Truth-Conditional Pragmaticists could all agree with that assessment: that he didn't ever have his own account seems like a classic particularized implicatum.

Now, there are some theorists—we'll meet a few in the next chapter—who would maintain that in all these cases, what was said was, in fact, false. These people have a radically expansive conception of saying, which goes far beyond that of Recanati or RT—and Horn succeeds in distinguishing his position from theirs. But, again, he does not succeed in his stated aim of defending neo-Gricean minimalism from views closer to it on the spectrum. To do this, he'd have to consider the kinds of examples we've already surveyed at length—and when we consider those, minimalism doesn't come out looking like a very good candidate to underwrite the distinction between lying and misleading.

And by avoiding these sorts of cases—featuring GCIs, implicatures, explicatures, and so on—Horn doesn't confront the even more vexing phenomenon that our investigations have revealed: the fact of variability in the degree to which different locutions' minimal

interpretations can constitute such a basis. For this we still require an explanation—and I'm afraid I don't have one. What we require is an account of the relationship between the semantic content of utterances, on the one hand, and the pragmatically enriched content of the speech acts performed in making those utterances. But even to ask for such an account immediately raises vexed questions, principal among which is: How do we distinguish between Semantics and Pragmatics? It is to this and related questions that we now turn.

### 3.0 WHITHER SEMANTICS?

As it's traditionally conceived, Semantics is supposed to satisfy a number of conditions, perform a number of tasks: (1) specify the stable foundation of communication, namely word- and sentence-meanings, the latter being composed out of the former systematically—which is to say, in a manner determined, or at least constrained by the arrangement of the words, the syntax of the sentence; (2) yield contents—propositions, truth-conditions, or what-have-you—for declarative sentences that are truth-evaluable; (3) account for our intuitions about speech acts, both about the content of what we say when we utter sentences (consulting 1), and about the truth or falsity of those utterances (consulting 2); (4) serve as a starting point for the larger communicative task of interpreting what a speaker means, which often goes beyond what she merely says (including, e.g., implicatures), and is determined by Pragmatics, the inferential processes of which use the output of Semantics as their input.

There are problems with this basic picture, some easily solved, others less so. One need only modify the traditional conception of Semantics slightly, for example, to accommodate the problem posed by indexical expressions, which systematically change their reference in different contexts; Kaplan taught us how to do this. Other problems, however, such as those presented by the kinds of phenomena canvassed in the previous chapter, can put considerably more strain on the traditional picture and lead a semantic theorist to abandon entirely one or more of the four desiderata listed above. Some, impressed with the importance of the third condition, such as

Recanati and the Relevance Theorists, call into question the propriety of the fourth. Others, who maintain the primacy of the first condition (particularly the syntactic constraint), such as Bach, significantly alter the traditional picture, dramatically softening the third condition—so that it’s only the minimal, locutionary sense of ‘saying’ that is correlated systematically with semantic content—and abandoning the second requirement entirely, so that semantic content needn’t be fully propositional. Still others cannot countenance non-propositional semantic content, and so make other adjustments. And so on.

The purpose of this chapter is to look at various attempts to re-conceive Semantics in light of the vexing phenomena examined in the last chapter. This is a very large topic, so I will narrow the focus somewhat. I will examine in particular detail a proposal to preserve condition 2 (that semantic contents must be propositional, or truth-evaluable<sup>164</sup>) by Herman Cappelen and Ernie Lepore, which they call Semantic Minimalism. For them, the cost of preserving condition 2 is an almost complete abandonment of condition 3 (that Semantics explain our intuitions about speech acts). I want to examine this proposal vis-à-vis that of Bach and others, which jettisons condition 2 instead. This examination will comprise the first, much longer, section of the chapter. In the second section, I will briefly look at a few other proposals.

---

<sup>164</sup> I intend to set aside all metaphysical questions about the nature of semantic content, and so use ‘proposition’ as a place-holder term for whatever-semantic-content-may-be, even though it’s often used in such a way as to carry some metaphysical import. The main point is: semantic contents (propositions) are truth-evaluable, whatever their *true natures*.

### 3.1 CAPPELEN AND LEPORE'S SEMANTIC MINIMALISM

In the last chapter, we used the designation ‘Minimalism’ to characterize the views of Grice and his more contemporary intellectual heirs (Bach, Horn, et al.). The topic of that chapter was the debate about the concept ‘what is said’, and so minimalists, in that context, were minimalists about that notion—i.e., theorists who maintain that contextual factors play only a minimal role in determining what we say (namely, disambiguating and resolving indexicals). There is an alternative use of ‘Minimalism’ that we must be careful to distinguish from the one we used in Chapter 2. This is ‘Semantic Minimalism’. It is a minimalist view, not about ‘what is said’, but about Semantics: it is, roughly, the view that context plays only a minimal role in the determination of semantic content. Now, since it is often maintained that what is said by an utterance *just is* its semantic content, these two versions of Minimalism are often conflated. However, it’s important to see that they can be distinguished. It is possible, for example, to maintain that context plays only a minimal role in determining the semantic content of an utterance, while at the same time acknowledging that contextual (pragmatic) factors play a major role in determining what is said by the utterance. This is the position of the Semantic Minimalists we will now discuss.

The neo-Griceans we’ve seen are happy to refer to themselves as semantic minimalists, too.<sup>165</sup> But their minimalism is too austere for many, in that it accepts the possibility that the semantic content of an utterance may not amount to a full proposition.<sup>166</sup> In their 2005, Cappelen

---

<sup>165</sup> See Bach 2006.

<sup>166</sup> Recall that for Bach, the semantic content of, e.g., ‘Rose isn’t tall enough’ (or, what’s the same thing, the content semantically expressed by an utterance of the sentence) is a mere “propositional radical,” not a truth-evaluable proposition.

and Lepore (hereafter, C&L) implicitly<sup>167</sup> reject this version of minimalism, and maintain that every grammatically complete, indexical-free sentence—even those, like ‘Rose isn’t tall enough’, that have been supposed by others not to be truth-evaluable—has for its semantic content a “full-blooded proposition with truth conditions and a truth value.”<sup>168</sup> This semantic content is constant: it does not vary with context. Contextual variability, according to C&L, is an extremely limited phenomenon: only a select group of truly indexical expressions—members of what they refer to as the “Basic Set,” which includes personal pronouns, demonstratives, various adverbs like ‘here’ and ‘now’, and other like expressions—give rise to contextual variation of semantic content. In taking this stance, they contrast their view with those of so-called Contextualists, who see context-sensitivity as a much more pervasive phenomenon. According to C&L, a range of views falls under this heading, from various forms of Moderate Contextualism, which expand the class of context-sensitive expressions beyond the Basic Set (with different versions of Moderate Contextualism distinguishable according to *which* expressions they add to the Basic Set), to Radical Contextualism, adherents to which “all hold some version or other of the view that *every* single expression is context sensitive....”<sup>169</sup> C&L’s argumentative strategy is show that (a) Radical Contextualism is an incoherent view; (b) Moderate Contextualism is an unstable position, the arguments for which lead inevitably to Radical Contextualism; and (c) the only viable alternative is their Semantic Minimalism.

Before we look at any arguments, though, we should mention the way in which C&L deal with the data that make various forms of contextualism so tempting. They’re the kinds of cases

---

<sup>167</sup> “Implicitly” because they don’t acknowledge the neo-Griceans’ position as a version of Semantic Minimalism; they call Bach, e.g., a “contextualist.” In his 2006, Bach rejects this label, insisting that his view is a type of semantic minimalism.

<sup>168</sup> Cappelen and Lepore 2005, p. 181

<sup>169</sup> *Ibid.*, pp. 5 – 6.



we've been examining all along, in which a sentence, when uttered in different contexts, intuitively has different truth values. This can happen for sentences that are (apparently) semantically incomplete or indeterminate: 'Rose isn't tall enough' is true when it's a question of whether she can play for the Lakers, false when the issue is whether she can ride the Tilt-A-Whirl. Or it can happen for sentences whose determinacy is not in question: 'Rose is tall' is true when uttered in her second grade classroom, false in the Lakers' locker room. The variability of truth conditions here suggests variability in the propositions expressed, which suggests a variability of semantic content—for what is the semantic content of a sentence other than the proposition it expresses? C&L's response is to carefully distinguish between semantic content on the one hand, and speech act content on the other. The semantic content of sentences (except those with genuine indexicals from the Basic Set) does not change: a single proposition is *semantically* expressed by every utterance of the same sentence, regardless of context. What is said, claimed, asserted, etc. in an utterance of a sentence, though, *can* vary from context to context; this is not semantic content, but speech act content. They call their position on speech act content "Speech Act Pluralism" (SPAP for short):

No one thing is said (or asserted, or claimed, or...) by an 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.<sup>170</sup>

This is how C&L's minimalism differs from the minimalisms of the last chapter: they are minimalists about semantic content, but they are *maximalists* about 'what is said'.

In what follows, I will first examine the details of C&L's positive proposals—Semantic Minimalism and Speech Act Pluralism—to see whether they are plausible on their own terms

---

<sup>170</sup> *Ibid.*, p. 4

and consistent with the insights gleaned from the last chapter's meanderings. I will conclude that they are not. C&L's minimal propositions—the contents *semantically* expressed by every utterance of the same sentence, regardless of context—are problematic, and so there is no reason to prefer their form of minimalism over those that allow semantic content to fall short of being fully propositional. Also, many of the conclusions that C&L draw from SPAP turn out to be absurd on their face and at odds with the linguistic practices canvassed in Chapter 2. A major source of these problems, we will see, is C&L's failure carefully to distinguish locutionary and illocutionary acts; the importance of that distinction was one of the principal lessons of the last chapter. This failure, in turn, undermines the core of C&L's overall argument: the claim that Moderate Contextualism is unstable, that it inevitably collapses into Radical Contextualism.

### 3.1.1 Minimal propositions: metaphysical problems

Perhaps the most striking claim that C&L make is that *all* indexical-free sentences have for their semantic content a fully truth-evaluable proposition. The claim is striking because it is so easy to elicit contrary intuitions about certain types of sentences—such as ‘Rose isn’t tall enough’ and ‘Rose is tall’—the truth-values of which seem sensitive to variations in circumstance. C&L bite the bullet: the semantic content of ‘Rose isn’t tall enough’ is just the proposition that Rose isn’t tall enough, which is true just in case Rose isn’t tall enough—never mind for what; the semantic content of ‘Rose is tall’ is just the proposition that Rose is tall, which is true just in case Rose is tall—never mind varying standards for what counts as being tall. On the face of it, these claims are just nuts. What is it for Rose to be *just plain* tall enough (never mind for what)—or to be *just plain* tall? When are these (alleged) propositions true? They’re supposed to be “full-blooded

proposition[s] with truth conditions and a truth value.” OK, so what’s the truth-value of ‘Rose is tall’, independent of any more specific standards for what counts as being tall? I have no idea, and neither do you.

Cappelen and Lepore admit that they don’t have any idea, either. But that’s OK; it’s not their job to tell us, say, in which possible worlds Rose counts as *just plain* tall, or whether not she is tall in the actual world, because they’re semanticists, not metaphysicians. *Qua* semanticists, C&L observe that the semantic value of ‘Rose is tall’ is the proposition that Rose is tall, which is true just in case she’s tall. What plain old tallness amounts to, as a property, is a metaphysical question, and C&L are not metaphysicians.

Not that they don’t dip their toes a bit into the metaphysical waters. There are a number of different ways the metaphysician might approach the questions of what it is for Rose not to be tall enough (never mind for what), and what it is for her to be tall *simpliciter*. C&L survey a number of them. When they’re discussing (apparently) incomplete sentences, in fact, they seem to stake out a metaphysical position of their own. Considering sentences of the form ‘x is ready’, with instances like ‘A is ready to rob a bank’, ‘B is ready to eat dinner’, and ‘C is ready to take the exam’, they remark, “Thinking about A, B, and C... they all have a common relation they stand in to their respective projects. ...What they have in common is that they are all *ready*.”<sup>171</sup> Similarly for ‘enough’: “Consider a bunch of people who have all had enough. For example, one person who has had enough wine, one has had enough turkey, and one has had enough cocaine. All these people have something in common: *They have all had enough*.”<sup>172</sup> The metaphysical position that C&L seem to be taking is something like this: someone is ready just in case there exists something for which she is ready; someone has had enough just in case there exists

---

<sup>171</sup> *Ibid.*, p. 167

<sup>172</sup> *Ibid.*, fn. 6

something she's had enough of. In general, when we're faced with an apparently incomplete sentence, a sentence which, out of context, prompts a follow-up question like 'Ready for what?' or 'Enough of what?', it is enough that there is some thing answering to the 'what?' for us to affirm that the person is ready, has had enough, etc., *simpliciter*.

Here's a problem for this sort of schematic metaphysical answer: it threatens to make the extensions of predicates like 'x is ready' too large—maybe universal. C&L give examples of people being ready to take on various projects, but it's not just people who can be ready for things. My dog Olive is ready for a walk. My bedroom walls, having been primed, are ready for a coat of paint. My trashcan, now full, is ready to be emptied; my trashcan, not yet full, is ready to receive more trash. My pencil, now sharpened, is ready to be used to write things down; my pencil, now dull, is ready to be sharpened, or to be used to point at something, or to be brandished threateningly at an easily intimidated opponent, or to be tossed across the yard for Olive to fetch, or.... Pick any object: there's something it's ready for, in some sense of 'ready'. This H<sub>2</sub>O molecule is ready to be attracted to another one via hydrogen bonding. Goldbach's conjecture is ready to be proved (or refuted). I'm starting to think that everything is ready for something or other—in which case everything is ready, *simpliciter*.

But perhaps I'm doing some violence to the concept of readiness, stretching it beyond its normal well-defined bounds into the realm of metaphor. A professional metaphysician, with help from a concept monger of some stripe (philosopher of language?), maybe, could set me straight, tell me at which point above I crossed the line. Let's try a different predicate: 'x is big enough'. On what appears to be C&L's metaphysical view, what it takes to possess the property corresponding to this predicate is for there to exist at least one thing for which x is big enough. Rose possesses the property, since she's big enough to ride the Tilt-A-Whirl; I possess it, since

I'm big enough to drive a car; Olive possesses it, since she's big enough to reach my lap when she goes up on her hind legs; the Titanic possessed it because..., well, for lots of reasons; but so too does the gnat, who's big enough to be seen with the naked eye, unlike, say, a paramecium—but those are big enough to be seen with a microscope; and protons are big enough to be decomposed into smaller bits, quarks—which are big enough to be detected using sensitive enough equipment.... This is looking familiar. It seems that anything with any size at all possesses the property of being big enough. I'm not a professional metaphysician, but that doesn't sound like a well-behaved property.

What's worse, though, is that everything seems also to lack the property, since one can always find something for which any given thing is *not* big enough: Rose isn't big enough to drive a car; I'm not big enough to (successfully) sumo-wrestle; the Titanic wasn't big enough to significantly distort the geometry of space-time; etc. But maybe I'm misreading C&L's metaphysical view here. Perhaps it's not enough, to *lack* the property, for there to exist one thing for which one isn't big enough; perhaps, in order to lack the property, it must be that one isn't big enough for *anything at all*. That's fine, but things are still awkward. Rose possesses the property of being big enough (because of her Tilt-A-Whirl eligibility), but she also possesses the property of being too small, since there's at least one thing for which she's too small, e.g., driving. So, according to C&L, Rose is both big enough *and* too small. The proposition semantically expressed by 'Rose is big enough and Rose is too small' is true. That's a tough bullet to have to bite.

Relations can misbehave in similar ways. Consider 'x is better than y', and the sentence 'LeBron James is better than Knachel', which, C&L would tell us, semantically expresses the proposition that LeBron James is better than Knachel, and is true just in case James is better, full

stop. It can be used to perform innumerable speech acts, like asserting that James is better than Knachel at basketball, but the *semantic* content expressed is *just* that he is better, *simpliciter*. Again, we wonder what it is for one thing to be just plain better than another, which is a metaphysical question, and C&L's official position is that they're not metaphysicians, and so it's not their job to answer such questions, but we can extrapolate from their metaphysical musings about the predicates above to guess at their answer. They would insist that ordered pairs, about which we can say truthfully that the first is better than the second (at something, for something, etc.) must have something in common—and what else could that be than that there is something or other at which (for which, etc.) the first is better than the second. So, LeBron James is better than Knachel at basketball, and Roger Federer is better than Knachel at tennis, so the pairs <James, Knachel> and <Federer, Knachel> share in this *just plain better than* relation; the semantic contents of 'LeBron James is better than Knachel' and 'Roger Federer is better than Knachel' are true propositions. Fine. Those guys would surely kick my butt at the games which they are, respectively, among the greatest of all time to play. But I can make really good homemade falafel. Really good. Took me years to get my technique just right. I bet my falafel is way better than any falafel LeBron James or Roger Federer could make. If I'm right, then there exists at least one thing at which *I'm* better than both James and Federer—which, by parity of reasoning, means that the propositions semantically expressed by 'Knachel is better than LeBron James' and 'Knachel is better than Roger Federer' are both true. The result is: I'm just plain better than them AND they're just plain better than me. This odd result seems to me to deprive the supposed relation *being just plain better than* of metaphysical respectability.

But perhaps I'm getting carried away, either by misinterpreting sketchy remarks and thereby saddling C&L with metaphysical views at odds with those they actually hold, or, what's

possibly worse, imputing to them metaphysical opinions in the first place. As noted, they expressly deny any obligation to articulate metaphysical theses, since they are mere semanticists, and are constantly at pains to resist any “attempt[s] to force metaphysics on [them].”<sup>173</sup> And the hints at metaphysical speculation cited above with respect to ‘x has had enough’ and ‘x is ready’ are contradicted when, for example, in considering the predicate ‘x is red’ and its true saturations, they explicitly remark, “We’re most certainly *not* claiming that something is red just in case it is *red in some way* (or *in some respect* or *under some conditions*).”<sup>174</sup> So perhaps it was a mistake to impute to them the view, e.g., that something is big enough just in case there is something for which it is big enough. C&L aren’t metaphysicians, so they can neither be criticized for giving incorrect answers to metaphysical questions, *nor* for failing to give any answers at all, since that’s not their job. But, as John MacFarlane points out, “Semantic Minimalism is problematic not because it does not *provide* an answer to questions about the intensions of its minimal properties and propositions, but because it requires that there *be* answers to such questions.”<sup>175</sup> The reflections above, even if misdirected at C&L’s non-existent metaphysical positions, at least give a sense of the difficulty of solving the puzzles that their *semantic* position lays at the metaphysician’s doorstep. There are reasons to think that the sensible metaphysician would reply to C&L’s buck-passing by demurring—or even by declaring the problems insoluble. Perhaps there’s no such thing as being just plain red or tall or ready, of having had enough, *simpliciter*.

C&L consider this kind of response and reject it out of hand. Their reasons are a bit puzzling, though. That there are such minimal properties as these is presented as the result of a

---

<sup>173</sup> *Ibid.*, pp. 164 – 65

<sup>174</sup> *Ibid.*, p. 158

<sup>175</sup> MacFarlane 2007, p. 243

rather stark choice, whose only other alternative is Nihilism: “[T]hose who claim[, e.g.,] that there is no such thing as being red *simpliciter*, should, in order to be consistent, say that there’s a puzzle about the very idea of two or more things sharing properties.... To raise [such] objections... is to be committed, in effect to a form of *Metaphysical Nihilism*.”<sup>176</sup> Well, nobody wants to be a Nihilist, so.... But why should we think that there’s no middle-ground between Nihilism on the one hand, and acceptance of C&L’s position on minimal properties on the other? “Think about what metaphysicians do. For at least two millennia, metaphysicians have been asking *What-Do-They-Have-In-Common Questions* (CQ, for short). Suppose you’re curious about what it is to be G. Then you ask...: (CQ) What do all G things have in common?”<sup>177</sup> This kind of just sounds like the one-over-many argument. Either there is a thing (property) that they all have in common, or you’re a Nihilist. About redness, they say, “Some things are red on the inside, some are red on the outside, some are red when scrubbed, some are red in the dark, [etc.].... What do they all have in common? What is this property of redness that they all share?”<sup>178</sup> C&L’s view seems to be that there must be answers to “CQ” questions: given a variety of true sentences featuring the same predicate, there must be a property corresponding to that predicate which all the subjects of the sentence have in common. In other words, syntax is, in general, a reliable indicator of the features of reality. If we reject Nihilism, must we accept this alternative? No. The two positions are contraries, but not contradictories; they can’t both be true, but it’s possible to reject both. One can allow the reality of properties in the world, while at the same time denying that every predicate picks out a unique property. All the Moderate Contextualist, for example, is proposing is that *sometimes* predicates fail to correspond to any

---

<sup>176</sup> *Ibid.*, p. 163

<sup>177</sup> *Ibid.*, p. 159

<sup>178</sup> *Ibid.*, p. 160



property, or that instead they correspond to a relational property (tall relative to some comparison class). It may be harder for a Radical Contextualist, for whom the designation of every predicate is potentially up for grabs depending on the context, to stake out such a moderate position, but C&L face a serious burden to show that the moderate position collapses into the radical one.<sup>179</sup> It's certainly not obvious why someone couldn't take the position that sometimes predicates relate neatly to properties, and sometimes they don't, without inadvertently embracing Nihilism.

Such a view is outlined in Mark Wilson's 2006, *Wandering Significance* (and adumbrated in his 1982). In light of the considerations put forth there, C&L have made an extremely poor choice of predicate ('x is red') to serve as the prime example of predicate/property alignment and foundation for general claims. The steps in their argument are: (a) to show that there's no metaphysical problem for an easy case like *redness*, and (b) to show that if there's a problem for apparently harder cases (*readiness*, *tallness*), it must also be a problem for *redness*. Wilson shows, though, that many predicates, 'x is red' prominent among them, are quite unruly in their behavior, and fail to exhibit the kinds of tidy word-world relations of the simple predicate/property alignment type that C&L demand *a priori*. Rather, attributions of *redness* are judged according to a variety of standards, which shift according to the various practical purposes speakers have in making such attributions, and it's very difficult, if not impossible, to discover underneath all these shifting usages a single physical property to which all uses of 'x is red' correspond.

Wilson's conclusions are difficult to summarize succinctly, partly because one of the upshots of his investigations is that conceptual behavior is so variegated that a unifying account,

---

<sup>179</sup> Thanks to James Shaw for stressing this point about the difference between Moderate and Radical Contextualists on this issue.

e.g., of predicate/property relations is impossible. Rather, one must proceed piecemeal, examining the individual personalities (his metaphor) of words and the processes according to which they're shaped over the history of their deployment. The picture that emerges contradicts the classical account of word-world relations, in which our predications are neatly attached to their corresponding properties (through a process he dubs "classical gluing") without remainder; instead, we get a dizzying variety of modes of attachment, depending on the predicate under examination. One can get a sense of the complications here just by looking at a list of metaphors Wilson deploys to characterize the different relationships: atlases, patchworks, façades, platforms, bridges, lifts, parking garages; compare these to the classical picture of the pasting together of congruent rectilinear sheets. The complications documented arise over the course of the history of a predicate's use, and usually result from the necessity of accommodating that usage to novel and unforeseen circumstances presented by the ever-beguiling Mother Nature, in such a way as to maintain the usefulness of those words in accomplishing our goals. In the case of color predicates like 'x is red', depending on whether our projects involve design or manufacture, depending on whether our materials are cloth or metal, depending on whether our instruments are "reduction screens" or darkened tubes, and so on, we will end up evaluating claims of the form 'such-and-such is red' differently. Wilson summarizes the situation thus:

"[T]he predicate 'is red' spreads itself over a rather complicated atlas of naturally connected sheets and locally corresponds to quite different forms of evaluations, to the degree that its target objects are not even of the same type (behaviors under illumination lie below the color tag sheet; material objects under the color class patch). Considered as a whole—and since its usages continue into one another naturally, it should—being red can't qualify as a true attribute at all, but more nearly corresponds to an informational package of quasi-attribute type."<sup>180</sup>

---

<sup>180</sup> Wilson 2006, p. 467

The important point is that this behavior is a feature, not a bug: there is no problem with our concept of *redness*; it is not confused or ambiguous; we don't need to undertake a program of conceptual sanitization to clearly distinguish separate usages or to discover some underlying, unifying conception of *redness*; things are in order as they stand; they're just messy and complex, driven inevitably by the shifting practical needs of language users and the surprising revelations of Nature. To think otherwise is still to be in the grip of the classical picture, according to which understanding a predicate like 'x is red' involves grasping a simple, constant conceptual content, which determines in advance proper future usage.<sup>181</sup> But this illicit "projection" of syntactic simplicity will only lead us to pursue impossible projects, like searching for stable properties in a world populated only with "ghost attributes." Wilson catalogues such misguided adventures, with respect to color terminology, and even more strikingly, concerning the concept *hardness*.<sup>182</sup> C&L, I'm sure, would demand that there be something that all hard things have in common, some single physical property. But despite heroic efforts through the years by scientists and engineers at unifying the varying uses of 'hard' under a single definition<sup>183</sup>, no such property is anywhere to be found. Yet the predicate marches on, doing useful work and characterizing objective features of the world.

With these Wilsonian insights in hand, we can see how C&L's slippery-slope argument—if there's a problem with tricky predicates like 'x is ready', then there's a problem with every predicate, even easy ones like 'x is red'—fails. It is undermined in two ways. First, the supposed disaster at the bottom of the slope (we're forced to admit that even good old 'red' is

---

<sup>181</sup> Wilson's conclusion in this regard are echoed by Travis, who says of predicates about which people like C&L have unrealistic expectations, "The right understanding of [them] would foresee every eventuality in or to which the description might be applied." (Travis 1997, p. 99) Wilson, however, does not draw the radical conclusion that Travis does, that talk of properties, in general, is hopeless.

<sup>182</sup> See *Ibid.*, Ch. 6 (ix) – (x) for the gory details.

<sup>183</sup> Wilson recounts, in particular, the work of a certain physicist, Samuel Williams, who dreamed of doing for hardness what Kelvin did for temperature, viz., to establish an absolute scale.

problematic!) is no disaster at all. Second, it suggests that the slope is not at all slippery, or that it's not even a slope, since different predicates display different personalities and it's impossible to treat them all the same way. Some predicates behave very well: for instance, biological species-terms tend to line up well with natural attributes; Wilson's favorite example is the neat alignment between 'x is a dog' and *belonging to Canis familiaris*. Others, as we have seen, are quite unruly. The kinds of complex relationships between predicate and property that Wilson describes are particularly common among the bits of language that aim to describe the physical world; it is here that Nature is constantly surprising us. Thus, C&L's attempt to establish the simplicity of 'red' is all the more inappropriate in that it relies on an analogy between 'x is red' and 'x is a dance'. There are different ways of dancing, just as there are different ways of being red, C&L argue; but there's no problem saying that all the dances have something in common (*being a dance*), so there's no problem saying that all the red things have something in common (*being red*). Whether or not something is a dance, though, is entirely a matter of human convention; negotiation with other language users is all that's required to set the proper boundaries to that concept. In the case of 'red' and other physical-world predicates, however, there's no negotiating with Mother Nature; we must accommodate our usage to her unwavering demands. It should come as no surprise, then, that 'x is red' might display irregularities not present in the case of 'x is a dance'.

The lesson here is that C&L's minimal propositions lack the metaphysical inevitability that they try to establish for them. Their argument—"Sure, they look strange, but it's either minimal propositions or Metaphysical Nihilism"—fails; the metaphysicians to whom they pass the buck of sorting out truth-conditions for these things have plenty of good reasons for refusing to take up a hopeless task. One happy consequence of this, as far as I'm concerned, is that it

shifts a great deal of labor back into the domain of Semantics, where it belongs. It's up to the semanticist to keep track of the complex and shifting relationships that obtain between the words we use and the world that we use them to describe and get around in; if Wilson is right, this is an arduous task. Contrast this occupation with the sinecure that is Minimal Semantics: it seems that, for C&L, the semanticists' only job is to pound the table and issue T-sentences ("‘It's red’ is true just in case it's red, dammit!"), while resisting the wheedling of Tom Sawyer metaphysicians trying to get them to do their whitewashing for them.

### **3.1.2 Minimal propositions: psychological problems**

Besides being metaphysically dubious, minimal propositions are, *prima facie*, psychologically problematic: if we're unable (without the help of metaphysicians) to determine their truth-values, in what sense can we be said to grasp them? How can they play any role in the understanding required to sustain our communicative practices? They seem dispensable—in favor of actually grasped and entertained speech act contents—in an account of communication, so why insist that they have any psychological reality at all?

For instance, suppose two people are discussing falafels. One says to the other, "Knachel is better than LeBron James." In context, it's clear that she means to assert that I'm better than LeBron *at making falafel*; it is this content that both interlocutors entertain. Further, suppose it's a well-known fact (ESPN just had a big exposé about it, or something) that LeBron's falafels are among the worst in the world; he's legendarily bad at making falafel. With this in the background, the original utterance might give rise to an implicature, to the effect that Knachel's no great shakes, falafel-wise, himself (a case of "damning with faint praise"). To calculate the

implicature, the hearer only needs the assertive content ‘Knachel is better than James at making falafel’. So, in neither the process of original utterance understanding, nor in the process of implicature derivation, does the minimal semantic content ‘Knachel is *just plain* better than LeBron James’ play a role. So what is it good for?

C&L’s answer is that minimal semantic content is indispensable, not in the sorts of communicative interactions described above, in which the interlocutors share a context and are apparently fully attentive and knowledgeable about one another’s intentions, beliefs, background assumptions, and so forth, but in cases where such ideal conditions are lacking. When we don’t know the speaker’s/hearer’s intentions or beliefs, when we’re ignorant of the relevant contextual features that help determine more robust contents, whether because of inattention, mistaken assumptions, or the simple fact that we’re not in the same context, minimal semantic content earns its keep:

“The proposition semantically expressed is that content the speaker can expect an audience to grasp... even if they have incomplete or mistaken communication-relevant information. ...[It] is the content the audience can expect the speaker to grasp... even if she has such mistaken or incomplete information. ...[It] is that content which can be grasped and expressed by someone who isn’t even a participant in the context of utterance. ...[It] is that content which speakers and audiences know can be transmitted through indirect quotation... to those who find themselves in contexts radically different from the original context of utterance.”<sup>184</sup>

So, even if the hearer mistakenly believes that we’re talking about basketball, rather than falafel-cookery, even if the speaker thinks that LeBron James is a famous Lebanese chef, both can grasp the content semantically expressed, that Knachel is better than James. Even someone who wasn’t present at the original conversation, and had no idea what the participants were talking about or what they believed, could be told, “She said that Knachel was better than LeBron James,” and thereby gain some measure of access to the content of the communication. Not full access, of

---

<sup>184</sup> C&L 2005, pp. 184 – 85

course, but a “starting point,” as C&L put it.<sup>185</sup> The range of possible interpretations is narrowed (she wasn’t, for example, making a claim about Generalissimo Franco or Father Guido Sarducci); the process of decoding the utterance can begin from there. This is enough for minimal propositions to have a cognitive function, for them to be “psychologically real.”

This is a fine response, so far as it goes; it just doesn’t go as far as C&L want it to. I have no quarrel with a level of content that, for example, anchors the practice of indirect quotation. This is an ideal role for ‘what is said’ to play, where saying is considered as a merely locutionary act.<sup>186</sup> But while C&L have given us reasons for supposing that there must be some level of psychologically real and efficacious content to play these kinds of roles, and reasons for thinking that Semantics provides such content, they have failed to establish that this content must be propositional. The job of providing a starting point for decoding illocutionary content and more robust speaker-meaning can just as easily be performed by Bach’s propositional radicals, for example. Consider the following story about a speech act and an indirect report of it: Paulie reports to Shep that Alice said that Rose isn’t tall enough; Shep doesn’t know anything about the context of the utterance, or about Paulie’s or Alice’s beliefs in that context, but he knows who Rose and Alice are, and, as an English speaker, he knows the meanings of ‘is’, ‘not’, ‘tall’, and ‘enough’ and has sufficient grammatical mastery to appreciate the upshot of their mode of combination in this case. This narrows considerably the range of possible speech acts Alice could have performed (it’s unlikely, for example, that she promised to donate a kidney to Vladimir Putin), and gives clues to how Shep should continue an investigation to narrow things further (e.g., inquire about what Rose isn’t tall enough for). So far, so trivial. But notice: Shep can do all of this without ever comprehending, understanding, or entertaining the supposedly

---

<sup>185</sup> *Ibid.*, 185

<sup>186</sup> cf. §2.5.1 above

full-fledged, truth-valued proposition that Rose is *just plain* not tall enough. In fact, it's hard to tell a plausible story in which that does happen. C&L try to, though. Their story is about the proposition that something is red, but it applies to our case, *mutatis mutandis* (and if their story about redness is implausible, then a similar story about our case is even less plausible, since according to C&L, *redness* is easier to handle than *enough-ness*):

Suppose, for the sake of argument, that the proposition *that A is red* is trivially true (not something we have endorsed, simply something we speculated about [when discussing metaphysical objections to our view]). The audience knows that the speaker is talking about A and its redness, not, for example, about oysters, France, or Relevance Theory. There's a lot of stuff 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 they have of the speaker. It is trivial that A is red on some surface or other under some condition or other. The audience can assume that the speaker knew that this was trivial and was not interested in conveying such trivialities with his utterance and can, therefore, infer that there's work to be done in order to figure out exactly what the speaker was trying to communicate. 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.<sup>187</sup>

This story is going well, and mirroring my own above, until C&L attribute knowledge to the speaker and hearer that they cannot possibly possess. We start the passage by supposing, for the sake of argument, that one particular metaphysical proposal gives us the proper account of what it is to be just plain red. Even if we set aside our concerns about the plausibility of this, it is not plausible to suppose, as C&L go on to do, that the upshot of such an account would be known to speakers and hearers, so that they could, as the story goes, form communicative intentions based on, and draw conclusions from, the triviality of the proposition semantically expressed. C&L stipulate, in their discussion of the metaphysical objection to their views, that neither they, nor anybody else they know, can say what it is to be just plain red; that's for some (future, very clever) metaphysician to decide. But despite the current state of general metaphysical nescience,

---

<sup>187</sup> *Ibid.*, pp. 185 – 86



communication goes along just fine. And it must go along *without* the sorts of intentions and inferences based on non-existent metaphysical knowledge that C&L include in their story about ‘A is red’.

There’s a natural way charitably to alter their story, though, removing the imputation of implausible mental states to the conversational participants: C&L might insist that they entertain, comprehend, or grasp the fully propositional content, where such an achievement needn’t involve secret metaphysical knowledge, but merely knowledge of the truth-conditions of the proposition. These, as we’ve been repeatedly reminded, are just given by the relevant T-sentences: ‘A is red’ is true iff A is red. Surely there’s no problem attributing this knowledge to speakers and hearers. I don’t object to this knowledge-attribution; I just wonder about the significance and efficacy of such knowledge. It seems to come too cheap: if someone tells Rose (who speaks no German) that General Scheisskopf said that *der Schnee ist weiss*, she doesn’t know much, but she does know that ‘*Der Schnee ist weiss*’ is true iff *der Schnee ist weiss*. Surely, though, this is not enough for her to qualify as grasping the proposition. Rose lacks something that I have.<sup>188</sup> She can’t, for example, formulate a correct T-sentence in which her native language is the meta-language: ‘*Der Schnee ist weiss*’ is true iff snow is white. So perhaps it’s this more robust kind of knowledge that communicators have, for C&L, when they comprehend the T-sentence truth-conditions of propositions. But what does such knowledge consist in? That’s a big question, and I don’t know the answer (it depends, to some degree, on what kinds of entities we take propositions to be—metaphysical issues that I, along with C&L, would prefer to avoid), but one natural and fairly neutral thing to say is that knowledge of truth-conditions is knowledge of what the world would have to be like in order for the sentence to be

---

<sup>188</sup> I have forgotten German three times in my life, but at least remember this much.

true. This kind of knowledge, though, in cases like ‘Rose isn’t tall enough’ and ‘A is red’, is knowledge that no one has, due to the metaphysical uncertainties already canvassed. So the kind of state we need is somewhere between the trivial ability to remove quotation marks from sentences and the more robust understanding of how things would have to be to make the proposition true. Borrowing some theoretical terminology, we could state the situation thus: let the meaning of a proposition (its “intension”) be a function from possible worlds to truth values; I, with respect to a proposition whose meaning I grasp, know which possible worlds map to the True, while Rose, with respect to, for example, ‘*der Schnee ist weiss*’, knows nothing of its intension; as far as she knows, *all* of the possible worlds are potential candidates for being mapped to the True.<sup>189</sup> Cases like ‘Rose isn’t tall enough’ fall somewhere between these extremes: they don’t involve knowledge of which possible worlds a particular proposition’s intension maps to the True, but grasp of them does provide a significant paring-down of the range of possible worlds under consideration—various subsets of which will be the mapped-to-True domain-members of the intensions of the different determinate propositions that one might have intended to express by uttering the sentence in question. Thus we can even more plausibly characterize the content entertained as incomplete—a propositional radical (*a la* Bach), or a propositional blueprint (*a la* Neale), or a semantically non-specific sentence-type (*a la* Jay Atlas), or a propositional skeleton (*a la* Scott Soames). Such incomplete contents (incomplete in the sense that some contribution from Pragmatics is required to generate a truth-evaluable proposition), as Soames and Atlas both put it, *constrain* rather than *determine* the contents of our assertions and other speech acts.<sup>190</sup> It’s just this sort of constraining-rather-than-determining role

---

<sup>189</sup> I intend to endorse neither possible-worlds semantics, nor a Fregean reification of truth values; they’re just convenient and vivid ways of speaking, which, I hope, make my point clear.

<sup>190</sup> Soames 2008, 2009(b), 2010. Atlas 2005 and 2007.

that C&L seem to have in mind for their minimal propositions; but it seems just the sort of role for something *less* determinate to play. There are multiple extant accounts involving such indeterminate semantic content, and so, absent some convincing argument that semantic content *must* be fully propositional—and as we’ve seen from our discussion of metaphysical issues, C&L are not convincing on this score—highlighting Semantics’ indispensability as a communicative starting-point gives us no reason to prefer C&L’s propositional minimalism over other, non-propositional versions.

### 3.1.3 Speech Act Pluralism

One advantage that minimal propositions might be supposed to have over their non-propositional competitors (radicals, blueprints, skeletons, etc.) is that it’s at least possible for them to play a role in the communicative phenomenon under investigation in the present work—namely, the practice of distinguishing lies from mere acts of misleading. Minimal propositions, being fully propositional, are truth-evaluable—and as we’ve repeatedly noted, the non-liar needs truth-evaluability, as his defense involves pointing to the truth of what he said (strictly speaking). Propositional radicals and the like cannot provide this fallback content, since they can’t be true or false; but maybe C&L’s minimal propositions could.

Alas, they cannot. And not just for the obvious reason that in many cases their truth-values are a metaphysical mystery. C&L expressly deny that they can play such a role—indeed, they deny that any sense at all can be made of the notion of ‘what is strictly speaking said’: one of the consequences of their Speech Act Pluralism is that no one of the indefinitely many things said/claimed/asserted by a speaker of an utterance has such a privileged status. I think it’s a

mistake to draw this conclusion, not only because it conflicts with the common practice of language users, who do use such a notion to distinguish lying and misleading, but because the conclusion ignores the distinction between locutionary and illocutionary senses of saying—the importance of minding which was one of the principal lessons of Chapter 2. C&L’s failure to acknowledge this distinction leads to all sorts of errors.

As we’ve seen, while C&L are minimalists about Semantics, they are maximalists about the notion at the center of Chapter 2—‘what is said’. They describe themselves as deliberately naïve about speech-act content: their method for uncovering it is simply to listen to utterances, then consult their intuitions about what has been said/claimed/asserted.<sup>191</sup> Using this procedure, they are evidently impressed by the massively permissive practices normal speakers use in reporting the contents of the speech-acts of others. They consider the infamous “Smoking Gun” utterance(s) by Richard Nixon from the tapes of his conversations in the Oval Office—the discussion with Robert Haldeman in which Nixon orchestrates the cover-up of the Watergate burglary—and note that one could legitimately characterize what is said by Nixon in any number of ways (“Nixon told Haldeman to tell the CIA to tell the FBI not to pursue their investigation into the Watergate burglary. ...Nixon told Haldeman to tell Helms that Nixon wanted him to stop the Watergate investigation. ...Nixon told Haldeman to break the law. ...And so on and so on.”<sup>192</sup>) They go on to note that this kind of flexibility in reporting speech-acts is typical, and conclude that indefinitely many things are said/asserted/claimed by speakers with their utterances. Furthermore, they claim, no one of the indefinitely many contents one could legitimately report as what someone said can be singled out as the one thing that they said *literally* or *strictly speaking*. There’s not much argument for this claim; they simply ask us to

---

<sup>191</sup> C&L 2005, p. 191

<sup>192</sup> *Ibid.*, pp. 192 – 193

reconsider the Smoking Gun Tape and ask, “Did Nixon *really* say what we attributed to him? Did he, *strictly speaking*, say it? Did he *literally* say it? Let’s see: We know this particular utterance was the cornerstone of the impeachment case against Nixon. ...Imagine the absurdity of a defense of Nixon that he didn’t, strictly speaking, ask the CIA to block the FBI investigation.”<sup>193</sup> We can’t even single out the minimal proposition expressed by an utterance as having some special speech-act status; minimal propositions are but one among the many things said: “One of the many propositions asserted by an utterance is the semantic content of the utterance (the proposition semantically expressed).”<sup>194</sup>

This is familiar ground for us by now. We took note of the permissibility of various reports of a single utterance at the very beginning of the Introduction (discussing the sign that read ‘10 for \$10’ and Camp’s example concerning Jim and his attendance at Democratic Party meetings at the Bluebird Diner). So I agree with C&L that our linguistic practices typically sanction, for any given utterance, a wide variety of reports of the form ‘S said that P’, and that these practices should be taken at face value; I disagree with their sweeping conclusion, however, to the effect that this, let us say “flexibility” (not “context-sensitivity”) in our use of ‘saying’ implies that no sense at all can be made of the concept of what is (strictly speaking) said. After all, if we’re supposed to “take our nontheoretic beliefs and intuitions about what speakers say, assert, claim, etc. at face value,”<sup>195</sup> then we should treat the common practice of distinguishing lies from mere deceptions with the same deliberate naivety, and hence take seriously the notion of strict saying on which the distinction is based. This doesn’t mean that we should demand or even expect a rigorous theory of such a notion—as Chapter 2 showed, there is

---

<sup>193</sup> *Ibid.*, pp. 196 – 197

<sup>194</sup> *Ibid.*, p. 200

<sup>195</sup> *Ibid.*, p. 191

every reason to be pessimistic about the prospects for a comprehensive account—but neither can we dismiss it as unintelligible or uninteresting. True, ordinary language-users accept a variety of contents as appropriate characterizations of what a speaker has said; but it's also true that, depending on the circumstance, ordinary language-users make distinctions among the variety of things that can be said to have been said, according to the appropriateness of attributing acts with the various contents to the speaker. One such act is that of lying, and an examination into the circumstances under which it's appropriate to label an utterance as such can reveal the subtle contours of the notion of strict saying. C&L tacitly acknowledge that this would be the way to get at such a notion, when they consider a hypothetical defense of Nixon along the lines familiar to us from Chapter 2—that he didn't, strictly speaking, say that Haldeman should break the law, for example. But they're too hasty when they apparently conclude, from the clear fact that such a defense in that particular case would be considered absurd, that the concept of strict saying is entirely inscrutable. There are plenty of cases, as we've seen, in which a defense against an accusation of having lied *does* succeed, and these cases involve arguing that some false content is not part of what one has, strictly speaking, said.

When part of what one says, or one of the things that one says, is false then one has, arguably, at least to some degree, lied. C&L's failure to countenance distinctions among the various things we say with an utterance makes liars of us all. Oftentimes the content semantically expressed by our utterances (C&L's minimal propositions) is false, even if the meaning we intend to convey with that utterance—one of the other things we say, by C&L's lights, but one which is on a par with each of the other innumerable things we say, including the minimal proposition—is something completely and uncontrovertibly true. "Rose isn't tall," I might say as she's getting dominated by Shaquille O'Neal in a game of one-on-one. According to C&L, one

of the things I've said is the minimal proposition that Rose isn't tall, *simpliciter*. As we've seen, it's hard to say what this means, what its truth-conditions are (that's a job for the metaphysicians), but according to one proposal floated by C&L, it turns out false if Rose has any degree of height at all. So I'm a liar, at least in part. And we needn't adduce metaphysically dubious minimal propositions to make this point. Less-weird minimal propositions can also be false when there's no lying going on: consider again the mother telling her barely injured son that he's not going to die; for C&L, one of the things she says is that he's immortal, and so she's a liar. But this is absurd.

Now, we might avert this conclusion by appealing to one of the results of Chapter 2—namely, that the sense of 'saying' at play when we're talking about lies is illocutionary, rather than merely locutionary. Lies are what Searle 1969 would classify as *assertive* acts, those that commit the speaker to the truth of their content. There is no such commitment in the mere expression of content involved in a locutionary act. So if we construe the expression of minimal propositions as merely locutionary saying, we're off the hook for having lied. This perfectly natural and obvious move, however, is not open to C&L. When they talk about the expression of semantic content, they codify their position as the second thesis of Speech Act Pluralism, under the heading "*What's said and what's semantically expressed*," thus: "One of the many propositions *asserted* by an utterance is the semantic content of that utterance (the proposition semantically expressed)." <sup>196</sup> We assert minimal propositions! We're committed to their truth. And so we lie all the time.

The occurrence of 'asserted' rather than 'said' in the second thesis of Speech Act Pluralism is not just a slip-up, an accident. C&L seem to recognize no distinction between saying

---

<sup>196</sup> *Ibid.*, p. 200. Emphasis mine.

and asserting: they consistently treat various speech-act verbs as interchangeable. Their formulation of the first thesis of SPAP is typical: “No one thing is said (or asserted, or claimed, or...) by any utterance: rather, indefinitely many propositions are said, asserted, claimed, stated, etc.”<sup>197</sup> The point is not that all these different *kinds* of acts—saying on the one hand, claiming on the other—are performed with an utterance; they treat ‘say’, ‘claim’, ‘assert’ and so on as equivalent. One can open the book to a random page and have a fine chance of seeing this: on page 91, we have “According to [Radical Contextualism], the two utterances of (1) assert (say, claim) radically different propositions.” Often, ‘say’ is out front, with ‘claim’ and ‘assert’ in parentheses; here they switch it up, clearly indicating their insensitivity to any distinctions in this area. C&L’s Speech Act Pluralism is a massively permissive (to borrow a phrase from Camp) one, which seems to acknowledge no difference between various distinguishable speech acts.

### 3.1.4 SPAP undermines other arguments

But C&L should be careful to draw a distinction between mere locutionary saying and its more forceful illocutionary cousins; without it, some of their central arguments are undermined. Recall the overall argumentative strategy of *Insensitive Semantics*: they seek to repudiate contextualism in all its forms, from moderate to radical varieties, first by showing that moderate versions collapse into radicalism, and then by showing that Radical Contextualism (RC) is incoherent. Contextualists are theorists who posit semantic context-sensitivity where there really isn’t any (according to C&L). The only genuinely context-sensitive expressions are members of their Basic Set, which includes the usual paradigmatic indexicals: pronouns, ‘here’, ‘now’, etc. The

---

<sup>197</sup> *Ibid.*, p. 199



centerpiece of C&L's attack on RC comes in Chapter 7, where they teach us how to distinguish between genuine and spurious context-sensitivity by formulating tests which members of the Basic Set pass, and which the contextualists' putative context-sensitive expressions fail. The reason that C&L ought to carefully distinguish locutionary saying from, e.g., asserting, is that these tests give different results depending on which speech-act verb one uses.

Recall that one of the primary reasons for embracing a locutionary sense of 'saying' is to explain metaphorical utterances: intuitively, these are instances of the speaker saying one thing ("My wife is a delicate flower") and meaning quite another (she's lovely, etc.); we need a non-assertive, force-neutral sense of 'saying' to account for the lack of commitment on the speaker's part to the literal content of the act. Camp 2007 argues that C&L's tests, when applied to metaphorical utterances, only give the correct result—or, at any rate, the result that C&L (and most other theorists) want—when formulated with a locutionary sense of 'saying'; if one uses illocutionary verbs, the tests end up showing that metaphor is semantically context-sensitive. Consider C&L's first test for genuine context-sensitivity: they point out that expressions in the Basic Set tend to block "inter-contextual disquotational indirect reports," while other allegedly context-sensitive expressions do not.<sup>198</sup> These reports are of the form "So-and-so said X," where X is the very same sentence that so-and-so used in the original utterance. When we shift to a relevantly different context, these reports fail when X contains an expression in the Basic Set: suppose I utter "It's going to be rainy tomorrow" on October 24<sup>th</sup>; on October 25<sup>th</sup>, Rose cannot report my utterance by saying "Dad said it's going to be rainy tomorrow," since in the new context, the reference of 'tomorrow' has changed. 'Tomorrow' is thus context-sensitive. Expressions that aren't in the Basic Set fail the test: change the contexts however you like and

---

<sup>198</sup> *Ibid.*, pp. 88 – 99

disquotational indirect reports will still succeed. So, for instance, about ‘tall’, they point out that “[a]ny utterance of ‘A is tall’ can be reported by ‘She said that A is tall’ and any two such utterances can be reported by ‘They both said that A is tall.’”<sup>199</sup> Thus, ‘tall’ is not context-sensitive. But, as Camp notes, if you change the verb from ‘said’ to ‘claimed’, the results of the test change. She considers two utterances in which a woman, Jane, is said to be a “long-stemmed rose.”<sup>200</sup> In the context of the first, it’s clear that the metaphor is meant to convey the speaker’s opinion that Jane is ostentatious and prickly; in the context of the second, the speaker thinks she’s classy. If we insert ‘Jane is a long-stemmed rose’ into C&L formulation above about tallness (a move suggested by C&L’s actual practice of re-deploying that schema for ‘tall’, ‘enough’, ‘know’, and all the other words whose context-sensitivity they regard as spurious), we get the perfectly unobjectionable observation that any utterance of ‘Jane is a long-stemmed rose’ can be reported by ‘She said that Jane is a long-stemmed rose’ and any two such utterances can be reported by ‘They both said Jane is a long-stemmed rose.’ This is true. But Camp reports that her intuition is that, if we report the utterances from a third, unrelated context using ‘claimed’, the following three reports are false: (1) “Alex claimed that Jane is a long-stemmed rose”; (2) “Charlie claimed that Jane is a long-stemmed rose”; (3) “Alex and Charlie both claimed the same thing: that Jane is a long-stemmed rose.”<sup>201</sup> My intuitions about (1) and (2) are not as strong as Camp’s, though I think those examples should remind us to make another important point: metaphorical utterances highlight the absurdity of C&L’s formulation of the second thesis of Speech Act Pluralism: “One of the many propositions asserted by an utterance is the semantic content of that utterance (the proposition semantically expressed).” The semantic content of the

---

<sup>199</sup> *Ibid.*, p. 95

<sup>200</sup> Camp 2007, p. 197

<sup>201</sup> *Ibid.*

utterances about Jane is that she's (literally) a certain kind of flower; according to C&L, Alex and Charlie assert this, and so are committed to it. Balderdash. But back to (1) – (3): since (1) and (2) aren't obviously false, it appears open to C&L to say that in these examples of intercontextual disquotational indirect reports are not blocked, so there is no context-sensitivity, just as one would expect. But while (1) and (2) aren't obviously false, I think there's a strong case that (3) is. If (3) is true, then it follows that Alex and Charlie made the same claim. But they didn't: one is a flattering assessment of Jane, the other is not. If they made different claims, then by *modus tollens* (3) is false. And the falsity of (3) undermines C&L's test: for them, the reason that disquotational reports are not blocked for expressions outside the Basic Set is that Radical Contextualists are wrong to think that different utterances of the same sentence can “assert (say, claim) radically different propositions.”<sup>202</sup> Rather, a single content is said/claimed/asserted, and so we can observe, e.g., that any two utterances of ‘Jane is a long-stemmed rose’ can be reported by ‘They both said that Jane is a long-stemmed rose.’ Since what underwrites this collective reporting is the sameness of the contents in both utterances, and since C&L are also committed to the interchangeability of ‘said’ and ‘claimed’, it's equivalent to put it as Camp does in (3): They both claimed the same thing: that Jane is a long-stemmed rose. And that is wrong, since they expressed different opinions of Jane—claimed different things about her—with their utterances. These utterances behave exactly the same way as two utterances featuring members of the Basic Set would: if I utter, on October 24<sup>th</sup>, “It's going to be rainy tomorrow,” and Rose utters the same sentence on October 25<sup>th</sup>, we have made different claims; for C&L, this is the reason inter-contextual disquotational reports are blocked. But if that's the case, then they would have to report the same intuitions as Camp about the falseness of the claim-reports (1) and (2)

---

<sup>202</sup> *op. cit.*, p. 91

above. It's easy to see how one can extend this procedure to other metaphorical utterances, and the end result is that metaphor seems to block inter-contextual disquotational indirect reports in the same way as members of the Basic Set, making metaphor semantically context-sensitive. This is an unattractive result for most theorists, and is certainly not a result that C&L would be happy with.

Camp is concerned strictly with metaphor, but her point generalizes. It's a simple fact of linguistic life that we use identical utterances in different contexts to perform different illocutionary acts. When I say "Rose isn't tall enough" in the Lakers' locker room, I make one assertion; when I say the same thing in front of the Tilt-A-Whirl, I make a different one. If we concentrate, then, on illocutionary acts, we will find context-sensitivity everywhere. C&L want to avoid that outcome, and so they need to distinguish carefully between locutionary and illocutionary acts, and formulate their tests using only the former notion. They don't; instead, they explicitly conflate the two, and so their Speech Act Pluralism leads to absurd results.

The solution—to embrace a merely locutionary sense of 'saying', alter the tests accordingly, and moderate SPAP—is not, however available to C&L, because it undermines their arguments against Moderate Contextualism (MC) from earlier in the book.<sup>203</sup> Again, recall that a crucial step in their overall argumentative strategy is to show that adopting MC puts one on a slippery slope to RC. They try to establish the slipperiness of the slope by showing that the moderate contextualist's methods for establishing the context-sensitivity of *particular* expressions can be used to demonstrate the context-sensitivity of *any expression whatsoever*. One method involves the deployment of so-called Context-Shifting Arguments (CSAs), in which

---

<sup>203</sup> Camp thinks this move is available to them: "Once we are armed with the distinction between saying and asserting, we can fairly easily revise Cappelen and Lepore's tests.... In the end, I doubt that shifting from 'what is said/asserted/claimed etc.' to 'what is said' construed restrictively will offer much comfort to the contextualist. ...But it does require abandoning the most wanton form of SPAP." (*op. cit.*, p. 212)

we imagine the same sentence uttered in two different contexts, note that the utterances have different truth-values in the different contexts, and conclude that the culprit is context-sensitivity of one of the expressions therein. To show that this pattern of argumentation can be used to make any expression appear context-sensitive (and thus that MC collapses into RC), they run through several examples of CSAs featuring expressions that, they imagine, only a radical would deem to be context-sensitive.<sup>204</sup> So they borrow an example from well-known radical Travis: ‘Smith weighs 80 kg.’ They imagine this sentence being uttered in two contexts: in the first, Smith has been on a diet and exercise regimen, and the speaker truly characterizes Smith’s success with the utterance, even though, as a matter of fact, while Smith’s morning weigh-in came to exactly 80 kg, at the time of utterance he is wearing lots of heavy clothing and has just eaten an enormous lunch; in the second, all the facts are the same, except that the utterance occurs at the doors of an elevator whose maximum capacity is exactly 80 kg short of being reached, making the utterance false.

One can imagine resisting the conclusion that ‘weighs 80 kg’ is context-sensitive<sup>205</sup> thus: while it’s true that the utterances are used to make different *assertions* with different truth-values, what is said, *strictly speaking*, on each occasion is the same thing, and it’s false. The content of the locutionary act in both cases is the same, and is delivered by Semantics, and its truth-conditions are such that it’s false unless Jones weighs exactly 80 kg. Whether or not one agrees with the semantic diagnosis, this kind of response is an option. C&L consider the possibility of an objection appealing to what is strictly speaking said, but since they lack the notion of a locutionary act, they are unable to assess it fairly. Considering questions about what

---

<sup>204</sup> *op. cit.*, pp. 44 – 48

<sup>205</sup> I’m deliberately setting aside the various Wilsonian complications involved with the concept of ‘weight’. See Wilson 2006, Ch. 6 (viii).

someone strictly speaking said, they say, “As far as we can determine, the only way to take this question so that its answer is ‘no’ is as a question about *direct* quotation.”<sup>206</sup> That is, the only way to specify what someone said, strictly speaking, is to quote their exact words. But we have an alternative sense of strict saying as a locutionary act, and I can correctly report the content of this act without resorting to direct quotation. For example, if Alice says “Smith weighs 80 kg,” I can report her utterance (supposing that Smith has a certain job) thus: “Alice said that the Regional Manager of Sales tips the scales at 176.37 lbs.” Co-referring expressions, synonyms, and so forth can be part of the specification of a locutionary act’s content. What we have then, is “a special class of indirect reports of the form ‘A literally (strictly) speaking said that *p*’ [that] creates a connection between indirect speech and semantic theory relevantly similar to the connection [that C&L have] been denying.”<sup>207</sup> That is, at least to a first approximation, Semantics provides the content of locutionary acts, and this content can remain stable across contexts when none of the expressions in an utterance are relevantly semantically context-sensitive. Cases in which the content of *illocutionary* acts performed using those utterances is unstable do not show that the moderate contextualist is committed to their semantic context-sensitivity.

We can even resist C&L’s conclusions without appealing to notion of a locutionary act. Note that the sense of ‘what is said, strictly speaking’ at work in the last paragraph—i.e., a purely locutionary sense—is not the sense in which we need to understand that notion as it is used to mark the distinction between lying and misleading. As we learned in Chapter 2, we need an illocutionary sense of ‘saying’ for that—and consideration of the nature of the relationship between the two types of saying was a source of significant vexation. But even according to the

---

<sup>206</sup> *op. cit.*, p. 51

<sup>207</sup> *Ibid.*

more robust illocutionary sense of strict saying, some of C&L's examples fail. For instance, they ask us to consider multiple utterances of 'John went to the gym.'<sup>208</sup> In one context, the topic of conversation is John's nightly walk; in another, it's his exercise regimen. C&L elicit the intuition that in the first case, the fact that John walked to the gym would vindicate the utterance, while in the second, it would not: if he only walked there, but didn't go in to exercise, the utterance would be false. But one can imagine, in such a context, a defense against lying appealing to the illocutionary sense of strict saying: "What I said was, strictly speaking, true, since John did in fact go to the gym. It was, I admit, misleading, though, since he didn't work out when he was there." Whether or not one agrees that the defense goes through in this case, this is certainly a standard use of the notion of strict saying, one that any competent speaker would recognize. I don't think C&L's expressions of puzzlement about the notion look at all coherent in light of such examples, the ubiquity of which is now familiar to us from our explorations in Chapter 2. Moreover, if we help ourselves to a merely locutionary sense of strict saying, the example is utterly straightforward: it seems quite plausible to say that 'John went to the gym', minimally construed, is clearly true in both cases, with Semantics demanding only that John changed his location from one not gym-proximate to one in the gym's general vicinity.

C&L deploy the same sort of slippery-slope argument against so-called Incompleteness Arguments as they do against CSAs. Incompleteness Arguments, as C&L understand them, are another way for moderate contextualists to establish the semantic context-sensitivity of non-Basic Set expressions, this time by pointing to the alleged fact that, without contributions from context, utterances containing them cannot express full truth-evaluable propositions. The examples are familiar: e.g., 'Rose isn't tall enough' doesn't express a proposition unless context

---

<sup>208</sup> *Ibid.*, pp. 44 – 45

tells us what she isn't tall enough for. C&L claim that this pattern of reasoning can be extended to any expression whatsoever, even those only a radical would consider context-sensitive—and so once again the adoption of the moderate contextualists' strategy leads inevitably to Radical Contextualism. They argue, for example, that even an utterance of the seemingly complete 'John went to the gym' could be queried in the same way as 'Rose isn't tall enough': "Went to the gym how? Walked to the vicinity? Did something in the gym? Did what in the gym? For how long? What if he went into the gym but was sleepwalking? Etc. We don't know how to evaluate [the utterance] without settling these questions...."<sup>209</sup>

Once again, this argument founders on the rocks of C&L's too-permissive Speech Act Pluralism—their inability or unwillingness to acknowledge real differences between different kinds of speech acts. Even if we grant that we require answers to further questions to determine the *illocutionary* content of an utterance of 'John went to the gym' (suppose it's a context in which his exercise regimen is under discussion, and so what's asserted by the utterance is that he performed his scheduled workout; it could be argued (see above) that this is not part of what is said (in the illocutionary sense), but rather a mere implicatum, but we'll set that aside for now), it's open to the moderate contextualist to appeal to the notion of the utterance's *locutionary* content, and say, quite plausibly, that Semantics delivers a perfectly truth-evaluable proposition, one that's true just in case John went to the gym—no matter how he got there, what he did there, or whether or not he was sleep walking at the time. The difference between this kind of case and those in which there is genuine semantic incompleteness is that for the latter, what we're given

---

<sup>209</sup> *Ibid.*, pp. 64 – 65



even at the level of locutionary content<sup>210</sup> requires that we answer further questions before we can evaluate its truth or falsity; it is not fully propositional.<sup>211</sup>

Now, C&L might object that there's no principled way to draw a distinction between sentences that are semantically incomplete and those that are not. 'Rose isn't tall enough' is maybe a relatively easy case, but what about 'It's raining'? This seems to require a specification of location in order to be truth-evaluable. Does that make it semantically incomplete? Perhaps its semantic content is the proposition that it's (currently) raining somewhere or other. Or, if a location is required, maybe it can be treated in the same way as sentences containing indexicals—as having an argument-place (in this case, covert) in its logical form demanding a specification of location. There are many difficult cases to consider (see Bach 2001, section 5 for a sample), but the lack of a clear criterion for deciding them all is not evidence that there's no such thing as semantic incompleteness. This sort of situation is (painfully) familiar to us: after an exhaustive (and exhausting) search in Chapter 2, we failed to find a principled way to distinguish lies from acts of mere misleading. There turned out to be a variety of difficult, borderline cases; the distinction resists theoretical codification. But that does not mean there's no such distinction! Similarly for semantic incompleteness: perhaps the best approach is to consider alleged examples on a case-by-case basis, consulting a wide range of data—from intuitions about what is said and

---

<sup>210</sup> I deliberately used the phrase 'what we're given at the level of locutionary content' rather than something like 'semantic content', because there are distinct views on this. For Bach, self-described radical semantic minimalist, there is no difference between the two phrases; locutionary content just is semantic content, and in the case of semantic incompleteness, that content is a propositional radical rather than a full-fledged proposition. Some moderate contextualists, though, would characterize as semantic the fleshed-out propositional content provided with input from context. C&L are blind to this distinction: to them, Bach is a moderate contextualist (a designation to which he objects in Bach 2006).

<sup>211</sup> At least, *arguably* non-propositional. C&L will maintain, of course, that there's a minimal proposition in such cases, but that claim, as we've seen, rests on some pretty dubious metaphysics. The thrust of the present objection, though, does not depend on metaphysical arguments against minimal propositions: the point is simply that there's a natural way to distinguish between the two examples in terms of locutionary content (one requires further questions to determine it, the other does not); failure to acknowledge this can be traced to C&L's SPAP, not their metaphysics. (Thanks to James Shaw for prompting this clarification.)

truth-values, to linguistic and psychological data. The difficulty of the enterprise does not show that it's hopeless. And if the only alternative is to accept C&L's minimal propositions—and they tell us that this is, in fact, the only alternative—then given the problems we've already identified with those, I'll take semantic incompleteness.

### 3.1.5 Summing up: how to use intuitions

C&L believe it's a mistake to assume that a semantic theory must explain our intuitions about speech act content. They codify this belief and call it the "Mistaken Assumption": "A theory of semantic content is adequate just in case it accounts for all or most of the intuitions speakers have about speech act content, i.e., intuitions about what speakers say, assert, claim, and state by uttering sentences."<sup>212</sup> This assumption may indeed be mistaken if we take it to demand of Semantics that it account for our intuitions about the contents of illocutionary acts—assertions, claims, and so on. Our linguistic practices allow such variability, such flexibility in the modes of expression we deploy in the performance of such acts that it seems too daunting a task for a formal semantic theory to account for it all. But it is not clearly mistaken if we demand of Semantics only the more modest task of accounting for our intuitions about locutionary acts—intuitions, for example, about the cancellability of generalized implicatures, which may in some cases (see Chapter 2) be part of asserted, illocutionary content, but whose non-minimal, (arguably) extra-semantic status can be revealed with the appropriate cancelling clause. C&L affect an extreme wariness about relying on intuitions: "[I]f you want to exploit intuitions about speech act content to fix semantic content, then you have to be extremely careful in so doing. It

---

<sup>212</sup> *Ibid.*, p. 53

can be done, and we'll show you how, but it's a subtle and easily corrupted process."<sup>213</sup> Despite this and similar declarations, though, I think a lack of care in handling intuitions is their undoing. Recall our methodological discussion from §2.1.1 above. There are in the literature two alternative approaches for getting at speech act content, what is said; according to one, we should rely solely on intuitions about the appropriateness of indirect reports; according to the other, it's better to focus on intuitions about the truth-values of utterances in various contexts. C&L emphatically declare their allegiance to the first school of thought; indeed, they find the second approach incomprehensible: "We don't know what it is to have intuitions about the truth value of utterances as such. If we are asked to have intuitions not about what an utterance says, asserts, claims, etc. but just about its truth value, we are at a loss."<sup>214</sup> But intuitions about the truth-values of utterances are robust and unavoidable—as C&L themselves demonstrate when they formulate their tests for genuine context-sensitivity (Test 3: an expression "*e* is context sensitive only if there is a true utterance of an instance of the following schema for Inter-Contextual Disquotation (ICD, for short; where *S* contains *e*): (ICD) There are (or *can* be) false utterances of '*S*' even though *S*."), and when they try to show that CSAs and Incompleteness arguments put us on a slippery slope to Radical Contextualism (we're asked, for example, to note that an utterance of 'John went to the gym' has different truth-values in different contexts). My own view is that the best approach is to combine the two sorts of intuitions—about propriety of indirect reports and truth-values of utterances—and that the best kinds of examples to consider involve the lying/misleading distinction. Moderated by a consideration of the truth-values of original utterances in context, intuitions about the appropriateness of various reports reveal a range of subtle, surprising, and vexing phenomena. Despite their implicit endorsement (through actual

---

<sup>213</sup> *Ibid.*, p. 4

<sup>214</sup> *Ibid.*, p. 98

practice) of the importance of truth-value intuitions, though, C&L steadfastly ignore them at crucial stages. This leaves them with an un-moderated set of intuitions about reporting practices, and from the liberality of those practices (the variety of reports of the form ‘S said that *p*’ that are often acceptable, the frequent inter-changeability of ‘say’, ‘assert’, ‘claim’, etc.) they derive conclusions that make them blind to distinctions among speech-act types, with disastrous results. The basic circumspection at the heart of Speech Act Pluralism is laudable: “We have been unable to figure out how to devise an algorithm that takes the proposition semantically expressed and delivers all the propositions said, asserted, etc. There might not be *any* systematic theory from which one can derive all of which [*sic*] is said by an utterance.”<sup>215</sup> Replace ‘proposition semantically expressed’ with ‘locutionary content’, and this passage becomes something we might have said to sum up the results of Chapter 2’s meanderings. But C&L’s wanton conflation of locutionary and illocutionary saying leads to many absurdities: minimal propositions are asserted despite being metaphysically mysterious, as are any contents—whether or not they’re known or believed by speakers—that can feature in indirect reports of their speech; so, for example, if I say, “Alice is playing with Shep,” and it turns out that, unbeknownst to me, Alice is the second-cutest 3-year-old in North America<sup>216</sup> and Shep, a doll we bought second-hand, used to belong to Chelsea Clinton, and was her favorite, then I have *asserted* that the second-cutest 3-year-old in North America is playing with Chelsea Clinton’s favorite doll—I’m committed to the truth of that content. But of course I’m not. Their basic error is to mistake the truth of a variety of *de re* assertion-reports for evidence of a corresponding variety of assertions; but there is no such variety, as the falsity of *de dicto* versions of the same reports makes clear. And besides committing them to such absurdities about speech act content, C&L’s flawed intuition-

---

<sup>215</sup> *Ibid.*, p. 200

<sup>216</sup> I had her first, naturally. Who is this other kid, I wonder?

mongering makes them blind to the sorts of distinctions that would allow Moderate Contextualists—and minimalists like Bach, et. al., who countenance sub-propositional semantic content—to resist a slide into Radical Contextualism. This has the disastrous effect of undermining the core of their argument.

### **3.2 ALTERNATIVES TO CAPPELEN AND LEPORE**

C&L try to resist the onslaught of the marauding pragmatic hordes by drawing a tight perimeter around their semantic stronghold, ceding some territory in order to shore up their defenses. They abandon almost entirely the traditional connection between semantic content and speech act content (condition 3 above) and try to protect the rest of their semantic redoubt. The maneuver is unsuccessful, however; their fortifications are insufficient, and their encampment suffers from internal strife.

I want to turn to an examination of some other ways of conceiving of the semantic project. The first, due to Emma Borg, represents a potential way to rescue the minimalist semantic project. The others, from Jason Stanley and Katarzyna Jaszczolt, respectively, are more expansive conceptions of Semantics (the latter more so than the former). I've chosen these three views to try to get a representative sample of the wide variety of positions within this logical space; there are many other views, and many variations of roughly similar approaches, but these three should give us a good idea of the choices that are possible, and the challenges that arise for each. The aim here is not a comprehensive evaluation of the alternatives; rather, I'll give a brief sketch and note some potential problems, noting for each any special issues that arise in the connection with our distinctive focus on the lying/misleading distinction.

### 3.2.1 Borg's semantic minimalism

Borg stresses the importance of the first of our conditions on Semantics, particularly the requirement that it account *systematically* for the relationship between the syntactic arrangement and semantic contents of the words in a sentence, on the one hand, and the sentence's semantic content on the other. She demands that Semantics be a *formal* theory: it abstracts from the features of the uses of expressions, considering only their formal properties (meanings of individual lexical items, mode of their combination, i.e., sentential syntax) and providing an account of how literal sentence-meaning (as opposed to use-based utterance-meaning) can be *computed* from these basic formal ingredients. The reason she demands such a conception of Semantics is that she's committed to a computational, modular conception of the mind, in the style of Fodor 1975 and 1983, according to which, very roughly, reasoning is a process of formal computation—manipulation of mental representation according to their syntactic properties—and the mind is divided into discrete modules carrying out distinct reasoning processes, one of which is a dedicated language faculty.<sup>217</sup> In Borg's words: "A module is then a probably innate, encapsulated body of information together with processes operating only over that information, which is responsible for realizing a given cognitive function."<sup>218</sup> A formal semantic theory dovetails with this picture of the mind.

If semantic content is the output of such a processing module, there are serious constraints on its character. Consider the processes by which we come to an overall interpretation of the communicative content of a speaker's utterances, including implicatures,

---

<sup>217</sup> The idea of an innate language faculty operating along these lines has its roots in the early work of Noam Chomsky, e.g. Chomsky 1957 and 1965.

<sup>218</sup> Borg 2004, p. 80

insinuations, indirect speech acts, and so on. We bring to bear on this interpretive task prodigious amounts of information, not only about lexical meanings and so forth, but also about the speaker's likely intentions, the current state and context of the discourse, and background information about the world, culture, etc. Putting all this together, we make an inference to the best explanation for what the speaker meant to convey with her words. This kind of abductive, all-things-considered reasoning, however, is not suited to the formal, computational, modular model that Borg envisions for semantic processing. Inferences in that mode must be deductive and draw only on the specialized information contained in the linguistic module. This condition severely constrains what can count as semantic content. In order to determine, for example, what's referred to by a demonstrative pronoun ('this' or 'that'), we must consult speakers' intentions.<sup>219</sup> Even the contents of so-called (by Kaplan) "pure" indexicals like 'here' and 'now' are not simply determined by objective features of context, the location of the speaker and the time of utterance; speakers can intend different areas of space and periods of time with these locutions, and so consultation of intentions again becomes unavoidable. (Compare 'Now I can surf the web on my phone', where I mean something like 'these days', to 'Now the steak is done', where I mean 'this very moment'.) So it turns out that even C&L are too liberal, on this conception, in granting the contextual determination of the referents of members of the Basic Set status as properly semantic. Borg says, "[T]he true test for minimalism is... how it deals with the existence of those expressions for which it seems uncontroversial that their content is bound up with current speaker intentions... [i.e.,] how it deals with overtly context-sensitive expressions like demonstratives and indexicals."<sup>220</sup> C&L fail the test.

---

<sup>219</sup> At least arguably. There's a literature on this question.

<sup>220</sup> Borg 2007, p. 357

To pass it, one must adopt an extremely austere conception of semantic content. According to Borg, the contents of demonstratives and other indexical expressions are singular (as opposed to general) concepts, grasped by the hearer at the time of utterance, the contents of which are the objects referred to—but, and this is the key point, the object referred to is not specified in anything but a descriptive manner. An example will help. When someone hears an utterance of ‘That is mine’, the semantic content entertained will be something like “ $\alpha$ , whatever it is, belongs to S, whoever that may be.”<sup>221</sup> Or the value of ‘This is red’ will be something like ‘ $\alpha$  is red’, or ‘the object ostended by the speaker of the utterance is red’. In these examples,  $\alpha$  is the singular concept, and it could refer to any number of objects.<sup>222</sup> Borg takes it that this sort of singular-concept account extends to singular terms in general, even proper names.

Rather than posit C&L’s comparatively more robust, allegedly more concretely specified, minimal propositions, Borg opts for what she calls “liberal truth-conditions.” She deals with apparent semantic incompleteness, for example, not by pounding the table and insisting that ‘Jane can’t continue’, for instance, expresses the (minimal) proposition that Jane can’t continue, and is true if and only if Jane can’t continue, but by offering the schematic, liberal truth-condition: “If  $u$  is an utterance of ‘Jane can’t continue’ in a context  $c$  and ‘Jane’ in  $u$  refers to  $\alpha$  then  $u$  is true iff  $\alpha$  can’t continue something.”<sup>223</sup> This doesn’t specify a single state of affairs, but rather a range of them that may satisfy the conditions. Specification of which precise state is the one indicated by the particular utterance is a task that falls outside the province of Semantics; it requires the consultation of information outside the ambit of the language faculty.

---

<sup>221</sup> Borg 2004, p. 197

<sup>222</sup> This is similar to John Perry’s notion of “reflexive content.” See Perry 2001, Korta and Perry 2011.

<sup>223</sup> *Ibid.*, p. 230, n. 23. Remember, even proper names are treated as having singular concepts as contents. Also, relativization to context here is due to the demonstrative aspect of the use of present tense.



That Borg's liberal truth-conditions are unlike what one normally thinks of as semantic content is apparent. She recognizes this, and discusses at length a series of objections that naturally arise in reaction to her view, all to the effect that liberal truth-conditions are not *real* truth-conditions—they're too schematic, they don't allow one to determine a truth-value in all possible situations, they fail to match intuitive truth-conditions in particular circumstances, and so on. Most of these objections, though, are just variations on the complaint that liberal truth-conditions are not the intuitive truth-conditions of utterances. But they were never meant to be. Borg is self-consciously separating the domain of Semantics from the domain of speech-act interpretation; condition 3 on Semantics—that it explain our intuitions about speech acts—is denied utterly. At most, Semantics provides a starting point for utterance interpretation, but “by the time we reach any judgment of what is said by a given utterance in a given context the peculiarly semantic contribution will have been swamped beneath the influence of other (non-semantic) factors.”<sup>224</sup> Determination of speech-act content is an almost entirely pragmatic affair. “‘Saying’ or ‘asserting’ are pragmatic notions.... [I]t is no longer (if it ever was) thought to be the job of a semantic theory to explain different speech acts or what a speaker says by a given utterance. The task of semantics is to capture the literal meanings of words and sentences, nothing more.”<sup>225</sup>

Borg's is an even more resolute minimalism than C&L's, and this is a mark in its favor, in my view. It allows Borg to avoid the absurd claim that one of the propositions we assert when we utter a sentence is the semantically minimal proposition associated with it—a claim that is in clear contradiction with the distinction that is our focus, between lying and misleading. If the semantic content of ‘Rose isn't tall enough’ is something like ‘Rose, whoever that might be, is

---

<sup>224</sup> *Ibid.*, p. 128

<sup>225</sup> *Ibid.*, pp. 130 – 31

not tall enough for something or other’, this content is easy to make true; she’s tall enough to look down at an ant, for example. But if we make this truth part of what we assert whenever we utter a sentence, we get a specious defense against lying (at least a partial defense) in circumstances in which an utterance of the sentence would clearly be false (standing in front of the Tilt-A-Whirl, e.g.). C&L are committed to this absurd result; Borg is not.

All this is not to say that I want to embrace Borg’s account of Semantics entirely. The purpose of this chapter is not to reach such definitive conclusions, but rather to chip away at the issues a little bit. And I have some reservations about Borg’s view. First, it’s not clear to me that we have much reason to prefer her version of minimalism to that of Bach et al., where we abandon condition 2 and allow sub-propositional content to count as the meaning of certain sentences. Liberal truth-conditions, being satisfiable by a range of different circumstances, are hard, on the face of it, to distinguish from propositional radicals or blueprints, which are also sufficiently indeterminate to allow for multiple ways of fleshing out. Moreover, as Bach points out, while adding an argument place to the syntax of a sentence to be filled with an existential quantifier is intuitively plausible in some cases, it fails to be in others. Compare ‘Jack is glad’ to ‘Jack is pleased’; in each case, we may specify something further that he’s glad or pleased about. But while ‘Jack is pleased about something or other’ seems fine, ‘Jack is glad about something or other’ seems fishy. It makes sense to say “Jack is pleased, but I have no idea about what,” but “Jack is glad, I have no idea about what,” is awkward. When we say someone is glad, we have to have something specific in mind, and so liberal truth-conditions are inappropriate here in a way that they’re not with ‘Jack is pleased’.<sup>226</sup> And Borg’s objections to sub-propositional content

---

<sup>226</sup> Bach 2007, pp. 305 – 06

commit the very error I praised her above for avoiding; namely, she takes semantic content to be something a speaker expresses as part of a speech act. She says:

If semantics gives us nothing but a gappy proposition we surely still want to ask which of these contextually provided fillings-in the speaker genuinely expresses or commits themselves to. Yet since we have given up the idea that what counts for literal content is what can be found at the lexico-syntactic level, it becomes quite unclear how we can isolate just one (or some) of all the possible expansions of a propositional radical as the one(s) that delivers literal meaning.<sup>227</sup>

Her preferred account, according to which we assign liberal truth-conditions as semantic content, is only an advantage over gappy propositions if it provides what Borg thinks is needed: a specification of what the speaker is committed to; this must be what she means here by “literal content/meaning.” But this is the content of an illocutionary speech act (the invocation of commitment makes it so), and liberal truth-conditions are no better positioned to provide this than propositional radicals. If she wants to claim that her minimal semantic content can serve such a role, she’s committing the same blunder as C&L did when they said that minimal propositions are always asserted.

My second reservation about Borg’s minimalism is its commitment to the computational, modular conception of mind. I’ve already noted that I’m reluctant to embrace that view (see the discussion of Relevance Theory above), though this is not the place to offer arguments against it. Instead, I just want to emphasize the degree to which Borg’s vision for Semantics is beholden to it. She freely admits that the semantic content generated by her linguistic processing unit has no essential role to play in the conscious process of utterance interpretation (her §2.3.4 is titled “There is limited conscious access to the intermediate representations of the language faculty”); rather, we often skip past these literal meanings, consulting pragmatic information to arrive at intuitive meanings for utterances. If this is so, then semantic content fails to satisfy Recanati’s

---

<sup>227</sup> Borg 2012, p. 527

“Availability Principle,” according to which conscious availability is a necessary condition for psychological reality (see §2.5.3 above). Now, as we saw in the case of locutionary content, failure always to be present to consciousness need not imply that there is no such thing, as one can train oneself to become aware of such content (in the case of locutionary content, by considering the cancellability of extra material). But Borg’s singular concept-laden liberal truth-conditions are alien to everyday linguistic practice in a way that simple locutionary contents are not, so it’s less clear that such a move is available to her. Consideration of cancellability and so forth seems more likely to lead to Bach’s more intuitive propositional radicals in the relevant cases. Only if, as a matter of empirical, psychological fact, there is a linguistic processing module operating exactly as Borg describes can we have any guarantee that her semantic contents are real.

### **3.2.2 Indexicalism**

Indexicalism is a comparatively bold proposal, in that it aims to satisfy all four conditions that tradition supposes a semantic theory should: it is compositional and syntactically constrained; it yields truth-evaluable propositions as contents of declarative sentences; these comport with our intuitions about speech acts; and they provide an input to the usual Gricean pragmatic processes. Its most prominent exponent is Jason Stanley (along with collaborators like Jeff King and Zoltán Gendler-Szabó). Though Stanley himself doesn’t refer to his view as “Indexicalism,” it’s a natural enough designation, as his account is merely an expansion of the standard Kaplanian treatment of indexical expressions. That is, he deals with all the apparent context-sensitivity of the contents of speech acts by positing indexical-like elements in the logical form of the sentences used to perform those acts, thus rendering the acts’ contents properly semantic, insofar

as their determination is triggered by syntactic considerations. Stanley does not share Borg's scruples about the need to consult speaker-intentions, since his theory is not so closely tied to a specific view about the nature of mental linguistic processing, and so he can allow a wide range of phenomena to fall under the semantic umbrella.

Some examples easily illustrate the general features of the approach. In their 2000, Stanley and Szabó take on the phenomenon of quantifier domain restriction—the ability of speakers to use sentences with apparently unrestricted quantifiers, like ‘Every bottle is empty’, to make claims in which the quantifiers’ scope is limited, e.g., ‘Every bottle I just bought is empty.’ In a context in which Lisa utters the former having just returned from the store with some bottles, and her friend Max, who needs empty bottles to fill with his homemade beer, has just asked her whether the bottles need to be emptied before he can use them, she successfully conveys the proposition expressed by the latter.<sup>228</sup> The challenge is to show that this intuitive content of Lisa's speech act is determined by Semantics. Their key idea is that the logical form of the uttered sentence contains a hidden element—hidden in that there's nothing phonetic in the sentence to which it corresponds—that triggers the consultation of context to specify the domain over which the quantifier is supposed to range. This element is a variable that takes different quantifier-domains as values and is associated grammatically with the noun.<sup>229</sup> Context tells us which, of all the bottles in the universe, we are to consider: the ones she just purchased. This syntactically triggered appeal to context is exactly analogous to the manner in which indexical expressions trigger such an appeal, except that the relevant element of syntax is hidden—it is, to

---

<sup>228</sup> Stanley and Szabó 2000, p. 231

<sup>229</sup> I'm deliberately setting aside various complications, such as the need to conceive of the variable element as a function, and the question of whether the variable co-habits a terminal node with the noun, or has one of its own (hence the intentionally vague ‘associated with’). Nothing I have to say depends on attending to these niceties, and so I won't.

use Neale's terminology, *aphonic*. This general approach allows Stanley et al. to deal with a variety of instances of context sensitivity. For example, the theory can be deployed to explain a phenomenon that we've returned to again and again—relative adjectives' apparent context-sensitivity. Different utterances of 'Rose is tall' can express different propositions in different contexts, depending on the salient comparison class. If one postulates a hidden variable in logical form that demands the specification of a comparison class, the phenomenon becomes properly semantic.<sup>230</sup>

There are various quibbles one could make with the formal details of Stanley's theory, though I'm not interested in those and suspect they can be dealt with. I'm willing to stipulate that the details of the (rather impressive) logico-syntactic apparatus are entirely on the up-and-up. The deeper concern one might have with this approach is that it's question-begging against the theorist who maintains that the relevant phenomena are pragmatic, not semantic. One might worry that hidden indexicals are a *mere* postulate, in the sense closer to that term's Latin root, *postulāre*—to demand—than to the sense in which it's synonymous with 'axiom'. That is, it's only a prior conviction that the observed context-sensitivity must have a semantic explanation that motivates the introduction of aphonic indexicals; but this is not evidence that the phenomena aren't pragmatic, especially since the explanans are invisible. What is required to ameliorate this concern is an independent argument for the existence of hidden indexicals. Stanley recognizes this challenge: "[S]yntactic structure cannot simply be postulated on semantic grounds. Rather, evidence of a syntactic sort must be available for the existence of domain variables."<sup>231</sup> The syntactic evidence Stanley musters for the existence of hidden variables is that they seem to interact with—be bound by—quantifiers. For example, the domain variable in 'Every bottle is

---

<sup>230</sup> Again, the details are considerably more complicated. See Stanley 2007, Ch. 3.

<sup>231</sup> Stanley 2007, p. 114.

empty' can be bound by the quantifier in 'In some rooms, every bottle is empty', which is naturally understood to convey something like 'In some rooms, every bottle *in that room* is empty'. The italicized material makes explicit the value that the domain variable must take, now that it's been bound; what was heretofore hidden is now manifest.

The Binding Argument is supposed to generalize beyond quantifier sentences. So, in his 2000, Stanley discusses 'It rains', and concludes 'rains' co-occurs with both temporal and locational variables.<sup>232</sup> Also, comparative adjectives have variables for comparison-classes associated with them.<sup>233</sup> Evidence for the existence of these variables in LF is provided by their bind-ability, based on the Binding Assumption: "If  $\alpha$  and  $\beta$  are within the same clause, and  $\alpha$  semantically binds  $\beta$ , then  $\alpha$  either is, or introduces, a variable-binding operator which is co-indexed with, and stands in certain specified structural relation to, a variable which is either identical to, or is a constituent of,  $\beta$ ."<sup>234</sup> The upshot is, where there's binding, there's a variable, possibly hidden in LF. So, for example, it seems that the natural reading of 'Everywhere I go, it rains' is 'for all locations  $x$ , if I go to  $x$  then it rains at  $x$ .'

There are reasons to be concerned about this argument from binding. First, as C&L argue, it over-generates; that is, it leads us to conclude that there are hidden variables present when, intuitively, there shouldn't be. Consider 'Everywhere I go,  $2 + 2 = 4$ '.<sup>235</sup> By parity of reasoning, this sentence provides evidence of a hidden variable for location in ' $2 + 2 = 4$ '. That's surprising! Furthermore, we can keep adding quantifier phrases to generate evidence for indefinitely-many hidden variables in a sentence's logical form: 'Everywhere I go, no matter when I go there,  $2 + 2 = 4$ ' indicates the presence of both location- and time-variables;

---

<sup>232</sup> Stanley 2000, p. 416

<sup>233</sup> *Ibid.*, p. 419

<sup>234</sup> *Ibid.*, p. 412

<sup>235</sup> Cappelen and Lepore 2005, p. 74

‘Everywhere I go, no matter when I go there, and no matter what I ate for breakfast,  $2 + 2 = 4$ ’ adds a variable for the speaker’s breakfast; we could keep going. There are moves that Stanley can make to avoid this conclusion. He can (and does, in his 2007), deny that ‘for all locations  $x$ , if I go to  $x$ ,  $2+2=4$  at  $x$ ’ is the proper reading (Jonathan Cohen and Sam Rickless make the same argument in their 2007). He can also argue that the objection misinterprets the Binding Argument, making it stronger than it was intended to be: C&L presume that it posits hidden variables as necessary to explain the grammaticality of the sentences in question; Stanley (again, see his 2007), can insist that bound variables are only arrived at via an inference to the best explanation for the binding phenomena he examines.<sup>236</sup> This is not the place to get into the merits or demerits of these moves and counter-moves; my only intention in presenting these objections is to indicate that the Binding Argument is not universally accepted—that there *may* be problems with it, and the issues are complicated.

Lenny Clapp registers a similar objection, this one focusing on relative adjectives. The following succession of sentences suggests an indefinite number of hidden variables: “(17) Most species have members that are small. (small for the species) (17’) Most species have some females who are small. (small for females of the species) (17’’) There is some region where for most species several females are small. (small for females of the species who live in the region)”<sup>237</sup> It seems that binding arguments force us to conclude that any given sentence contains a plethora a hidden indexical variables—possibly infinitely many!<sup>238</sup> This is a troubling

---

<sup>236</sup> Stanley 2007, p. 221

<sup>237</sup> Clapp 2012, pp. 452 – 453, n. 21

<sup>238</sup> It is not open to Stanley to say, e.g., “We don’t need an infinite number of variables to be there ahead of time—just a recursive method for adding them in proportion to the number of quantifiers in any given sentence.” The hidden variables identified by the binding test aren’t just there when the expression appears in the scope of a quantifier; they’re there in the LF of the expression even if it occurs in isolation. If they weren’t, Stanley’s position wouldn’t be an alternative to an unarticulated-constituent view, or a pragmatic approach; it’s crucial for Stanley that the expansion of ‘It’s raining’ to ‘It’s raining at such-and-such’ be *mandatory*—dictated by the syntax of the stand-



conclusion about the grammar of our language. How does one learn such a language, for example?<sup>239</sup>

If any of the various objections to the Binding Argument stick, then, deprived of independent motivation for the postulation of hidden variables in logical form, Indexicalism must face the charge of begging the question. And the charge can be made vehemently. Neale is particularly scathing: “The only substantive difference between the way the [advocate of a pragmatic theory] sees the process of identifying the proposition expressed and the way someone postulating aphonics in syntax sees it is that the latter is just *insisting* that the search for the integration of contextual information in the interpretation process is triggered syntactically.”<sup>240</sup> It gets worse in a footnote:

In effect, then, [Stanley’s] proposal is nothing more than a pointlessly formal and absurdly syntactic way of saying that interpreting an utterance of, say, ‘Every philosopher explained several theories to every linguist’ involves identifying which class of philosophers, which class of theories, and which class of linguists are being talked about. But that is precisely what [the pragmatic theorist] has been saying all along, only without the syntactic palaver and dogma.<sup>241</sup>

It’s useful to recall Borg’s reservations, and Stanley’s lack thereof, about countenancing as even potentially semantic the sorts of richly inferential processes involved in, say, determining the referent of a demonstrative pronoun. If we allow it, then we need strong arguments in favor of syntactic triggering to distinguish the view that such interpretation is a semantic matter from the view that it’s pragmatic. I’m not sure the Binding Argument is up to the task.

---

alone expression—rather than optional (Recanati would call it “free enrichment”). This only works if the variables (whose presence is revealed with binding tests) were there all along. (See Recanati 2004, Ch. 7) Thanks to James Shaw for suggesting the objection addressed here.

<sup>239</sup> Cappelen and Lepore 2005, p. 75

<sup>240</sup> Neale 2007, p. 82

<sup>241</sup> *Ibid.*, p. 83, n. 12

My final concern with Stanley’s program has to do with its ambition. His aim seems to be to maintain the traditional roles of Semantics in determining both the contents and truth-conditions of our speech acts—of what we say or assert. Semantics, moreover, must save the phenomena, in that its output must correspond to our intuitions about content and truth-conditions. “[A]ccounting for our ordinary judgments concerning the truth-conditions of various sentences is the central aim of semantics.”<sup>242</sup> “What results from a successful application of [the semantic] stage of interpretation is a unique proposition, a fully truth-evaluable entity. Furthermore, if the process is successful, the proposition it yields is one the interpreter would recognize as the proposition expressed by [the speaker’s] assertion.”<sup>243</sup> Now, even if we stipulate that Stanley has successfully accounted for some phenomena within a plausible semantic framework—quantifier domain restriction, relative adjectives, and a few others are explicitly tackled in his writings—there remains a host of others to consider. We were exposed to the overwhelming variety in the last chapter. Stanley is optimistic: “I... hope to provide convincing evidence of the promise of the project of reducing all apparent effects of context on semantic content to a small number of sources.”<sup>244</sup> I don’t share Stanley’s optimism. It’s not just the sheer variety of phenomena to be accounted for that makes me pessimistic; it’s their evident unruliness. If there was one definitive result in Chapter 2, it was this: the sorts of linguistic constructions that exhibit apparent sensitivity to context are a motley bunch resistant to systematization; and even within subsets of the motley of phenomena that can be given a unified analysis (Q-type GCIs, for instance), our intuitions about content and truth-conditions exhibit no regularity. It’s these sorts of intuitions that Stanley aims to account for. But as we’ve seen,

---

<sup>242</sup> Stanley 2007, p. 90

<sup>243</sup> Stanley 2002, p. 150

<sup>244</sup> Stanley 2007, pp. 111 – 12

there's a prior question to be answered about how to tap into those intuitions. Should one focus on indirect reports of speech? This can lead to the excesses of C&L's Speech Act Pluralism. Or should one only consider intuitions about truth or falsity of utterances? This approach runs the risk of making us blind to merely locutionary sayings, which may be intuitively more closely associated with literal linguistic meaning. Stanley's use of 'assertion' above suggests that he's more interested in illocutionary content. My hybrid approach, relying on intuitions about cases in which the lying/misleading distinction is salient, and which seeks to retain the virtues of both alternatives, also sought an illocutionary sense of 'saying' (only this, we saw, was suitable to underwrite the lying/misleading distinction), but revealed along the way that there was much of interest in a merely locutionary notion, and many vexing questions about the relationship between the two. For Stanley's project to succeed, it must reckon with all these issues. Chapter 2 showed that "our ordinary judgments concerning the truth-conditions of various sentences"—accounting for which is the "central aim of semantics"—are an extraordinarily complex, subtle matter, with ambivalence often being the prevailing intuition. Stanley may have succeeded in providing a plausible syntactic/semantic account for a few phenomena—I can't, for instance, imagine a scenario in which one would convey a falsehood by restricting a quantifier domain in one way, and then plausibly defend himself against lying by claiming to have made a claim involving a less-restricted domain—but the relative modesty of this achievement compared with his stated ambitions is thrown into sharp relief against the background of the data presented in Chapter 2. That's why I'm pessimistic.

### 3.2.3 Default Semantics

Having spent so much time considering various minimalist approaches to Semantics, it's appropriate that we finish our discussion with look at a view that may as well be called *Sematic Maximalism*. Jaszczolt's Default Semantics is such a position. She throws off the shackles of the first traditional condition on semantic theories—that they be constrained by syntactic considerations—and re-conceives Semantics as open to input from a variety of non-traditional sources.

The 'defaults' in Default Semantics (DS, hereafter) are "salient meaning[s] intended by the speaker, or presumed by the addressee to have been intended, and recovered (a) without the help of inference from the speaker's intentions or (b) without conscious inferential process altogether."<sup>245</sup> They are the sorts of speaker-meanings that Generalized Conversational Implicatures are a prime example of—the sorts of implicatures that go through automatically, or by default. They are aspects of speaker-meaning that are assumed to be present *unless* the context is unusual in some way, such that the implicature is cancelled. For Jaszczolt, the proper subject matter of Semantics is not sentence-meaning, but rather (contextualized) utterance-meaning, the primary meanings that speakers intend to convey with their utterances. The guiding idea is that defaults provide an important input to the determination of such meanings. Specifically, defaults from two sources—innate properties of the mind and background knowledge about the world—conspire with the outputs of other processes—computations over lexical meanings and syntactic structure, and certain pragmatic inferences—to deliver an interpretation of a speaker's primary intended meaning. Defaults of the first kind, "cognitive

---

<sup>245</sup> Jaszczolt 2010(a), p. 1

defaults”, include the preference for referential over attributive interpretations of definite descriptions, and the preference for *de re* over *de dicto* readings of propositional attitude ascriptions. Defaults of the second kind, “social, cultural, and world-knowledge defaults”, yield the following interpretations (of the a-utterance as expressing the b-proposition), for example:

(3a) The temperature fell below -10 degrees Celsius and the lake froze.

(3b) The temperature fell below -10 degrees Celsius and as a result the lake froze.

and

(5a) A Botticelli was stolen from the Uffizi last week.

(5b) A painting by Botticelli was stolen from the Uffizi Gallery in Florence last week.<sup>246</sup>

The former default results from our knowledge of physical laws, the latter from our knowledge of Botticelli’s medium of choice and the nature and location of the Uffizi. The problem arising from the fact that not all interlocutors are likely to know where the Uffizi is, or for that matter who Botticelli is (or, sadly, the temperature at which water freezes) is taken care of by identifying the meanings under investigation as those intended and recoverable by the Model Speaker and Model Addressee, respectively—an idealization that allows for assumptions about background knowledge. Also, the fact that conscious, inferential processes are allowed to contribute to overall meaning on equal footing with unconscious default sources, gives the model flexibility to account for differences in background knowledge among different people.

This point about the different sources of information contributing to the overall interpretation is important and novel, and should be emphasized. Background knowledge about the world, innate properties of our cognitive system, the current situation of the discourse, and knowledge of word meanings and syntax are all co-equal contributors to semantic content, none

---

<sup>246</sup> Jaszczolt 2010(b), p. 7

more fundamental than the others. The sum of their contributions, the “merger representation”  $\Sigma$ , is the so-called “primary meaning” of an utterance, the main, intended meaning conveyed by the Model Speaker and recovered by the Model Addressee. The most novel, surprising and controversial claim here is that the process normally conceived of as the purely semantic one—that which processes utterances based on knowledge of words’ meanings and syntactic combination—is conceived as only one process among many, and no more fundamental than the others. The principle of syntactic constraint on semantic content is abandoned. This, along with the allowance for pragmatic inference (from world knowledge, the discourse situation, etc.) as a contributing factor in semantic interpretation, gives us even more novel results. Jaszczolt’s analysis of Bach’s famous example of the mother saying, “You’re not going to die” to her son who’s cut himself, identifies not the usual completion ‘You’re not going to die from this cut’ as the content; freed from syntactic constraint, she says the primary meaning is closer to something like ‘There’s nothing to worry about.’<sup>247</sup> Indeed, freed from the bonds of syntax and armed with the tools of pragmatic inference, Jaszczolt even allows particularized implicatures to serve as primary meanings: in Grice’s famous dialogue

A: Smith doesn’t seem to have a girlfriend these days.

B: He’s been paying a lot of visits to New York lately.

she claims that B’s primary meaning is that Smith may have a girlfriend.<sup>248</sup>

Now, there are a number of questions one might have about Jaszczolt’s picture, a number of objections that spring immediately to mind. If the gates are thus opened to pragmatic contributions, doesn’t everything count as semantic content? Jaszczolt’s answer is to distinguish

---

<sup>247</sup> *Ibid.*, p. 4

<sup>248</sup> Jaszczolt 2012, p. 25 [My references are based on pagination of a copy of the paper available on Jaszczolt’s website: <http://people.ds.cam.ac.uk/kmj21/>]

primary meaning from secondary meaning, the latter of which is extra-semantic, additional content conveyed over and above (or along with) primary meanings. “What about compositionality?” one might cry, worried that the demotion from pre-eminence of syntactical processing might make an account of the systematic construction of meaning impossible. Jaszczolt’s claims that “compositionality need not be as strict as the traditional truth-conditional semantics requires” and insists that if one “places the methodological requirement of compositionality on the representation of *utterance* meaning rather than *sentence* meaning” the way is open to predicate compositionality of “of merger representations – the  $\Sigma$ s of primary meanings of utterances, as intended by Model Speakers and recovered by Model Addressees” while admitting that, given the processes that inform such constructions, “semantic composition is necessarily largely ‘pragmaticky’....”<sup>249</sup> One might also worry about cancellability: what distinguishes pragmatic content from semantic, according to the original Gricean picture, is that the former can be felicitously cancelled; even opponents of the basic Gricean approach admit as much, must account for the phenomenon. But if we allow cancellable content to be a part of semantic content, what are we to make of the phenomenon? Jaszczolt argues at length in her 2009, marshalling various examples, that in fact the data more clearly indicate that the cancellability criterion corresponds to the primary/secondary meaning distinction than to the old Gricean distinction between what is said and what is implicated; primary meanings, even if they’re implicatures in the classic sense, resist cancellation more strongly than what is allegedly explicitly said.

I don’t want to get into the details of all these objections and replies and assess their merits (or demerits). I’ll just note that DS is a novel, surprising, and controversial view, at

---

<sup>249</sup> Jaszczolt 2010(b), p. 13

significant odds with many well-developed extant accounts in the literature, and that Jaszczolt has a difficult task in defending it. But maybe she can. My personal concern, from the point of view of the present work, is the effect of Jaszczolt's re-conceiving of Semantics on the distinction between lying and misleading. For her results are deeply at odds with many of the intuitions we relied on in Chapter 2. Let's consider a paradigmatic example of merely misleading (as opposed to lying): I and my wife are driving home from an all-too-rare night on the town, during the course of which we got rather carried away, booze-wise, so that we're both drunk; my wife, let's suppose, is a prominent politician, for whom a DUI conviction would be disastrous; I, on the other hand, am a low-level philosophy teacher, for whom a DUI, while inconvenient, would not be seriously disruptive, career-wise; it is therefore decided that I should do the driving. We get pulled over, and the patrolman comes to the car window and asks me whether I've been drinking tonight. I make a nod in my wife's direction and say, "She's been drinking tonight; that's why I'm driving." The way I understand it, DS would say that my primary meaning in this context would be something like, 'I have not been drinking' (remember, we're not considering actual speakers' intentions, but idealized Model Speakers'; in this example, I'm not a model of anything). This is a particularized implicature, according to the traditional analysis; according to DS, it's the semantic content of my utterance. Here's the problem. My intuition is that what I said to the patrolman is, strictly speaking, true; I didn't lie, but I did mislead. But Jaszczolt makes the falsehood I conveyed the *semantic* content of my utterance—its very *meaning*. This is a serious tension: the meaning of my utterance is a falsehood, but I have a strong intuition that what I said was true, that I didn't lie. This, for me, is the most serious drawback of re-conceiving Semantics as broadly as DS does. It renders the common practice of distinguishing between lying and merely misleading incomprehensible. Implicit in all of our



explorations so far has been a conception of Semantics according to which it provides the fallback position of the mere misleader—or at least, the point toward which a stricter interpretation of his utterance tends. If we adopt DS, we must abandon this conception; there is no semantic redoubt to which the misleader can retreat. This would involve conceiving of lying much more broadly, so that the example above and the other examples of mere misleading in Chapter 2, feature lying instead. All deceptive speech acts become lies. This is perhaps a way of putting a finer point on the concern raised above: if we re-conceive Semantics so broadly, it's not really Semantics anymore; to borrow Jaszczolt's word, it becomes too "pragmaticky." If such a move does such violence to so commonplace a distinction as that between lying and misleading, then perhaps we should not make it.

It is not an impossible move to make, though. We just need to redefine 'lying'—or, I should say, since we haven't explicitly defined lying, relying instead on rather hazy intuitions about the concept, we need to adopt an expansive definition of lying, of which there are many in the literature. It is to an exploration of this literature on how to define lying, and how different definitions are compatible with different positions in the philosophical literature(s) we've been exploring, that we now turn.

#### **4.0 LYING, ASSERTING, MORALITY, AND THE LAW**

So far, we've been preoccupied with theoretical issues in the philosophy of language and linguistics. But our focus on the lying/misleading distinction naturally leads us to consider a variety of other philosophical questions, as well. For example, we have not yet, as just noted, given any sort of definition of lying, relying instead on our pre-theoretical intuitions about the concept. There is, however, a vast literature on the question of how best to define 'lie', and it would be well for us to examine it. We will start with the most common dictionary definition, and look at various minor and major revisions to it, evaluating each in turn. Consideration of these will highlight a lacuna in the literature: while it's often said that lying is a kind of (insincere) asserting, rarely are such pronouncements accompanied by any sort of rigorous account of that speech act. There are, though, many to choose from in the rich literature on assertion. In the second section of the chapter, I will start to fill this gap in the literature on lying, by examining various accounts of assertion and asking how well they account for our focal phenomena. I identify one account in particular (due to Robert Brandom) that seems best suited to the job. After this, in the third section, we will consider another issue that arises quite naturally when one begins to think about the lying/misleading distinction: what is its moral significance? This is too large a topic to deal with adequately in the space we have, but it's a topic we'd be remiss in ignoring. We'll deal with it briefly, in light of a proposed expansion of the definition of lying to include what we've been referring to as merely misleading acts. This definition has some

support in the literature, but is obviously at odds with the usage we've adopted. We'll see that, absent compelling linguistic arguments for making this change, the rationale must include the claim that there is no moral difference between lying and misleading. The alternative case—that the lying/misleading distinction *is* morally significant—turns out to be surprisingly difficult to make. Finally, as a purely practical matter, the lying/misleading distinction *does* have *legal* significance. The legal definition of perjury, for example, sets a very high standard, one that only lies—in a narrow, not-merely-misleading sense—meet. The Supreme Court has ruled explicitly, in *Bronston*, that a false implicature cannot subject a defendant to conviction on perjury charges. Some legal scholars have argued that the *Bronston* case was wrongly decided; others support the ruling. In the last section of this chapter, we will bring to bear any insights we've acquired so far on this debate. We'll examine the *Bronston* decision, as well as a prominent, more recent, case involving a deceptive political advertisement by a member of the Wisconsin Supreme Court.

## 4.1 THE DEFINITION OF LYING

In his 2008(b), Mahon identifies what he calls “the most common definition of lying”:

To lie =<sub>df</sub> to make a believed-false statement to another person with the intention that that other person believe that statement to be true.<sup>250</sup>

This definition is meant to be an improvement over the dictionary (OED) definition of lying, according to which a lie is “a false statement made with the intent to deceive.” This OED definition needs improving because it is subject to a number of counterexamples. First, there are

---

<sup>250</sup> Mahon 2008(b), §1

cases in which the conditions are not met, but which are intuitively instances of lying. Siegler 1966 considers the situation faced by the character Pablo in Sartre's story *The Wall*, who was captured by the authorities and questioned about the whereabouts of his comrade Gris. Pablo tells them Gris is hiding in the cemetery, trying to throw them off his trail, as he believes Gris to be somewhere else. It turns out Gris actually was in the cemetery (and tragic consequences ensue), but the intuition is that Pablo's statement was still a lie, despite its truth; what's crucial is that he believed that it was false. In addition, there are cases which meet the conditions of the dictionary definition, but intuitively are not lies: suppose my neighbor, the lousy cook, has invited me over for lunch the following day; to get out of it, I tell him that I have an important meeting that day, a claim that I believe to be true, but which is intended to mislead my neighbor, since the meeting is very early in the morning and would be over in plenty of time for me to come over for lunch. Now, suppose further that I was mistaken about the meeting: had I checked my e-mail, I'd have known that it had been rescheduled for the following week. Hence, the statement I made to my neighbor is, in fact, false—and since my intention in making it was to deceive, I have lied according to the dictionary definition. But there is a strong intuition (among many) that since I believe that what I was saying was true, it can't count as a lie.<sup>251</sup> A second counterexample of this type involves eavesdroppers: suppose I'm a criminal, and I know the cops have tapped my phone; I may try to throw them off by calling up a co-conspirator (he's in on the deception) and saying, "I'll meet you at the hideout at noon." Neither of us has any intention of being anywhere near the hideout at noon (in fact, we're planning to lam it); I said what I did to get the police to have false beliefs. But since they weren't the (ostensible) audience

---

<sup>251</sup> Fallis 2009 quotes George Costanza (a character from the television show *Seinfeld*, played by Jason Alexander): "It's not a lie if you believe it." This intuition is not universally shared, however; there are those who categorize any deceitful speech act, even those believed-true or outright true, as a lie. More on them anon.

of my utterance, many have the intuition that I didn't lie to them; and since I certainly didn't lie to my co-conspirator, there's no lie in this case.<sup>252</sup>

One can challenge this “most common” definition of lying in a number of ways, some more radical than others. One quite radical objection, which we may as well start with, if only because it's the hoariest, comes from philosophers whose moral convictions dictate that lying be wrong by definition—that it be an analytic truth that it's always, indefeasibly morally wrong to lie. Accounts of this kind are often motivated by an adherence to background religious traditions which feature exceptionless proscriptions on lying, and so among the philosophical elaborations of such positions are some of quite antique vintage: Augustine, Aquinas, Grotius, and Kant are prominent historical figures in this tradition. The principal challenge to the advocate of an exceptionless prohibition of lying is the commonplace intuition that in some cases, it is morally permissible to lie—to spare someone's feelings or to prevent some evil act. Kant (in)famously bites the bullet on this point, maintaining that it is impermissible to lie even to a murderer at one's door, inquiring as to the whereabouts of one's friend, whom he intends to murder.<sup>253</sup> To avoid these sorts of unintuitive consequences of their view, theorists of this sort must carefully define lying so that the apparent counterexamples to the universal proscription turn out not to be lies at all. Thus, Grotius' definition (as glossed by Mahon):

To lie =<sub>df</sub> to make a believed-false statement to another person, with the intention that that other person believe that statement to be true, violating that person's right of liberty of judgment, with the intention to harm that other person.<sup>254</sup>

---

<sup>252</sup> Chisholm and Feehan 1977 present an example like this.

<sup>253</sup> Kant, 1797. Though see, e.g., Mahon 2009 for a discussion of the nuances of Kant's views on lying. As is always the case with Kant, it's not clear that the received interpretation of his views is the correct one; there are many other texts, with many fine distinctions, to consider.

<sup>254</sup> Mahon 2008(b), §1.6

This definition builds in the wrongness of lying, with the violation of the victim's rights and the intention to harm her, while leaving an out: certain listeners are not entitled to the truth, so that deceiving them with falsehoods does not violate their "right of liberty of judgment," and hence is not lying. The murderer's right of liberty of judgment is over-ridden by "the opposition of another right which, in the common judgement of all men, is much more cogent."<sup>255</sup>

I'm not in favor of defining lying in this way. The first, to me most persuasive, reason not to is that it conflicts with standard, everyday usage of the term. Every child is familiar with the idea of a "white lie"—a kind of deception-via-falsehood that is morally permissible for whatever reason (it spares someone's feelings, it's harmless, etc.), but is nevertheless a *lie*. The difference between me and the advocate of an exceptionless proscription of lying is merely terminological, in this sense: we both agree that giving the murderer at the door false information is morally permissible (with the possible exception of Kant, but see Mahon 2009); we just disagree about whether to call this act a 'lie'. And since my terminological choice is in keeping with standard usage, it is to be preferred. A second reason to avoid the hard-line on lying is that the standard move to avoid counterexamples—in those cases, the listener isn't entitled to the truth anyway, so it's not a lie—is not always available. Consider a typical white lie: suppose I've severely injured myself by carelessly operating a piece household maintenance equipment; when I get home from the hospital, my neighbor has thoughtfully prepared food for me and my family. The dish in question, though, is gumbo, and, unbeknownst to my neighbor, we're all allergic to shellfish; so I give the entire pot to my brother-in-law. When my neighbor asks me if we enjoyed the gumbo, I lie and tell her that we did, that it was delicious. She obviously went to a lot of trouble; what's the point of telling her the truth? My lie seems morally permissible. Disanalogy: my neighbor is

---

<sup>255</sup> Grotius 2005, p. 617

not a murderer at the door. She doesn't have evil intentions that deprive her of her "right of liberty of judgment"; she's just as entitled to the truth as anyone else. So, unless we want to bite the bullet and claim that I did wrong on this occasion, we would need an alternative explanation, possibly involving a different definition, for why this is not a genuine lie. Perhaps we could provide such, but why bother? This conception of lying is already inconsistent with conventional usage; adding further complications can only make things worse.

Less radical criticisms of the "most common" definition of lying come from the consideration of a variety of supposed counterexamples, with various authors reporting conflicting intuitions and adding or subtracting necessary conditions from the definition accordingly. So, for example, I'm not sure I share Siegler's intuition about the case of Pablo from *The Wall*. Part of me wants to say he *thought* he was lying, but he was mistaken; he was actually telling the truth. (Or: he *tried* to lie to his captors, but, tragically and ironically, failed.) I'm inclined to embrace Thomas Carson's view: "Showing that a statement is true is always sufficient to counter the accusation that one has told a lie."<sup>256</sup> Imagining defenses against lying-accusations is of course in line with my own methodology, and it helps crystallize my intuition in the Pedro case: it would be odd indeed for the authorities to return from killing Gris in the cemetery and accuse Pedro of lying. Given these intuitions, I would add a necessary condition to the most common definition to the effect that the statement must be false (retaining the condition that it must also be believed-false, to handle the case of misfired misleading discussed above; I find it hard to argue with the Costanza Doctrine: it's not a lie if you believe it<sup>257</sup>).

---

<sup>256</sup> Carson 2006, p. 284

<sup>257</sup> Although, as Kieran Setiya (p.c.) points out, there may be complications here involving self-deception: if you believe *p*, but have deceived yourself that you really believe its opposite, is it possible to lie by asserting *p*? I don't have a good answer to that, but since I don't think much of importance hangs on the Costanza Doctrine, I'll just set that question aside.

I don't think very much of importance hangs on my preference for a falsity condition, though. I don't begrudge others their conflicting intuitions, and I suspect such disagreements are unlikely to be settled by argument. Lots of the moves and counter-moves in the debate over the definition of lying strike me as concerning similarly minor issues, concerning non-standard scenarios about which intuitions might be expected to vary from person to person. So, for example, Roderick Chisholm and Thomas Feehan take it to be a virtue of their account that it accommodates the intuition (shared by Augustine, from whom they borrow the example) that the following scenario does not involve lying: Harry is about to take a trip down a road on which, Sally believes, there are bandits; Sally is concerned about Harry's safety, but since the two have just had a quarrel, Sally knows that Harry will distrust anything she says; so Sally tells Harry that there are no bandits on the road. This is not a lie, according to Chisholm and Feehan, because, while Sally disbelieves her statement, and is in some sense being deceptive (outwardly misrepresenting her actual beliefs), she has not (according to their definition of 'assertion') asserted that there are no bandits, since a necessary condition on genuine assertions is that the speaker state a proposition under conditions which she believes justify the hearer in believing that she accepts the proposition. The conditions of Sally's statement are not like this: she knows that Harry distrusts her, so she does not think that Harry will believe that she accepts the proposition that there are no bandits; in fact, she's counting on uptake to the contrary. Since asserting is necessary for lying, and there is no assertion in this case, Chisholm and Feehan conclude that there is no lie.<sup>258</sup>

If all that seemed confusing, that's because it is. Chisholm and Feehan themselves have trouble keeping track of the relevant features of the situation. At one point they say: "[Sally] has

---

<sup>258</sup> Chisholm and Feehan 1977, p. 154



acted with the intention of deceiving [Harry]; for, while believing that there *are* bandits on the road, [s]he has acted with the intention of causing [Harry] to believe that [s]he, [Sally], believes that there are no bandits on the road.”<sup>259</sup> But that’s not right: Sally knows that Harry doesn’t trust her, so she can’t intend to cause him to believe that she thinks there are no bandits by saying *that very thing*. Sally is being deceptive here, but not in the way Chisholm and Feehan describe. Probably, they describe it that way—in terms of Sally’s intentions with respect to affecting Harry’s beliefs—because that description fits neatly with their rather convoluted definition of assertion. One counts as genuinely asserting what one states just in case the conditions are such that he believes that they justify the hearer believing that (i) the speaker accepts the proposition he states, and (ii) the speaker intends to “contribute causally” to the hearer’s coming to believe that he (the speaker) accepts the proposition.<sup>260</sup> So Sally is deceptive, on the incorrect description of the scenario, because she believes one thing and acts with intention of causing Harry to believe that she believes its negation (though, again, she hasn’t lied because she hasn’t asserted: she doesn’t believe the conditions are such as to justify appropriate beliefs by Harry).

I dwell on this example not to belabor Chisholm and Feehan, but because it highlights some key features of the literature on the definition of lying. First, many of the scenarios considered are so atypical that I’m not sure what my intuitions are. Did Sally lie to Harry above? I don’t know what to say. She said something she didn’t believe, and she was being deceptive, but then again.... Is it possible, in general, for someone who lacks the trust of one’s audience to lie to them? Those who insist on countenancing so-called bald-faced lies—where both speaker and hearer know that the claim is false—would say yes. But this situation differs from those, in that the falsity is not out in the open. In fact, there is no falsity in the actual utterance: the

---

<sup>259</sup> *Ibid.*

<sup>260</sup> *Ibid.*, p. 152

intention of the speaker is to cause the listener to believe the very claim made. Can there be a lie in such a case? I suspect opinions are likely to vary on these questions, and on the status of many of the other sorts of examples and counterexamples discussed in this literature. All of which makes one wonder what the point is of focusing on such non-standard scenarios in the first place, and doubt whether it's possible to articulate a set of necessary and sufficient conditions capable of capturing the concept of lying. Might we not be better off considering more prototypical cases and asking what they have in common?

This is the view of Linda Coleman and Paul Kay, who argue that 'lie' should be given a "prototype semantics" according to which the "applicability of a word to a thing is in general NOT a matter of 'yes or no', but rather of 'more or less'."<sup>261</sup> Rather than having as its meaning a set of necessary and sufficient conditions, "[t]he meaning of the word *lie* ('prevaricate') consists in a cognitive prototype to which various real or imagined events may correspond in varying degrees."<sup>262</sup> There are certain features that prototypical cases of lying share—according to Coleman and Kay, a prototypical lie is an assertion of a proposition (i) that is false and (ii) believed to be false, with (iii) the intention to deceive the listener—and to the extent that various scenarios lack one or more of these, they will be *less* of a lie. Coleman and Kay present the results of an experiment to back up their claims: they presented subjects with eight fictional scenarios, from the prototypical lie in which all three conditions are present, to the prototypical instance of truth-telling, in which none of them are present—and then every possibility in between, in which some of the conditions are present and others are not; subjects were asked to rate the scenarios on a lie/non-lie scale, and the results were consistent with Coleman and Kay's expectations—for example, that scenarios in which one doesn't expect to deceive with one's

---

<sup>261</sup> Coleman and Kay 1981, p. 27

<sup>262</sup> *Ibid.*, p. 26

falsehood, or those in which one seeks to deceive with a truth, score somewhere on the scale between the prototypical lie and truthfulness.

The very idea of prototype semantics, of course, raises a whole host of difficult questions in philosophy of language and mind (So Semantics must be use-based? What sorts of cognitive structures correspond to prototypes? What about compositionality? Etc.), and is controversial, to say the least. This is not the place to debate the merits and demerits of the overall approach to word-meaning. But if we set these larger issues to one side, we can perhaps draw some lessons from Coleman and Kay's results—and suggest ways in which they might be expanded and improved upon.

The first thing to note is that the results are consistent with some of the reservations about particular cases that we have registered already. They present a story with the same features as the Pablo example from *The Wall*, in which the speaker intends to deceive by asserting a proposition she believes to be false, but which turns out to be true. Subjects have mixed feelings about the example, scoring it a 5.17 on a 7-point scale (where 7 is an outright lie), consistent with my own reservations, and suggesting perhaps that the case for abandoning the falsity condition is not as clear as some authors have suggested. They do not, however, present a story with the same features as the Harry and Sally example, in which the speaker utters a falsehood that she disbelieves, *and* intends to deceive (so that all three of Coleman and Kay's conditions are apparently met), but intends to deceive in a non-standard way—i.e., not about the content of the proposition uttered, but about her beliefs. Perhaps Coleman and Kay would agree with Chisholm and Feehan that this is not even a genuine case of assertion, so that it doesn't even merit consideration as a potential lie (Coleman and Kay are explicit that lies are a type of assertion), but since they don't give us an account of assertion, it's hard to tell. We'll return to

this issue (the importance of an account of assertion) below. Anyway, I'd be curious to see where subjects would put that case on the 7-point scale.

Another thing to note about the results is that the case closest to the concerns of the present work—in which the speaker tells the truth with the intention to deceive—is also one about which subjects report mixed feelings. It rates 3.48 on the non-lie/lie scale—slightly closer to telling the truth than lying. It's a classic case of false (particularized) implicature:

John and Mary have recently started going together. Valentino is Mary's ex-boyfriend. One evening John asks Mary, 'Have you seen Valentino this week?' Mary answers, 'Valentino's been sick with mononucleosis for the past two weeks.' Valentino has in fact been sick with mononucleosis for the past two weeks, but it is also the case that Mary had a date with Valentino the night before. Did Mary lie?<sup>263</sup>

I think the subjects' overall ambivalence about this example is telling—both about our concept of lying and about Coleman and Kay's methodology. After each story in the survey, subjects were presented with the following choices: "It was {a lie / not a lie / I can't say}. I am {very sure / fairly sure / not too sure} most others would agree with the choice I just circled."<sup>264</sup> If the subject circles "I can't say", then the story gets scored a 4 regardless of how sure she is that others would agree; if "not a lie" is circled as the first choice, then the story gets a 1, 2, or 3, depending on certainty (1 for most certain, 3 for least); and if "a lie" is circled, the scores are 5, 6, or 7 depending on certainty (with 7 corresponding to the highest degree of certainty). Again, this false-implicature story got an overall score of 3.48—on the non-lie side, but close to maximum ambivalence. But there are a number of ways a story could get such an overall score;; lots of people could respond with "I can't say", along with a few rather tepid endorsements of non-lie, providing evidence of widespread ambivalence; or, the subjects could be almost evenly split between those who think it's a lie and are confident others will agree, and those who think it

---

<sup>263</sup> *Ibid.*, p. 31

<sup>264</sup> *Ibid.*, p. 30

isn't and are also confident in their judgment. Alas, Coleman and Kay don't provide a complete breakdown on the results, but they do tell us that, among the 67 total subjects, 7 responded "I can't say", with 42 calling it a non-lie and only 18 calling it a lie. So there's pretty overwhelming support for calling this a non-lie; though from the fact that the score ends up so close to 4 we can infer that there's not a whole lot of confidence in that judgment. I'd be curious to see if this is a case in which people's intuitions could be "sensitized", as Bach once put it.<sup>265</sup> That is, I wonder what the results would be if subjects were given three choices: lie, truthful claim, and misleading claim—after being exposed to examples of the last. Since it is such an everyday notion, my guess is that after it was made salient, most people would agree that Mary misled rather than lied.

This suggests a way that Coleman and Kay's results might be made more robust—by discarding the overly simple bipolar distinction between lie and non-lie, or by making a particular kind of non-lie (misleading with the truth) salient. With this kind of foundation, the natural way to extend this kind of study is to focus on the sorts of cases encountered in Chapter 2, in which we encountered what appeared to be a range of intuitions about the particular kind of case—misleading with the truth—represented in Coleman and Kay's John and Mary story. It might be revealing to probe subjects' intuitions about various scenarios involving false GCIs, implicatures, explicatures, etc. This gets to what I consider a more fundamental issue raised by consideration of the differences between lies and non-lies: what portion of a speaker's intended communication counts as genuinely asserted, as opposed to merely implicated or implied? Coleman and Kay, and most other participants in the debate over how to define lying, bracket this question (if they consider it at all), focusing instead, for example, on whether the intent to deceive is a necessary condition for lying, or whether one's statement must be false. My focus is

---

<sup>265</sup> Bach 2002, p. 2

more radical: what's the content that we're evaluating for truth or falsity? How is it separated from other contents conveyed in a speech act? It's no surprise that subjects exhibit ambivalence about a false implicature, since these sorts of difficult questions are in the background and unaddressed. I suggest bringing them to the fore.

## 4.2 ASSERTION

Such a move means addressing head on an issue that has only lately been given much attention in the literature on defining 'lie': what kind of a speech act is it? It is often averred that lying is a kind of assertion.<sup>266</sup> Rarely, however, are these averments accompanied by any kind of detailed account of this speech act.<sup>267</sup> Often, definitions of lying seem to ignore distinctions among various speech acts.<sup>268</sup> Lately there has been a shift toward taking more seriously the question of just what constitutes an assertion. This is, in my view, a salutary development, since, as I've noted, I think the deepest issue one confronts when trying to circumscribe the concept of lying is the problem of specifying precisely what (false, believed-false, etc.) content the liar is being held

---

<sup>266</sup> E.g., by Mannison 1969, Chisholm and Feehan 1977, Simpson 1992, Fallis 2009, Stokke 2013(a).

<sup>267</sup> Mannison 1969 and Simpson 1992 provide none; Chisholm and Feehan 1977 and Fallis 2009 provide sketchy *ad hoc* glosses (pp. 149 – 152 and p.35, respectively; to be fair: Chisholm and Feehan are much less sketchy (though, as noted above, their account is quite unwieldy, and Fallis is much more self-conscious about his sketchiness). Stokke is the exception; he recognizes the importance of the point I'm pressing—that a definition of lying that characterizes it as a type of assertion ought to be accompanied by a serious account of assertion.

<sup>268</sup> E.g., Carson 2006 and Mahon 2008(a) both defend definitions of lying in which it's characterized as "making a statement" that meets certain conditions (intended to deceive, false, believed-false, etc.). Mahon's survey of the literature on defining 'lie' (Mahon 2008(b)) also avoids any debates about speech-act types, opting for the neutral but unexplained "making a statement" to explicate the various competing definitions—even the definitions of those who are more careful about making distinctions among speech acts.

responsible for; that is, if lying is to be a kind of assertion, we must say what it takes for a proposition to count as being asserted. And this is a difficult task, as the results of Chapter 2 demonstrate: speakers communicate much with their speech acts, and it's hard to say which part of this they are fully responsible for the truthfulness of.

The shift in the literature toward taking assertion more seriously has been occasioned by the consideration of a particular kind of purported counterexample to the most common definition of lying—so-called “bald-faced” lies.<sup>269</sup> One tells a bald-faced lie when one states a falsehood, the falsity of which is common knowledge to both speaker and hearer. One of Carson's examples (modified slightly to add color) is illustrative: consider a (reluctant) eye-witness at the murder trial of a notorious *mafioso*; he witnessed this gangster killing someone, but he knows that if he testifies to that effect, there will be swift and terrible retribution from the Family. So he gets on the stand and tells the jury that he didn't see anything; he wasn't even at the scene of the crime—he was in, uh, Buffalo at the time. Yeah, Buffalo. Suppose further that the jury had just viewed video evidence (surveillance footage from the parking garage in which the assassination occurred) clearly showing the witness at the scene. So everybody knows that what the witness is saying is false—he could even be brought up on perjury charges (bad, but better than “sleeping with the fishes”)—and the witness knows that everybody knows his claim is false; he doesn't think he's fooling anyone. The intuition is that it's nevertheless appropriate to say that he's lying. Thus, contrary to the most common definition, the liar needn't intend to deceive his audience. Various alterations to the definition are proposed in the literature. Carson, for example, proposes that one lies when one (insincerely) warrants the truth of a falsehood. Roy Sorenson objects that this definition fails to account for bald-faced lies, because

---

<sup>269</sup> See especially Sorenson 2007, Carson 2006, Fallis 2009, and Stokke 2013(a).

those take place only in contexts in which it's impossible to warrant the truth of the false statement: "One can warrant p only if p might be the case. When the falsehood of p is common knowledge, no party to the common knowledge can warrant p because p is epistemically impossible."<sup>270</sup> Sorenson counters with a pleasingly spare definition of lying, according to which it amounts simply to asserting what one believes to be false, but, as Carson points out, "the plausibility of this definition depends on his account of asserting a proposition."<sup>271</sup> And Sorenson doesn't provide much of an account. He says only that an assertion must have "narrow plausibility," so that "someone who only had access to the assertion might believe it," as opposed to "wide plausibility," i.e., "credibility relative to one's total evidence."<sup>272</sup> Noting that this isn't a whole lot to go on, Carson takes up the challenge to more fully explicate his notion of warranting the truth of a statement, saying that it involves "promising or guaranteeing that it is true" and "invit[ing] others to trust it or rely on it."<sup>273</sup> Don Fallis adopts Sorenson's definition (lying = asserting what is believed-false), and provides a gloss of assertion in Gricean terms: "[Y]ou assert something when you say something and you believe that Paul Grice's first maxim of *quality* (viz., "Do not say what you believe to be false") is in effect as a *norm of conversation*."<sup>274</sup>

These are all moves in the right direction, toward taking more seriously the task of specifying just what kind of speech act lying is. It's usually said to be a kind of assertion, and some attempts are being made to spell out what that means. These attempts, though, tend not to make contact with the large and mature literature on the very question of what constitutes an

---

<sup>270</sup> Sorenson 2007, p. 254

<sup>271</sup> Carson 2009, p. 20

<sup>272</sup> Sorenson 2007, p. 255

<sup>273</sup> Carson, *op. cit.*

<sup>274</sup> Fallis 2009, p. 6 [My references are based on pagination of a copy of the paper available on Fallis' website: <http://sirls.arizona.edu/fallis>]



assertion. As a result, they fail to deal adequately with the very difficult questions raised by our investigations in Chapter 2: Under what conditions are the things we *say* genuinely asserted? What sense of ‘saying’ is operative here? That term can be used more or less minimally. It is a commonplace that we communicate more than we say; is any of this additional content part of what we assert, and if so, what criteria do we use to determine that it is? And so on.

The need to attend to such issues is nicely illustrated by a critique of Fallis’ definition by Andreas Stokke. In the course of his criticism, Stokke makes many of the same points about *saying* that were established in Chapter 2. Fallis freely admits that, on the face of it, ironic utterances appear to be counterexamples to his definition of lying—cases in which all three conditions (saying something, believing it false, and believing Grice’s first maxim of quality is in effect) are met, but in which there is no lie. He considers a sarcastic line delivered by the character Han Solo in the movie *Star Wars*: after he and some others, on another character’s suggestion, have escaped danger by sliding down a garbage chute, Solo remarks, “The garbage chute was a really wonderful idea. What an incredible smell you’ve discovered!” Solo thinks it’s false that the garbage chute was a wonderful idea—indeed, that is what he intends to convey with his utterance—and there’s no apparent reason to think Grice’s maxims aren’t in effect. And yet he’s obviously not lying. Fallis notes that this case can be analyzed as a conversational implicature, in which the speaker succeeds in conveying something other than what he has literally said by openly flouting one of the maxims. Fallis says that in this case Solo is not lying because the relevant Gricean maxim, having been flouted, is not in effect: “I contend that, by flouting this norm of conversation, Solo turns it off.”<sup>275</sup> As Stokke points out, this analysis reflects a fundamental misunderstanding of the way in which implicatures are supposed to arise.

---

<sup>275</sup> *Ibid.*, p. 31

Flouting maxims doesn't "turn them off." On the contrary, that the maxims are in effect is one of the premises a hearer must rely on to make the appropriate inferences in cases of implicature; she infers *from* the presumption that the speaker is being cooperative—from the presumption that the maxims are in effect—to the conclusion that he intends to convey something other than what he says. Stokke does a nice job of explicating the inference in the case of Solo's utterance (note that "FMQ" abbreviates 'First Maxim of Quality'):

- (a) Solo said that the garbage chute was a wonderful idea.
- (b) Solo believes that the garbage chute was a bad idea.
- (c) Solo is observing FMQ.
- (d) Unless Solo intended to convey that the garbage chute was a bad idea, he would not be observing FMQ.
- (e) Therefore, Solo intended to convey that the garbage chute was a bad idea.<sup>276</sup>

So Fallis cannot say that, in cases of irony, FMQ is not in effect; according to his definition, then, such remarks count as lies. The only other option available to rescue the definition is to deny that ironic utterances count as genuine *sayings*. There is Gricean precedent for this: as we noted in Chapter 2, for Grice, what we mean is factorable, without remainder, into what we say and what we implicate; saying something thus entails meaning it. We don't mean the literal contents of ironic utterances, so we don't *say* them; rather, we "make as if to say" them. Putting the point in terms we're familiar with: the sense of 'saying' according to which the literal content of ironic utterance counts as having been said is a more minimal, merely locutionary sense; Grice's conception of saying is a more robust, illocutionary one. As Stokke points out, the adoption of this Gricean sense of 'saying', while it allows Fallis to deny that Solo's utterance is a lie, is still not an option, since it opens the door to another counterexample in the form of bald-

---

<sup>276</sup> Stokke 2013(a), p. 4 [My references are based on pagination of a copy of the paper available on Stokke's website: <http://andreasstokke.net/>]

faced lies, which Fallis also wants his definition to capture.<sup>277</sup> On Grice's view, we mean what we say, and meaning is analyzed in terms of reflexive intentions: a speaker means something iff he intends the audience to *believe* it, by way of recognizing that very intention. The intention to cause belief is the key: bald-faced liars have no such intention, and so in a Gricean framework they can't count as meaning, and hence saying, the falsity they utter. If Fallis adopts this more expansive conception of saying, bald-faced liars don't count as having lied. This is Stokke's observation; I would add another: the proposal being considered is that Fallis adopt the strong, illocutionary Gricean sense of 'saying' as part of his definition of assertion. But to do so, great care must be taken to avoid circularity. Self-consciously following the Gricean tradition, Bach and Harnish 1979 define assertion in terms of reflexive intentions much the same as Grice defines his sense of saying; indeed, they list 'say' as a synonym for 'assert'.<sup>278</sup> So if Fallis were to adopt this notion of 'saying', as Stokke suggests is a possibility, then it looks as though he'd just be defining assertion in terms of itself. To avoid that, he'd need a way of distinguishing illocutionary saying from asserting; not an impossible task, perhaps (it's a slightly weaker act in some sense?), but not one Fallis seems equipped to take on.

Stokke extracts, I think, the proper lesson from the failure of Fallis' *ad hoc* definition of assertion: if we need to define asserting before we can define lying, then why not look to one of the well-developed accounts of assertion to be found in the vast literature on that speech act? Stokke opts for an account inspired by that of Robert Stalnaker: to assert something is to (i) say it, and (ii) propose to add it to the conversational common ground.<sup>279</sup> This common ground is a body of information available to conversational participants, a background against which the

---

<sup>277</sup> *Ibid.*, p. 7

<sup>278</sup> Bach and Harnish 1979, p. 42

<sup>279</sup> The theory first appears in Stalnaker 1978, with further refinement and development in various subsequent writings. Stalnaker himself says very little about 'saying'; Stokke stresses the importance of this condition.

conversation develops—where development involves updating this body of information, typically with the contents of assertions, but also, for example, with contents which much be presupposed for an utterance to be felicitous (if I assert that the aliens who designed the pyramids had almond-shaped eyes, the common ground may have to be updated to include the presuppositions that there were such aliens, and that they designed the pyramids).

Stokke notes that we must be careful about how we define the common ground, lest the definition of lying it's meant to underwrite succumb to counterexamples. It seems natural to say that the common ground is a shared set of beliefs, and so an assertion is a proposal to add to this set of beliefs. But bald-faced liars do not propose to add a belief to the common ground; they and their audience *disbelieve* the content of the lie, and everybody knows this. So, following Stalnaker (who had distinct reasons for making this move), Stokke defines the common ground in terms of a non-factive attitude weaker than belief: *acceptance*. Interlocutors can accept propositions that they're unsure of, or that they know or believe to be false (as is sometimes required to facilitate conversational smoothness). The common ground, then, is a body of mutually accepted information, available to all conversational participants.

Stokke's definition of lying in terms of the common ground conception of assertion is this:

S lies to X if and only if

(L1) S says that *p* to X, and

(L2) S proposes that *p* become common ground, and

(L3) S believes that *p* is false.<sup>280</sup>

---

<sup>280</sup> Stokke 2013(a), p. 14

He argues that this definition is able to handle all the standard cases in the literature. An ironic utterance, for example, will be one for which the first and third conditions are met—one says something one believes to be false—but for which the second condition is not: we do not propose to update the common ground with the believed-false content (Solo does not propose that he, Luke, Leia, and Chewbacca all accept that the garbage chute was a good idea). The definition correctly rules out cases of false implicature (acts of misleading, in our terms), Stokke claims, since in such cases the falsity does not count as having been *said*—it is merely implicated. And the definition is able correctly to handle bald-faced lies, as they meet all of the conditions, thanks to the characterization of the common ground in terms of acceptance rather than belief.

So Stokke's definition appears to do what it needs to. There are, however, complications, which Stokke notes and addresses, with varying degrees of success. First, Stokke discovers the same difficulty that occupied us for most of Chapter 2: it's hard to pin down precisely the notion of 'saying' we need to capture intuitions about lying. It must not be so robust that we end up counting genuine cases of mere misleading as lies (that is, it must not include certain clear cases of false implicature); and yet it cannot be so minimal as to only count truth-conditional content as having been said, for in that case we miss capturing some intuitive lies. Here Stokke is dipping his toes into murky waters, the imponderable depths of which we explored in Chapter 2. Stokke's first example of non-truth-conditional content that should be counted as part of what is said involves conventional implicature: we imagine an utterance of 'Ames, who stole from the FBI, is now behind bars', which conventionally implicates that Ames stole from the FBI; if we suppose that Ames is in fact behind bars, but that he did not steal from the FBI, Stokke reports that his intuition is that we should nevertheless count this utterance as a lie.<sup>281</sup> I agree with the

---

<sup>281</sup> *Ibid.*, p. 16

intuition, but disagree with the conclusion he draws from the case: “I take this to be in favor of a line... according to which conventional implicatures are part of what is said.... Since conventional implicatures are clearly proposed for common ground inclusion, the Common Ground Definition would therefore count cases [in which they are false] as lies, in accordance with this intuition.”<sup>282</sup> It’s a hasty generalization to conclude, on the basis of this one example, that conventional implicatures are part of what is said—and that if they’re false, utterances giving rise to them ought to count as lies. We briefly discussed conventional implicatures above (§2.5.3), asking whether they ought to be part of what is said. The results were mixed, because conventional implicatures are a motley bunch: there seemed to be a stronger case for including the implicata arising from ‘therefore’ as part of what is said; the case for those arising from ‘but’ was much weaker. I’m sure that a detailed examination of a wider variety of conventional implicatures would reveal the same sort of variability of intuitions that by now we’ve become inured to; cases like Stokke’s will seem clearly to be part of what is said, others clearly not, with many more falling in the grey area in between. And of course, consideration of conventional implicatures as candidates for inclusion in ‘what is said’ is only the beginning. As we saw, GCIs, implicatures, explicatures, and the like provide countless examples of subtle phenomena, the correct classification of which with respect to ‘saying’ resists tidy solutions. This is a concern about Stokke’s overall approach: while he’s given a more precise, theoretically robust specification of assertion (in terms of the common ground), there still remains significant uncertainty about what contents count as asserted in particular cases, for the fuzziness of the concept of *saying* so thoroughly documented in Chapter 2 remains. Now perhaps there’s not much one can do about this; there are all sorts of unavoidably vague distinctions, and that

---

<sup>282</sup> *Ibid.*

between what is said and what is otherwise communicated may be one of them. But it would be an extremely attractive feature of a definition of assertion (and a definition of lying in terms of it), if it could at least provide *some* guidance for making judgments in difficult borderline cases—or if not guidance, exactly, at least more basic definitional concepts about which we’re likely to have robust intuitions, intuitions that dovetail with those about judgments we make further downstream, such as whether or not a certain speech act counts as a lie. If Chapter 2 proved nothing else, it’s that our intuitions about *saying* are a bit of a mess in this respect. And the other more basic concept in Stokke’s definition—the common ground—doesn’t seem likely to be much help either: it’s a set of mutually accepted (not believed) propositions; our speech acts are proposals to update it. The notion of *acceptance* is not one about which anybody is likely to have robust pre-theoretical intuitions; indeed, it’s odd on the face of it to say that someone accepts a proposition that is false, and yet we need this possibility to be part of the theory. ‘Updating the common ground’ has some intuitive appeal: I make a claim, and we add it to the stock of things we accept. But questions remain. Adding a proposition to the set via assertion is a straightforward enough idea; presupposition accommodation also makes good sense. Is *subtracting* a proposition also straightforward?<sup>283</sup> And what about implicatures? Stokke thinks that it’s clear that conventional implicatures are to be added to the common ground. Are conversational implicatures also clearly to be so added? If implicatures are all proposed additions to the common ground, then the entire weight of drawing the distinction between lies and false

---

<sup>283</sup> MacFarlane 2011 gives reasons for thinking it’s not. The formal apparatus for representing propositions in the common ground involves sets of possible worlds, so that additions to the common ground actually have the effect of *reducing* the number of possible worlds in the set (one takes the intersection of the common ground and the set corresponding to the added proposition); it’s hard for technical reasons to picture how retraction works on this model, since it can’t be a simple union. And even if we set these technical concerns aside, retraction only seems like a straightforward matter if it occurs immediately: “Once an assertion has been accepted and its content integrated into the common ground, and a few more assertions have been made and accepted, it’s no longer obvious how one could ‘undo’ the assertion.” (p. 16) [My references are based on pagination of a copy of the paper available on MacFarlane’s website: <http://johnmacfarlane.net/>]

implicatures falls on the first condition of Stokke's definition of lying—on the concept of *saying*. If implicatures are not proposed updates to the common ground, why not? If they are, is there any difference between how GCIs and PCIs are added? It's easy to see why presuppositions need to be accommodated: those propositions are very closely connected to the ones that presuppose them. Arguably, conventional implicatures—since they arise from the meanings of terms—are also so closely related that they must also, obviously, be included in updates to the common ground. But conversational implicatures—the computation of which is mediated by the maxims—are relatively more distant from the propositions that give rise to them; and arguably PCIs more so than GCIs (given the latter's default status). So are all of these proposed for inclusion in the common ground? I'm not suggesting that this question has no answer within the Stokke/Stalnaker framework, only that the fact that the answer isn't *obvious* is evidence that the idea of updating the common ground is not as intuitive as it first appeared. I'll return to this concern in a moment, after first considering another potential complication for his definition that Stokke addresses.

This complication involves, once again, the case of bald-faced lies. As noted, the notion of common ground must be spelled out in terms of mere acceptance, rather than belief, so that a proposal to add to doesn't involve proposing one's audience believe what one says; that's not what the bald-faced liar does. This complication has been pre-emptively avoided. The further complication is this: if it's known to all—speaker and audience—that what the bald-faced liar is saying is false, shouldn't we say that this common knowledge is part of the common ground? And if so, isn't there something extremely odd about proposing to update the common ground with a proposition when its negation is already accepted? Does this amount to proposing that everyone accept a contradiction? Stokke says two things in response. First, it doesn't follow from



something's being common knowledge to its being part of the common ground; in fact, its negation might be common ground. We're asked to consider Keith Donnellan's old chestnut, 'The man drinking a martini is a philosopher'. It may be that both speaker and hearer of this utterance know that it's not a martini, yet, for the purposes of the conversation, they have both agreed to accept that he *is* drinking a martini.<sup>284</sup> So we can have a proposition be part of the common ground even though its negation is common knowledge; this is what happens with bald-faced lies. Furthermore, Stokke says, even if we insist that common knowledge must be part of the common ground, this does not present a problem in the case of bald-faced lies. Suppose it's common knowledge, and hence common ground, that the witness was at the scene of the crime. When he testifies to the contrary, the witness is proposing to update the common ground with the proposition that he was not there. This needn't involve a proposal that everyone accept both  $p$  and  $\sim p$ ; rather, the witness is suggesting that the update involve removing one proposition and replacing it with its negation.

I don't have any serious objection to any of this; I think Stokke handles the potential problems with aplomb. Rather, I want to express (again) a vague sense of unease. My unease stems from the feeling that a lot of this is quite unintuitive. I think the very fact that the objection sounded so plausible as initially stated (Isn't common knowledge common ground? How can you update a set of propositions with one of their negations?) is evidence that the ideas here—common ground, updating, acceptance (as opposed to belief)—are not ones that come naturally. Again, this is not a real objection; sometimes we need to resort to unintuitive, novel concepts in order to develop a comprehensive account even of an ordinary notion like assertion. But the notion, for example, that I can accept something while knowing it to be false is just odd. What

---

<sup>284</sup> *Ibid.*, p. 18

we accept appears to be shaped by practical (as opposed to theoretical) considerations, and to depend on contextual demands.<sup>285</sup> It's easy to see how acceptance can be the appropriate attitude in a case like the following:

The three of us need jointly to decide whether to build a house together. We agree to base our deliberations on the assumption that the total cost of the project will include the top of the estimated range offered by each of the sub-contractors. We facilitate our group deliberations and decisions by agreeing on a common framework of assumptions. We each accept these assumptions in this context, the context of our group's deliberations, even though it may well be that none of us believes these assumptions or accepts them in other, more individualistic contexts.<sup>286</sup>

That situation involves mere uncertainty about what is accepted; we don't *know* that the contractors won't all come in at the top end of their estimates, but it's useful to operate on that assumption. Appeal to such practical considerations is perhaps also plausible in the case of the non-martini-drinker: what he's actually drinking is some strange concoction we don't know the name of, so to facilitate a smoother conversation, we agree to accept that he's a martini-drinker, even though we know that's false. But the case of bald-faced liar in court seems different; sure, the witness goes on the record with his claim to have been in Buffalo, but in what sense does anyone *accept* that claim? What practical end is served by the alleged acceptance? Perhaps there's something Stokke could say about this example to put my mind at ease; I'll even stipulate that he can. My concern is how easily this kind of worry can arise in the first place. I'm just highlighting the unintuitive nature of some of Stokke's fundamental concepts. I'd prefer an account of assertion with a more intuitive foundation.

There are alternatives to the Stalnakerian Common Ground approach in the literature. MacFarlane 2011 provides a useful taxonomy of the views on offer. He divides them into four categories, and frames them as answers to the question "What is assertion?"

---

<sup>285</sup> See Bratman 1992.

<sup>286</sup> *Ibid.*, p. 7

1. To assert is to express an attitude.
2. To assert is to make a move defined by its constitutive rules.
3. To assert is to propose to add information to the conversational common ground.
4. To assert is to undertake a commitment.<sup>287</sup>

It remains for us to see whether any of these alternatives provides a satisfying, intuitive account of assertion—one that is, in particular, suitable for characterizing lies and differentiating them from acts of mere misleading.

The first type of account—assertion as an expression of an attitude—does not fare particularly well. Typically, the attitude in question is taken to be belief, so that to assert something is, roughly, to express one’s belief in it. This is obviously too simple; for one thing, we can assert things we don’t believe (lying!). A sophisticated definition of assertion along these lines appears in Bach and Harnish 1979:

In uttering *e*, *S* asserts that *P* if *S* expresses:  
     i. the belief that *P*, and  
     ii. the intention that *H* believe that *P*.<sup>288</sup>

As MacFarlane has also noted<sup>289</sup>, this definition seems to count implicata as part of what someone asserts.<sup>290</sup> For example, if my wife asks me whether I remembered to do the dishes, and I respond, “My fingers look like prunes,” I implicate that yes, I remembered to do the dishes, and have the dish-pan hands to prove it. But in fact I’m trying to get away with something: true, my fingers are prune-y, but only because I’ve spent the last few hours luxuriating in the hot tub, ignoring my chores—the dishes among them. This is a clear case of misleading with false implicature, which we want to distinguish from an outright lie. But in producing my utterance, I satisfy both of Bach and Harnish’s conditions, even if we let *P* = ‘I did the dishes’. The key to

---

<sup>287</sup> MacFarlane 2011, p. 2

<sup>288</sup> Bach and Harnish 1979, p. 42

<sup>289</sup> MacFarlane 2011, p. 4

<sup>290</sup> Bach apparently acknowledges and accepts this consequence: *ibid.* p. 5 (n. 5)

seeing this is to understand what they mean by ‘express’. It is not necessary for *S* to have the belief that *P* in order to express that belief; if it were, it would be impossible to assert things we don’t believe. Rather, to express an attitude is “for *S* to R-intend the hearer to take *S*’s utterance as a reason to think *S* has that attitude.”<sup>291</sup> An “R-intention” is a reflexive intention—an intention to produce an effect in the hearer, the fulfillment of which “consists in [its] recognition.”<sup>292</sup> That is, an assertive act succeeds just in case the hearer recognizes that the speaker intends that she take his utterance to be a reason for thinking that he has the relevant belief. Clearly, then, both of the conditions are met for the proposition ‘I did the dishes’: (i) I’ve expressed the belief that *P*, in the sense that I’ve R-intended my wife to take my utterance as a reason for thinking I believe that *P* (I’m trying to get away with something, so naturally I want her to believe that I believe that I did the dishes); and (ii) I’ve also expressed an intention that she believe that *P*. Now, if we assert not only what our utterances literally mean, but also whatever they implicate—and if we define lying as a form of assertion, so that asserting something believed-false is a sufficient condition for having lied—then we lose the distinction between lies and false implicatures. This account of assertion won’t do for our purposes.

The second alternative—according to which assertion is an act governed by constitutive rules—comes in many varieties.<sup>293</sup> What they have in common is the claim that what is essential to an assertion is just that it is subject to evaluation by a particular rule or norm; the accounts differ with respect to which norm they identify as the constitutive rule. The most frequently defended norm is the so-called knowledge norm of assertion, according to which it’s constitutive of assertive practice that we follow a norm something like “one ought to assert only what one

---

<sup>291</sup> Bach and Harnish, *op. cit.*, p. 15

<sup>292</sup> *Ibid.*

<sup>293</sup> See, e.g., DeRose 1996 and 2002, Hawthorne 2004, Koethe 2009, Lackey 2007, Weiner 2005, Williamson 1996, Unger 1975.

knows.” Other possible norms replace the knowledge-condition with, e.g., truth, belief, reasonableness, etc.; some are skeptical of the whole enterprise of providing such norms.<sup>294</sup>

While talk of norms dovetails nicely with my own preoccupation with lying and misleading, I’m among those skeptical of this general approach. I lack fully developed arguments to support my skepticism, and even if I had them, it would probably take us too far off track to spell them out. I’ll limit myself to (i) the expression of a vague suspicion, and (ii) pointing to a feature that these accounts lack, and that my preferred account does not. First, the suspicion. Arguments for these constitutive-norm accounts of assertions rely on intuitions about particular cases—about the correctness or incorrectness of assertions in different circumstances. Exponents of these views are careful to distinguish the constitutive norm governing the practice of assertion from other norms we might use to judge individual assertions. Constitutive norms are supposed to be like rules of a game—rules that serve to define a kind of practice.<sup>295</sup> If one violates such a rule, one is subject to criticism on grounds that are internal to the practice. So, if one lies, one is sanctioned not only from the perspective of moral norms, which lie outside the practice of asserting, but from within: it’s a defective assertion, against the rules of the game, as it were. Here’s my problem: I know what it’s like to have an intuition about the moral permissibility of a particular assertion, or about its politeness, aptness, etc.; I *don’t* know what it’s like to have an intuition about the correctness of a particular assertion qua assertion, independent of any of the other usual evaluative standards. Peter Pagin apparently shares my concern, and puts it more clearly:

Even more rare, I think, are evaluations in such cases when it is also made clear that the assertion is not judged from a moral point of view, or from the point of view of general

---

<sup>294</sup> DeRose, Hawthorne, Williamson and Unger are in the knowledge camp; Lackey goes for a “reasonable to believe” norm; Weiner favors truth; Koethe is skeptical.

<sup>295</sup> Williamson puts it this way, using the analogy to games: *op. cit.*, p. 234

prudence or etiquette, but strictly from a point of view that only deals with semantics, speech act principles and epistemic conditions. In fact, it is not so clear that it happens at all outside philosophical contexts. Finally, when the question comes up, it seems to be up for debate what is to be considered as right or wrong, and this indicates that there is no firm empirical basis for such assessments. It is then not easy to see what the empirical basis could be for the claim that assertion is truly characterized by one norm rather than another.<sup>296</sup>

Indeed, when authors defending various norms over others cite their intuitions, they find it hard not to express them in moral terms. Analyzing a particular assertion in context, Matthew Weiner says about it, “This sounds more acceptable than the assertion of (1) in the original set-up. It may be banal, unkind, and pointless to assert (11), but it is not likely to mislead.”<sup>297</sup> The assertion in question is vindicated—despite various unattractive features—on the grounds that it’s not misleading; that doesn’t sound like a purely semantic assessment, as Pagin puts it. Again, I don’t take this to be a knock-down argument; only an expression of unease with this particular approach, one reason I feel more comfortable with an alternative. Another such reason is simply that the constitutive-norm accounts of assertion lack a certain directness; summing up Stalnaker’s approach, MacFarlane puts it like this:

Stalnaker’s account of assertion differs from the two accounts we have examined so far in focusing neither on what is expressed by an assertion nor on the norms for when an assertion may be made, but on what he calls the “essential effect” of an assertion. As an answer to the question “what is it to make an assertion,” this is attractively direct. It has the form: “to assert is to  $\Phi$ .”<sup>298</sup>

My preferred account of assertion (and MacFarlane’s) shares this directness.

It also shares the main attraction of the rules-based approach: it defines assertion in explicitly normative terms, dovetailing with the present focus on assertions (lies) we judge to be wrong. I prefer a definition of assertion according to which to assert is to undertake a certain

---

<sup>296</sup> Pagin MS, p. 11

<sup>297</sup> Weiner 2005, p. 16 [My references are based on pagination of a copy of the paper available on Weiner’s website: <http://mattweiner.net/>]

<sup>298</sup> MacFarlane 2011, p. 14

kind of commitment. In this, I'm following the general approach outlined by Brandom (especially in his 1983 and 1994). The basic idea is to characterize assertions in terms of their normative consequences: in making an assertion, one undertakes a commitment, imposes upon oneself certain obligations. Specifically, in asserting one "commits oneself to justifying the original claim."<sup>299</sup> This produces an obligation "to justify the claim if challenged."<sup>300</sup> So lying, within this framework, is, roughly, committing oneself to justify something one believes to be false (and hence, unjustifiable). Given the explicitly normative terms in which this account is framed, it's easy to see why lying is often wrong: one shirks one's responsibilities and hangs one's interlocutors out to dry, for the effect of asserting on them is to entitle them to assert the same thing (and other things that follow from it), deferring the justificatory responsibility to the original speaker; when one can't meet that obligation, one has (often) wronged one's audience, having put them into a position of vouching for unjustifiable claims.

One of the main problems we had with the other accounts of assertion was the difficulty of drawing the distinction between lying and misleading within their frameworks. For the Bach and Harnish view, false implicatures ended up counting as asserted contents, making it impossible to draw the distinction. For the Stalnaker/Stokke view, the difficulty was that the terms in which assertion was defined were insufficiently intuitive, so that they provided little guidance in deciding difficult cases. The notion of the common ground was not one about which we had any pre-theoretical intuitions, so the burden of distinguishing lies from non-lies fell on the notion of 'saying', about which there are too many unanswered questions. If what is said is to be specified according to some prior semantic, pragmatic, or mixed account, then, as we saw in Chapter 2, the difficulties may be insuperable. But here we find another potentially attractive

---

<sup>299</sup> Brandom 1983, p. 641

<sup>300</sup> *Ibid.*, p. 642

feature of the Brandomian framework: his approach reverses the explanatory order, specifying content not in terms of some prior semantic account, but in terms of the practices of participants in the *Sprachspiele* of assertion:

Whether or not one claim justifies another, for example, is not determined by some objective semantic content or relations the sentences have and which the community must try to live up to or reflect in their social practices of recognizing some claims as justifying others. Rather, a justification is whatever the community treats as one—whatever its members will let assertors get away with. It is from the communal responsive dispositions to recognize some claims as justifying others that the sentences involved first acquire their semantic contents, paradigmatically their behavior as antecedents of conditionals. The meanings do not determine the appropriate inferences, but what inferences are socially appropriate determines the meanings of the sentences involved in those inferences.<sup>301</sup>

I quote Brandom at great length here in order to include both what I like about his overall approach and what I'd rather not embrace. The latter category includes his “inferentialist” Semantics, according to which semantic content is determined by inferential role; this is the view gestured at in the last two sentences of the quoted passage (and developed fully in his 1994). I'm not prepared to embrace that kind of account of meaning. But I think it's separable from the basic idea that a focus on practices can be helpful in determining speech act content—especially the content that is my narrow focus: what counts as said or asserted for the purposes of distinguishing between lies and merely misleading acts? Given the results of the investigations of Chapter 2, I'm skeptical about the prospects for a definitive answer to this question coming from some prior semantic or pragmatic theory. I think we can maintain that those sorts of considerations play *some* role in determining the contents of our assertions, but that the ultimate tribunal is community practices—“whatever its members will let assertors get away with.” What counts as part of what is (strictly speaking) said depends on the kinds of practices that Brandom focuses on: if a challenge is issued to an assertor to vindicate some allegedly asserted content by

---

<sup>301</sup> *Ibid.*, p. 644



justifying it, and the challenge is considered appropriate, then the speaker has an obligation to defend that content, and it counts as part of what he said. If he can't vindicate it, he may be judged a liar. On the other hand, if the speaker is able to sidestep such a challenge (legitimately, according to the community's reckoning), on the grounds that he does not have an obligation to justify certain content (such as, say, a false implicature), then that content is not part of what he said. In such cases, he may be judged to have merely misled (if, again, something like false implicature is involved). Clearly, Brandom's overall approach to assertion is most compatible with the methodology we employed in Chapter 2: we appealed to intuitions about whether an accusation of lying would be judged appropriate, based on whether or not the alleged liar could succeed in avoiding the accusation by legitimately claiming that he is not responsible for vindicating the falsehood in question; these are just the sorts of scenarios that are at the foundation of Brandom's theory. We discovered that our intuitions about certain difficult cases were not firm. This is just what we should expect if community practices are the ultimate standard of evaluation; people will disagree about borderline cases. This is as it should be. Community practices can be the subject of debate and disagreement, and can (and should) change over time. An account of assertion, and hence of lying, in terms of such standards is attractive for just this reason: it makes room for reflection on and revision of our concept of lying, which we may want to adjust in light of various considerations, be they linguistic or moral.

### 4.3 MORAL QUESTIONS

One way we might adjust our concept of lying, for either linguistic or moral reasons (or both), is by defining it in such a way that *any* deceptive speech act counts as a lie, whether or not one believes that its content is false; it may be true (and believed-true). Thus, Sissela Bok writes, “I shall define as a lie any intentionally deceptive message which is *stated*.”<sup>302</sup> According to this definition, then, the acts that I’ve been characterizing as merely misleading would count as lies, since they’re deceptive statements. So we lie not only when what we say is false (or believed false, or what have you), but also when the implications of what we say are false—even if what we say is true. At least, this is how we would put it if we assumed a standard Gricean framework. But as we noted, this conception of lying dovetails nicely with a broader conception of semantic content, such as that provided by Jaszczolt’s Default Semantics. For her, implicatures can be included within a speaker’s primary meaning—the semantic content of her utterance—if they are part of what ideal interlocutors would intend to communicate (as speaker) and interpret (as hearer). It seems quite natural to say that if the *meaning* of a speaker’s utterance is false, then she has lied. This gives the advocate of the expanded definition of lying at least some cover from criticism to the effect that it competes with our intuitions about lying (and misleading). However, one needn’t adopt a radically contextualist position like Default Semantics to go along with a more expansive definition of lying. In his 2005, Jörg Meibauer, a mainstream linguist, “analyses falsely implicating from the point of view of Gricean theory of implicature, ...[and] argue[s] that

---

<sup>302</sup> Bok 1978, p. 13. She does not use ‘stated’ as a synonym for ‘asserted’; rather, she is contrasting deceptive speech acts (statements) from other deceptive acts, which may not involve language—deception “through gesture, through disguise, by means of action or inaction, even through silence.” (*Ibid.*)

the case of falsely implicating should be included within a general definition of lying.”<sup>303</sup> Meibauer offers an extended, disjunctive definition of lying: someone lies if and only if he asserts something he believes is false *or* conversationally implicates something he believes to be false.<sup>304</sup>

I am disinclined to accept such a broad definition of lying because it makes no distinction between true and false deceptive claims—that is, between what we’ve been calling acts of misleading and lies proper. Now, I can’t just register this complaint and leave it at that, lest I beg the question against Bok, Meibauer, and others. But I do not think that my position and theirs are on equal footing; rather, the burden of proof is on the advocate of the expanded definition of lying. She must provide compelling reasons for abandoning the distinction between lying and misleading. This is her burden since that distinction is so well-entrenched: it is a commonplace everyday usage; there is a theoretical, linguistic framework within which we can at least roughly draw it; and there is a long history of philosophical reflection on its moral importance. The distinction between lying and misleading is easy to explain, even to a child. Simple examples are enough to make it clear; the misleading grocery store sign under cans of soup, with which we began this essay, always works for me. Normal users of language make the distinction all the time. Indeed, even the Supreme Court of the United States, in *Bronston*, distinguishes between (perjurious) lying and non-felonious misleading. Grice gives us the tools to distinguish between what we say and what we implicate, which is helpful (though, as we saw in Chapter 2, hardly decisive) in making the distinction somewhat more precise. Finally, the distinction between lying and misleading has long been thought to be morally relevant. Saint Augustine, for example, defends Abraham who, in *Genesis* 20 tells king Abimelech that Sarah is his sister, which is true,

---

<sup>303</sup> Meibauer 2005, p. 1373

<sup>304</sup> *Ibid.*, p. 1382

but meant to mislead the king into thinking that Sarah is not his wife; Abraham is praiseworthy, according to Augustine, for having avoided a lie in this case.<sup>305</sup> Aquinas concurs with Augustine's defense of Abraham, and offers further that "it is lawful to hide the truth prudently, by keeping it back...."<sup>306</sup> Kant relates a personal anecdote in which he avoids lying to Friedrich Wilhelm II, but nevertheless deceives him.<sup>307</sup> The debate about the moral relevance of the lying/misleading distinction continues in contemporary philosophical discussions: Jonathan Adler (1997) sides with tradition in maintaining that it does make a moral difference; Bernard Williams (2002) and Jennifer Saul (2012 and 2013) express doubts.

In light of all this, I think we should demand some very compelling reasons for abandoning the lying/misleading distinction and embracing an expanded definition of lying. Bok doesn't offer such reasons. She is concerned exclusively with questions in moral philosophy—when and why it might be justifiable to lie—so she doesn't offer a discussion about the linguistic considerations that underwrite our distinction. This appears to be intentional: "[My] purposes are best served by concentrating on choices between truth-telling and clear-cut lying, rather than on other forms of deception such as evasion or the suppression of relevant information."<sup>308</sup> Apparently the difficult, borderline cases that have been our focus are not useful in bringing out the moral points she's interested in making. That's well and good, but it's not enough reason to abandon our focal distinction. Meibauer, on the other hand, at least attempts to provide an argument for preferring a more expansive definition of lying. Alas, it's pretty weak tea. First, Meibauer notes that lies differ from deceptions in that the latter need not be verbal (think of Kant's example of the man silently packing as if going on a trip) and, if verbal, need not involve

---

<sup>305</sup> Augustine 1887, §23

<sup>306</sup> Aquinas 1920, 2a2ae.110.3

<sup>307</sup> Reported in MacIntyre 1995, p. 336 – 337

<sup>308</sup> Bok 1978, p. xxi

assertion (one can ask a question, e.g., with a false presupposition). He then presents an argument for the expanded definition of lying, which I quote in full to avoid any appearance of being uncharitable:

In my view, however, it follows from the two differences between deception and lie listed above that false implicatures should not be treated as deceptions. To be sure, false implicatures are deceptive, but lies are deceptive, too. An advantage of the extended definition of lie... is that it makes no use of the very broad and general notion 'deception'. The use of the notions 'assertion', 'lie', and 'conversational implicature' is sufficient. Thus, in the case where the assertion p is false and the implicature q is false, we do not have to say that there is a lie, and, in addition, a deception. It is a lie, because it is false on the basis of the assertion and the implicature. A further advantage of our extended definition of lie... is that it comprises the intimate connection between the assertion and its implicature. The crucial point is that a false implicature only comes about through a verbal act of assertion to which it is bound. In contrast, a characterisation of false implicature as deception misses this connection and remains purely terminological.<sup>309</sup>

I don't really understand this argument, let alone find it convincing. Meibauer seems to think that it's best not to think of lies as sub-types of deception—despite the fact that, as he says in his second sentence above, they're deceptive! Is this because the concept of deception is too "broad" and "general"? That claim is hinted at, but I don't know. What's wrong with having broad, general notions, anyway? He claims that the expanded definition of lying allows us more simply to handle cases in which one asserts a falsehood, and in so doing also implicates a falsehood. We can just say the person lied, rather than saying that he both lied and misled. But don't we still have to note the falsity of both the assertion and the implicatum? The only advantage of Meibauer's definition is that we can call them both lies; is that really so much better than saying there's a lie and another deception? It seems worse, in that it obscures the differences between the two falsehoods. Granted, there are lots of different ways of deceiving others (Chisholm and Feehan (1977) distinguish eight), even if we restrict ourselves to verbal means of deception, and

---

<sup>309</sup> Meibauer 2005, p. 1384

lying and falsely implicating are the only ones that necessarily involve assertions, but this alone is not enough reason to lump them into the same category. We can still distinguish the two without “miss[ing the] connection” between an implicatum and the assertion to which it is bound. And again, the distinction between lying and misleading is so well-entrenched in everyday usage that Meibauer himself has trouble avoiding it. He repeatedly characterizes misleading with false implicatures as “lying while saying the truth.” At one point he says, “[T]here is a striking parallel between *genuine* lies... and lying while falsely implicating.”<sup>310</sup> In light of all this, I’m inclined to agree with an anonymous reviewer of Meibauer’s article, who characterized it as “merely a terminological contribution to the theory of lying.”<sup>311</sup>

Meibauer tries to make the case for abandoning the distinction between lying and misleading with false implicatures on purely linguistic grounds, and, in my view, he fails. It’s worth reflecting, though, on how one might make a compelling case for his conclusion. I think one would have to muster more than linguistic considerations and focus on moral ones as well. One could, perhaps, take Meibauer’s point about the close connection between assertion and implicature, combine it with observations about the difficulty in many cases of deciding how to even draw the line between what’s asserted and what’s merely implicated (pointing to my Chapter 2 for multiple examples), and add an argument that, from a moral point of view, there is no real difference between lying and merely misleading. The distinction is impossible to draw on principled linguistic grounds, and it’s ethically pernicious—giving cover to deceivers whose deeds are just as blameworthy as liars—so it should be abandoned.

There has been some discussion of the moral half of this argument in the literature. What we might call the traditional view, since it goes back at least to Augustine, is that the distinction

---

<sup>310</sup> *Ibid.*, p. 1383. Emphasis mine.

<sup>311</sup> *Ibid.*, p. 1384, n. 18

between lying and misleading is morally relevant, such that the latter is always to be preferred to the former. This often goes along with the view that lying is always wrong, without exception<sup>312</sup>—and complements it, by giving us a deceptive practice that’s allowable in the kinds of circumstances where our intuitions tell us that deceiving is the right thing to do (the old murderer-at-the-door scenario, e.g.). There is a straightforward way to challenge the traditional view; Bernard Williams puts the question pointedly: “If someone deliberately brings it about by what he says that you have a false belief, getting you to rely on something which he knows is untrue, what moral difference, if any, does it make whether he does this by lying or by relying on an implicature?”<sup>313</sup> It is, *prima facie*, strange to say that two actions with identical consequences—the hearer comes to have false beliefs—and undertaken with the same deceptive intention—to cause those false beliefs—are not equally blameworthy. The means of achieving the end shouldn’t matter. According to a simple act-consequentialism, we should say there’s no moral distinction between a lie and merely misleading speech act. Things are murkier if we shift to a rule-based approach: since the broader consequences of a general moral preference for misleading over lying might “encourage deviousness and a legalistic attempt to get away with what one can...,”<sup>314</sup> it seems doubtful that a general rule preferring misleading to lying could be justified on consequentialist grounds. There may be some special circumstances, involving, for example, tact, in which such a rule might arguably be defensible (we’ll look at such a case soon), but I suspect such circumstances are rare; they certainly can’t justify a general principle according to which misleading is better than lying. From the point of view of a kind of simple-minded virtue ethics, one might be tempted to attribute moral significance to the distinction:

---

<sup>312</sup> As in Augustine, Aquinas, Grotius, Kant

<sup>313</sup> Williams 2002, p. 101

<sup>314</sup> Adler 1997, p. 452

choosing to mislead rather than lie evinces a kind of respect for the truth or laudable desire to avoid wrongdoing; it's tempting to talk this way, say, about Abraham's misdirection and other traditional anecdotes about the saints.<sup>315</sup> But on further reflection, this looks dubious: it's certainly not the case that all non-lying misleaders act as they do out of saintly scruples; much more common are those who choose their means of deception strategically, to maintain "plausible deniability" in the face of accusations of lying. This is hardly evidence of virtue. In a more sophisticated development of a virtue-based approach to these questions, Williams defines Sincerity as "the virtue of trustworthiness in speech," and since misleaders are no more worthy of our trust than liars, then possessing the relevant virtue must involve more than merely avoiding outright lies.<sup>316</sup> Finally, a rather crude deontological approach might point to the Ten Commandments or the Five Precepts of Buddhism, and point out (or argue), that only outright lying is proscribed; this is unlikely to be universally convincing. A more sophisticated deontologist might offer more argument. Kant (*the* sophisticated deontologist) puts things in typically Kantian terms:

The human being as a moral being (*homo noumenon*) cannot use himself as a natural being (*homo phenomenon*) as a mere means (a speaking machine), as if his natural being were not bound to the inner end (of communicating thoughts).

Someone who does not believe what he says to another... has even less worth than if he were a mere thing. ...But communication of one's thoughts to someone through words that yet (intentionally) contain the contrary of what the speaker thinks on the subject is an end that is directly opposed to the natural purposiveness of the speaker's capacity to communicate his thoughts, and is thus a renunciation of his personality, and such a speaker is a mere deceptive appearance of a human being, not a human being himself.<sup>317</sup>

---

<sup>315</sup> MacIntyre 1995 relates the following story: "St. Athanasius was rowing on a river when the persecutors came rowing in the opposite direction: 'Where is the traitor Athanasius?' 'Not far away,' the Saint gaily replied, and rowed past them unsuspected." (p. 336)

<sup>316</sup> Williams 2002, p. 97

<sup>317</sup> Gregor 1996, p. 552. Kant regards lying as a violation of a duty that the agent has to himself, hence the formulation in terms of using oneself as a means.



This is a familiar Kantian formulation. But notice that it does not provide an argument that lying is worse than misleading. His formulation does restrict itself to lying (the speaker doesn't believe what he says). Why is it worse than misleading? Kant believed that the distinction was morally significant: as we noted above, he once carefully misled the King of Prussia without lying to him; apparently, shortly before Friedrich Wilhelm II's death, the King enjoined Kant to refrain from denigrating Christianity by expressing heterodox views; since Kant's views were quite heterodox, this injunction required Kant to remain silent on religious matters, so Kant promised, "as your Majesty's faithful subject, that I shall in future completely desist from all public lectures or papers concerning religion, be it natural or revealed."<sup>318</sup> Kant resumed lecturing and writing on religion after the King's death, explaining that he was no longer his "Majesty's faithful subject" at that point, since there was a new king. Alasdair MacIntyre attributes to Kant the view that one is only responsible for the content of what he asserts; anything conveyed over and above that content (implicatures, etc.) is out of his hands. "Kant therefore places himself among those who hold that my duty is to assert only what is true and that the mistaken inferences that others may draw from what I say or what I do are, in some cases at least, not my responsibility, but theirs."<sup>319</sup> The thought seems to be that the audience is (at least to some degree) responsible for making inferences, and this shifted (or shared) responsibility mitigates any badness in the deception. Jonathan Adler points out that Kant's choice of example to distinguish lying from misleading, where the onlooker infers that the man packing is going on a trip, supports this interpretation.<sup>320</sup> But there are several problems with this kind of reasoning. As Adler notes, there's a difference between communicative and non-communicative misleading. In the case of

---

<sup>318</sup> Reported in MacIntyre 1995, p. 336 – 337

<sup>319</sup> *Ibid.*

<sup>320</sup> Adler 1997, p. 444

the latter, when implicatures are involved, the speaker invites the hearer to make an inference in a way not present in the former case: implicatures, after all, are *required* to make sense of assertions as being as consistent with conversational maxims.<sup>321</sup> And any competent speaker knows this about his assertions, making him culpable at least to some degree for the hearer's inferences. In that case, responsibility for inferences from assertions must be shared between speaker and hearer. In order to maintain that the distinction between lying and misleading is morally significant, then, we must argue that in cases of shared responsibility, badness of actions is mitigated. But this does not hold in general, as Saul points out. She asks us to consider two mugging victims, one who was walking in a safe part of town in broad daylight, and another who was, as is her wont, walking in a dangerous neighborhood in the middle of the night. The second victim, it seems, is partially responsible for her fate, since she was so reckless. But this shared responsibility does not mitigate the badness of the muggers' actions.<sup>322</sup> Hence, shared responsibility does not in general have morally relevant implications, and we need—and lack—and argument that it should make a difference in the case of mere misleading.

So it seems, at least on a cursory look at the various normative-ethical approaches one might take, that it's difficult to make the case that the difference between lying and misleading ought to carry any moral weight. But again, we've only briefly considered some relatively unsophisticated and/or old arguments. A more sophisticated and recent attempt to defend the moral relevance of the distinction appears in Adler 1997. His aims in that paper are relatively modest, to show “that there is a moral asymmetry favoring falsely implicating or deceiving to lying, while holding neither that this is true in each case nor that a general deceptive strategy to

---

<sup>321</sup> *Ibid.*, p. 445

<sup>322</sup> Saul 2012, p. 5

avoid lying is moral progress.”<sup>323</sup> He gives an example of a situation in which lying is morally inferior to misleading, involving tact: a man at a party asks after the whereabouts of his friend’s absent wife. The friend replies that she’s out of town on business, which is true, but leaves it at that, concealing (by inviting a contrary inference) the fact that they’re having marital problems (and she wouldn’t have come to the party with him anyway). Adler says this “answer is appropriately tactful, and even admirable if an effortful attempt to avoid an outright lie. So the opportunity to deceive smooths over social frictions without the naked transgression of a lie.”<sup>324</sup> He argues that in cases in which we’re inclined to think that misleading is morally preferable, we’re sensitive to a feature of the circumstances under which implicatures arise, viz. a diminished demand (as opposed to expectation) of truthfulness.<sup>325</sup> We’re less likely to demand a rationale for an implicatum than the content of an assertion, since the utterance that gives rise to the implicature often states the reasons. (“Ted: ‘Whose turn is it to walk the dog?’ Marcia: ‘I did it yesterday.’”<sup>326</sup>) Adler concludes that this “difference in demands of truthfulness for assertions compared to implicatures provides a salient rationale for a corresponding ethical norm.”<sup>327</sup> This norm falls short of providing permission to deceive, though, since the *expectation* of truthfulness is still present even if the *demand* is diminished.

Adler’s argument is very subtle, and his conclusion is admirably modest, but I’m not convinced on a number of points. Adler doesn’t make it clear exactly how we move from a conversational regularity to an ethical norm. The regularity in question (utterances giving rise to implicatures often state reasons for the implicatum) isn’t all that regular (lots of utterances don’t

---

<sup>323</sup> Adler 1997, p. 452

<sup>324</sup> *Ibid.*, p. 447

<sup>325</sup> *Ibid.*, p. 449

<sup>326</sup> *Ibid.*

<sup>327</sup> *Ibid.*, p. 451

give the reasons for the implicata they generate), so it seems hard to extend this rationale for distinguishing lying and misleading to other examples that don't involve the regularity. Also, the intuition that he aims to explain—that in cases where we have some legitimate reason to deceive (such as tact), lying is worse than misleading—is not one that I necessarily share. It's at least possible to explain the prevalence of such intuitions as a vestige of the sorts of hard-line views about lying canvassed above, according to which lying is always wrong; after all, many of us grew up heavily influenced by such manichean doctrines.<sup>328</sup> It's even possible to reject the notion that misleading is better than lying, but still to have an intuition that the person who chose to mislead rather than lie did something praiseworthy: he's mistaken about the moral significance of his choice, but that mistake is widespread, and he tried to do what he (and many others, possibly including his interlocutor) believe to be the better thing; this is perhaps revealing of the quality of his character, and should be lauded. Adler sometimes argues as though his suggested moral norm—derived from the supposed linguistic regularity involving a reduced demand for truthfulness on implicatures—is needed not only to explain our intuitions in cases involving a laudable choice to mislead rather than lie, but more generally from the frequency of social situations in which there is pressure on us to deceive: borrowing a phrase from John Rawls, he says, “[A] norm corresponding to the lessened demands of truthfulness for implicatures would be desirable for all... [g]iven the previously mentioned ‘strains of commitment’ generated by the numerous situations pressuring us to deceive.”<sup>329</sup> But if this is the motivation, we don't need a norm that rests on the dubious linguistic principle of a lessened demand of truthfulness for implicatures; as Saul points out, a plausible moral principle might

---

<sup>328</sup> This Catholic boy certainly did.

<sup>329</sup> *Ibid.*

simply involve “a lessened demand for truthfulness when one has legitimate reason to deceive.”<sup>330</sup> Saul puts the rationale so nicely, I’ll just quote her:

If the reason for allowing some kinds of deception is that sometimes we have legitimate reason to deceive, this norm makes much better sense than one that focuses on method of deception. Compare the case of violence. In general, we think that violence is bad; but sometimes, e.g. in self-defence, we think it is legitimate. We accommodate this fact not by allowing that, say, shootings with one brand of gun are better than shootings with another, but by allowing that violence for the sake of self-defence is better than other violence.<sup>331</sup>

This alternative derivation of the relevant moral maxim is appealingly direct: rather than attempt to justify intuitively innocent deceptive practices from tenuously related and narrowly applicable conversational principles, we can point directly to the moral features of the circumstance at hand. This does not, of course, allow for any moral difference between lying and misleading, and it does not explain intuitions we may have in certain cases that the latter is preferable to the former.

I suggested an explanation above; Saul offers a similar one:

A decision to mislead may reveal an admirable desire to mitigate the wrong of one’s deception: Since many people hold the false belief that misleading is better than lying, many people think they do something better by misleading than by lying. A person concerned with being moral, and troubled by deceit, will often make the effort to craft a merely misleading utterance in the hope of doing something less bad. That what they do is in fact equally bad does not undermine the fact that a person like this, who tries to act morally, is more admirable than one who simply does not try because they do not care.<sup>332</sup>

I thus tentatively conclude, with Saul and Williams, that the lying/misleading distinction has little if any moral significance. I do so tentatively, though, since this is a large topic deserving much more sustained examination—and I have barely begun that task.<sup>333</sup> It may be that a compelling case can be made, but my preliminary investigation suggests to me that the task would be quite difficult. I only considered the question, though, as part of an imagined

---

<sup>330</sup> Saul 2012, p. 6

<sup>331</sup> *Ibid.*

<sup>332</sup> *Ibid.*

<sup>333</sup> Saul 2013, I gather, takes up the challenge.

argument for the abandonment of the distinction: because it is theoretically unfounded (see Chapter 2) and morally otiose, it should be abandoned. Even if the second premise of this argument looks like it may be true, I've already considered and rejected the idea of abandoning the distinction between lying and misleading on the grounds that it can't be made in a theoretically robust, principled way (cf. §3.1.4 above): that a distinction resists theoretical codification, that it may be vague—allowing for indeterminate borderline cases—does not mean that there is no such distinction. And the fact that the lying misleading distinction continues to play a role in our lives is reason to hold on to it. Saul and Williams both point to similar examples: there is a reduced expectation of truthfulness (to put it in Adler's terms) in certain specialized contexts, so that the difference between an outright lie and a false implicature takes on legal, if not moral, force—namely, in courtrooms (Saul's example) and in the British Parliament (Williams'). In courtrooms, as we've seen, the norm against deception is provided by perjury law, which requires outright falsehood for conviction; mere misleading, as we've called it, is perfectly legal on the witness stand. This lack of expectation of truthfulness, especially in the questioning of defendants, is a natural adjunct to the principle—codified in the Fifth Amendment to the U.S. Constitution—that citizens have a right against self-incrimination. While it may be difficult to find moral justification for such a principle (there is a large literature on this topic, featuring widespread skepticism about moral arguments), it is nevertheless a deeply held, widespread legal norm. Similar norms (with even more dubious prospects of moral justification, it seems to me) apply to members of the British Parliament. As Williams notes, "[M]inisters may not lie when answering questions or making statements, but they can certainly omit, select, give answers that reveal less than the whole relevant truth, and generally give a misleading impression. (There is indeed an offence of 'misleading the House' which falls short of straight

lying, but the general idea is on the lines of the traditional distinction.)”<sup>334</sup> So, despite its shaky moral status, I’m still disinclined to give up our focal distinction; its pervasiveness in both everyday life and in legal settings makes it worth holding on to. And its formal legal status makes it, as a purely practical matter, a distinction worthy of serious consideration, lest it be abused to the detriment of justice.

#### 4.4 LEGAL MATTERS

Many legal scholars and jurists have subjected the lying/misleading distinction to scrutiny, with just that concern in mind. The most obvious statutory application of the distinction is to perjury, and there has been debate over the correct way to interpret the law proscribing that particular form of deception. The relevant U.S. statute defines perjury as occurring when a witness “states or subscribes any material matter which he does not believe to be true.”<sup>335</sup> The most prominent case in which interpretation of this language was at issue is *Bronston v. United States*, which we have seen a couple of times already. A reminder of the relevant facts: in Bronston’s original trial, the following exchange occurred: “Q. Do you have any bank accounts in Swiss Banks, Mr. Bronston? A. No, sir. Q. Have you ever? A. The company had an account there for about six months, in Zurich.”<sup>336</sup> Bronston’s final answer was true; what he didn’t say was that he had also had a personal account in Zurich. Bronston was convicted of perjury after the original trial, but the Supreme Court overturned that ruling, maintaining that the literal truth of Bronston’s answer precludes a perjury conviction. That answer is a classic instance of false particularized

---

<sup>334</sup> Williams 2002, pp. 108 – 109

<sup>335</sup> 18 USC. § 1621.

<sup>336</sup> *Bronston v. United States*, p. 354

implicature, and so clearly the Supreme Court's view is that the lying/misleading distinction does and ought to have legal significance. Some scholars share that view; others disagree, arguing not only that *Bronston* was decided incorrectly, but that the decision sets a dangerous precedent not only in perjury law, but in general.

Stuart Green is among those who agree with the Court. In his 2000 and especially in his 2001, he argues first that the lying/misleading distinction has *moral* significance, and that this maps neatly onto its legal significance. Those who merely mislead—paradigmatically, by uttering truths that give rise to false implicatures—“should be regarded as less fully culpable than if [they] had lied.”<sup>337</sup> A principle of *caveat auditor* applies. And just as there is a moral difference between the two kinds of deceptive acts, there is a corresponding legal difference: perjury involves lying, while other offenses—like the various kinds of criminal fraud—can involve merely misleading. The Court, therefore, gets it right in *Bronston*. The principle of *caveat auditor* is even more clearly applicable in a courtroom: one of the main differences between acts of lying and misleading “is that the latter afford the listener the opportunity for more precise questioning, which... lies generally do not. This distinction applies *a fortiori* in the courtroom. A lawyer who fails to clarify evasive or nonresponsive statements from a witness bears even more responsibility for improper inferences than does a listener in everyday conversation.”<sup>338</sup> The Court itself avers that “it is the lawyer's responsibility... to flush out the whole truth with the tools of adversarial examination.”<sup>339</sup>

We can criticize Green's work, I think, from a number of angles. First, he seems to rely on a rather naïve view about meaning to draw his distinction between lying and misleading. The

---

<sup>337</sup> Green 2001, p. 160

<sup>338</sup> *Ibid.*, p. 177

<sup>339</sup> *Bronston*, pp. 358 – 359



difference between the two is just “the difference between (1) asserting what one believes is literally false, and (2) leading the listener to believe something false by saying something that is either true or has no truth value.”<sup>340</sup> To determine whether an assertion is literally false, one checks its literal meaning, which “is derived, roughly speaking, by determining the meaning of the individual words... and applying the grammatical rules of the language to those words.”<sup>341</sup> And the literal meaning of an utterance is to be distinguished from what a speaker intends to convey by uttering it. If only things were so simple. It is painfully clear to us by now that this rough sketch elides any number of vexing complications; I won’t even bother listing them at this point. Second, I think Green does too little to establish his claim that “*ceteris paribus*, lying is more wrongful than merely misleading.”<sup>342</sup> He spends only a few pages on the topic, citing, e.g., the dubious proposition (considered above) that in cases of misleading the listener shares some responsibility for drawing the inference (this seems to be the thrust of his appeal to a *caveat auditor* principle). He points to the Judeo-Christian tradition of making the distinction, and the Jesuits’ practice of “mental restriction”; again, we’ve considered these already, and found them wanting as justifications for a real moral distinction. It’s quite difficult to make the case that the lying/misleading distinction has moral significance, and Green has not done nearly enough to establish that claim. However, we should say, again, as already noted above, that a lack of moral significance need not imply that the distinction lacks *legal* significance, especially in the context of perjury, where the speech acts under consideration take place in a special, formalized context—a courtroom, in which, arguably, there is a reduced expectation of truthfulness and a

---

<sup>340</sup> Green, *op. cit.*, p. 164

<sup>341</sup> *Ibid.*

<sup>342</sup> *Ibid.*, p. 168

principle of *caveat auditor* is appropriately applied, e.g., to attorneys with the power and responsibility to ask probing follow-up questions.

The more serious objection, though, is the first: with a too-simplistic conception of literal meaning, one runs the risk of defining perjury so narrowly as to allow clearly specious defenses against lying to win the day in court. Some scholars maintain that the *Bronston* decision opens the door to this kind of pernicious “Literalism,” as William Simon calls it, exemplified in recent history by Bill Clinton’s legal maneuverings during the Monica Lewinsky affair, the Paula Jones trial, etc. Clinton, famously, quibbled over the proper understandings of the word ‘is’ and the phrase ‘sexual relations’, the latter of which, he claimed, he understood to preclude fellatio. Peter Tiersma, in his 1990, argues at length that the Court’s “literal falsity” requirement for perjury is an untenable standard, and “that the central issue in determining whether a false statement has been made ought to be what the witness meant by his statement, rather than what the words of the witness literally mean.”<sup>343</sup> That is, he argues that speaker meaning, rather than merely what is said—to put it in Gricean terms—should be evaluated for truth or falsity in perjury cases. Tiersma explicitly appeals to the Gricean framework to make his case: he argues that speakers in a courtroom ought to be expected to adhere to the various conversational maxims, and that exploitation of them to communicate falsehoods ought to be considered perjurious. His overall strategy is familiar to us: he considers a variety of scenarios—some real, some invented—in which it’s difficult to say what counts as falling on either side of the Gricean divide between saying and implicating; he also considers scenarios in which there’s a strong intuition to include implicata as part of what is said. In other words, he does much the same thing that we did in Chapter 2. So, for example, he considers the question “Do you drive a Chevy?” and imagines the

---

<sup>343</sup> Tiersma 1990, p. 375

response “I drive a Ford.”<sup>344</sup> Supposing the person in question drives both a Ford and a Chevy, but wants to conceal (for some reason) the latter, has she lied/perjured herself in this case? What she said is arguably literally true. He looks at the case of *United States v. Earp*, in which the defendant is an incompetent racist, who, on a number of occasions had *attempted* to burn crosses in the lawns of inter-racial couples, but had failed to get them to light. When asked whether he had ever burned any crosses, he answered that he hadn’t; his conviction on perjury charges was overturned, appealing to the *Bronston* standard.<sup>345</sup> But in *Harrison v. State*, an Indiana case, a certain trustee reported that he had received \$500 from a bankruptcy settlement, when in fact he had received \$1500; the court was not impressed with his defense that it was literally true that he had received \$500.<sup>346</sup>

From these and many more examples, Tiersma concludes that the notion of literal truth cannot be satisfactorily defined, and so it should be abandoned as a legal standard. This is a tempting conclusion to draw; it mirrors our own reasoning in Chapter 2, in which we concluded that there is no precise way to draw the distinction between what is said (strictly speaking) and what is otherwise conveyed. The sorts of examples Tiersma cites, and the variety of intuitions they trigger, are just the sorts of examples we considered. But while the conclusion that there is no clear line to be drawn between lying and misleading is justified, I think it is hasty to conclude from this, as Tiersma does, “that whether a witness makes a false statement should depend on what the witness intended to communicate-what the witness meant by the statement.”<sup>347</sup> This would amount to throwing out the distinction between lying and misleading entirely, and counting false implicatures as perjurious. I’ve considered this possibility before (cf. §3.1.4 and

---

<sup>344</sup> *Ibid.*, p. 389

<sup>345</sup> *Ibid.*, p. 395

<sup>346</sup> *Ibid.*, p. 397

<sup>347</sup> *Ibid.*, p. 393

§4.3 above) and rejected it: the fact remains that the lying/misleading distinction, even if it can't be underwritten by a satisfactory semantic or pragmatic theory, even if it is vague, is one that normal speakers naturally make. Among those normal speakers are legislators, and they have made it clear that they intend the distinction to have legal force: as Lawrence Solan points out, "legislators know how to write a statute that includes both false statements and misleading statements [as is the case in statutes covering fraud]..., but did not write the perjury statute that way."<sup>348</sup>

It is a separate, and much more difficult, question whether the perjury statute *ought* to be written that way. Many of Tiersma's arguments seem to be directed at this issue. Thus he notes, I think compellingly, that witnesses in court take an oath not only to tell the truth, but the "whole truth."<sup>349</sup> This is, on its face, a very strong promise; arguably, it's a promise to be a maximally cooperative conversational participant, one who holds nothing back, answers questions straightforwardly, and follows all of the Gricean conversational maxims. This seems to be the standard that Tiersma envisions for witness behavior: "[Jurors] will interpret what is said by a witness roughly as they would utterances heard outside the courtroom. Therefore, in evaluating the statements of a witness for purposes of the law of perjury, there is no reason to apply distinct rules of interpretation."<sup>350</sup> That is, jurors should assume that Gricean maxims apply in the courtroom, and that witnesses are adhering to them. This would arguably have a salutary effect. Simon argues that overly literalistic standards for testimony threaten the rule of law, in the sense that "requires that enforcers have material information available to them. Literalism threatens

---

<sup>348</sup> Solan 2002, p. 190

<sup>349</sup> Tiersma, *op. cit.*, p. 395

<sup>350</sup> *Ibid.*, p. 403

this pre-requisite by permitting the withholding of material information.”<sup>351</sup> The probability that justice will be served varies directly with the amount and quality of factual information judges and jurors have access to; a regime under which witnesses were held to a higher standard of truthfulness-expectation would thus promote fair outcomes. There are, however, countervailing norms that militate toward granting some latitude to witnesses, especially defendants. As Simon himself admits, “there is an asymmetry of stakes peculiar to criminal prosecution. We deem erroneous conviction far more costly than erroneous acquittal.”<sup>352</sup> This may explain and partially justify a higher tolerance for misleading in criminal, as opposed to civil, contexts. Simon also allows that it is plausible to say that “implicit deception is less blameworthy than explicit deception because the deceiver is less active and because the victim's sense of betrayal will be weaker. The claim rests on the omission/commission distinction that, though sometimes hard to justify in principle, has strong support in intuition and convention.”<sup>353</sup> Thus the felt difference between outright lying and (many instances of) mere misleading. Simon speculates: “If we ask why implicit deception seems less bad, the answer is likely to be that it is closer to the situation where the subject is entirely silent.”<sup>354</sup> This seems right, and is closely related to traditional norms according to which it is permissible, or at least less-blameworthy than it otherwise would be, for accused persons to avoid incriminating themselves. This principle is codified in the Fifth Amendment to the U.S. Constitution, which explicitly permits silence. And there is a related norm, traditionally accepted but lately repudiated (by the Supreme Court in 1998, in *Brogan v. United States*), that even allows for false denials: the so-called “exculpatory no” doctrine. According to this principle, one could lie to investigators, provided the falsehood consisted of a

---

<sup>351</sup> Simon 2002, p. 1893

<sup>352</sup> *Ibid.*, p. 1898

<sup>353</sup> *Ibid.*, p. 1895

<sup>354</sup> *Ibid.*

simple denial of wrongdoing, and not be subject to prosecution for making false statements. As Green points out, prior to 1998, there was wide agreement among lower courts about the legitimacy of this norm.<sup>355</sup>

Thus it is a difficult question whether perjury law ought to be written the way it is. We cannot decide the issue here. But given that it is in fact written explicitly to proscribe only falsehoods, I think it's hard to fault the Supreme Court for its decision in *Bronston*. There are cases of clear-cut lies and cases that are clearly acts of mere misleading; and so there are cases of clear-cut perjury, and cases in which, though misleading, the defendant clearly did not "state [anything] which he does not believe to be true." There are also, of course, many borderline cases in between. But the law clearly puts the onus on the courts to decide difficult cases. People can have various opinions, I think, about *Bronston*'s utterance, but given the statute, and given the criminal standard of proof (beyond a reasonable doubt), it seems like the Court made the proper call. And concerns that their talk of "literal truth" will lead to widespread unjust exonerations are, I think, overblown. The Court, in its decision, is explicitly sensitive to the fact that "literal truth" defenses have limits. They consider an example proffered by the District Court: "[I]f it is material to ascertain how many times a person has entered a store on a given day and that person responds to such a question by saying five times when in fact he knows that he entered the store 50 times that day, that person may be guilty of perjury even though it is technically true that he entered the store five times."<sup>356</sup> The Court rejects the lower court's reasoning: "[I]t is doubtful that an answer which, in response to a specific quantitative inquiry, baldly understates a numerical fact can be described as even 'technically true'."<sup>357</sup> This recalls

---

<sup>355</sup> Green 2001, p. 199

<sup>356</sup> *Bronston*, p. 355 n. 3

<sup>357</sup> *Ibid.*

our discussion of GCIs—specifically, Q-Implicatures of the scalar variety, one of which involves number-terms, where the literal sense is supposed to be ‘at least’ and the implicatum ‘at most’. We argued that in cases of deception, a defense against lying that appealed to the alleged literal meaning being ‘at least’ would fail. The Court agrees. Again, this is entirely proper: there are a variety of phenomena to consider on the boundary between lies and other kinds of deceptions, and judgments will have to be made on a case-by-case basis. Given the statutory definition of perjury, courts will have to make such judgments. The Supreme Court seems to recognize the need for flexibility.

A recent case from my current home state of Wisconsin provides another nice case-study highlighting the need to be sensitive to Chapter 2 considerations about the viability of defenses against accusations of lying in various circumstances. This is not a perjury case, but rather a judicial ethics investigation. The defendant, then a candidate for but now a member of the Wisconsin Supreme Court, Michael Gableman, was accused of committing an ethics violation by running a deceptive television ad against his opponent for the seat, then-Justice Louis Butler. Here’s the text of the advertisement:

Unbelievable. Shadowy special interests supporting Louis Butler are attacking Judge Michael Gableman. It’s not true! Judge, District Attorney, Michael Gableman has committed his life to locking up criminals to keep families safe—putting child molesters behind bars for over 100 years. Louis Butler worked to put criminals on the street. Like Reuben Lee Mitchell, who raped an 11-year-old girl with learning disabilities. Butler found a loophole. Mitchell went on to molest another child. Can Wisconsin families feel safe with Louis Butler on the Supreme Court?<sup>358</sup>

Butler had been the defense attorney for the rapist in question, from 1985 – 1988. He had indeed successfully appealed his client’s conviction. That ruling, however, was overturned by the

---

<sup>358</sup> *Wisconsin Judicial Commission v. Gableman*, Judicial Conduct Panel Decision, p. 8 [Hereafter, ‘Panel Decision’]

Supreme Court, and Butler's client was sent to prison. He was released in 1992, when Butler was no longer his attorney. He was convicted of sexual assault in 1995.

Gableman agreed to the finding of fact that “[n]othing that Justice Butler did in the course of his representation of Mitchell caused, facilitated, or enabled Mitchell’s release from prison in 1992,” and further that “[n]othing that Justice Butler did in the course of his representation of Mitchell had any connection to Mitchell’s commission of a second sexual assault of a child.”<sup>359</sup> Obviously, the advertisement suggests otherwise.

The ethics rule that Gableman was alleged to have violated says this:

*Misrepresentations.* A candidate for a judicial office shall not knowingly or with reckless disregard for the statement’s truth or falsity misrepresent the identity, qualifications, present position, or other fact concerning the candidate or an opponent. A candidate for judicial office should not knowingly make representations that, although true, are misleading, or knowingly make statements that are likely to confuse the public with respect to the proper role of judges and lawyers in the American adversary system.<sup>360</sup>

It’s important to note that the difference between the words “shall” and “should” is significant in this context: proscriptions of conduct beginning with “shall not” are mandatory; those beginning with “should not” are merely “aspirational.” Hence, although it’s arguable that Gableman violated the second condition (he argued before the Panel that he had), since it is aspirational, he was not subject to sanction for that violation. Only the first condition is in play. The question is how to interpret that condition. Citing the need to interpret statutes in such a way as to avoid superfluity, the Panel concluded that the first condition cannot be interpreted to proscribe true but misleading claims, since the second condition explicitly addresses those. They conclude that “it stands to reason that the first sentence must apply to statements that, standing alone, are

---

<sup>359</sup> Panel Decision, p. 10

<sup>360</sup> Supreme Court Rule (SCR) 60.06(3)(c)



false.”<sup>361</sup> This interpretation of the rule makes this case directly comparable to those involving perjury: there is a false-statement requirement. Two of the three judges on the Panel argued that, since the individual sentences in the advertisement were all true, there was no false statement, and so Gableman had not violated the mandatory portion of the rule. Enough members of the Supreme Court agreed with this reasoning, and Gableman was exonerated.<sup>362</sup>

I would argue that the Panel and (half of) the Court got this one wrong. Considerations from Chapter 2 can guide us here. The offending portion of the advertisement is this: “Butler found a loophole. Mitchell went on to molest another child.” The Panel (along with three Justices) notes that both sentences are true, and so Gableman cannot be said to have made a false statement. The linguistic phenomenon on display in this case is parataxis—the concatenation of clauses without conjunctions. Levinson argues that this gives rise to I-implicatures:

“John turned the switch. The motor started.”

I++> ‘John turned the switch and hen as an intended result the motor started.’<sup>363</sup>

He points out that many have argued that parataxis should be interpreted as “elided conjunction,” but that the range of implicata is even wider for parataxis, as this example demonstrates: “John fell and broke his leg. He lost his grip on the cliff.” If ‘and’ were inserted between the two sentences, the natural reverse-temporal reading would be unavailable. Levinson also notes that many languages get along without conjunctions, relying instead on parataxis.<sup>364</sup> We considered conjunctions above (§2.2.2), and noted the phenomenon of “buttressing,” whereby it’s natural to read them as carrying temporal, causal, and other types of information. Further, we concluded

---

<sup>361</sup> Panel Decision, p. 14. Some Justices of the Supreme Court, who subsequently took up this matter, disagreed with the Panel’s reading of the rule here, but I’ll set that issue aside, since the Panel’s reading makes it possible to bring Chapter 2 considerations to bear on this case.

<sup>362</sup> The Court was split 3 – 3, unsurprisingly along partisan lines (judges in Wisconsin declare their political party affiliations). Since neither side had a majority, the Panel’s decision was upheld by default.

<sup>363</sup> Levinson 2000, p. 124

<sup>364</sup> *Ibid.*, p. 125

that in many cases, the evidence was compelling that these sorts of I-implicatures ought to be included as part of what is said. I think these arguments apply to the Gableman ad: it seems quite natural to say that it *said* or *claimed* that there was a close temporal and causal connection between Butler's finding a loophole and Mitchell going on to molest another child. I am therefore with the one member of the Panel, Ralph Fine, who dissented from the others' conclusions on this question, writing, "The 'fact' asserted in the advertisement, by its language and the juxtaposition of that language, is that Justice Butler did something when he was a lawyer representing Mitchell that permitted Mitchell to commit another sex crime. There is in my view no other way to read the advertisement's two key sentences...."<sup>365</sup> An interpretation of sentences according to which the falsity is merely implied "is a crabbed reading, lashed to the mast of a sentence-by-sentence literalism, and ignores the way we use language...."<sup>366</sup> I agree, and wish I had turned that phrase myself.

The situation in the Gableman case is exactly analogous to that faced in *Bronston* and any other perjury trial. A determination must be made as to whether or not the defendant has lied; this requires that what he *said*, strictly speaking, was false. Judgments about such matters can be difficult to make, as there are a wide variety of linguistic phenomena which the would-be deceiver can take advantage of. Chapter 2, while not exhaustive, surveyed a broad range of these, and its conclusions can be helpful in guiding our judgments. It is satisfying to conclude our sometimes abstruse theoretical investigations with a clear example of their potential for practical applicability.

---

<sup>365</sup> Panel Decision, p. 23

<sup>366</sup> *Ibid.*, pp. 23 – 24

## BIBLIOGRAPHY

- Adler, J. 1997. "Lying, Deceiving, or Falsely Implicating." *The Journal of Philosophy* 94: 435 – 452.
- Aquinas, T. 1920. *The Summa Theologica of St. Thomas Aquinas*. Literally translated by Fathers of the English Dominican Province. London: Burns, Oates, and Washbourne.
- Arico, A. and Fallis, D. (forthcoming). "Lies, Damned Lies, and Statistics: An Empirical Investigation of the Concept of Lying." *Philosophical Psychology*.
- Atlas, J. 2005. *Logic, Meaning, and Conversation: Semantical Underdeterminacy, Implicature and Their Interface*. New York: Oxford University Press.
- Atlas, J. 2007. "Meanings and Semantical Underdeterminacy." in Preyer and Peter 2007: 217 – 239.
- Augustine, A. 1887. "Against Lying." in Browne, H. (tr.), Schaff, P. (ed). 1887. *Nicene and Post-Nicene Fathers, First Series, Vol. 3*. Buffalo, NY: Christian Literature Publishing Co.
- Austin, J. 1962. *How to Do Things with Words*. Cambridge, MA: Harvard University Press.
- Bach, K. 1994. "Conversational Implicature." *Mind and Language* 9: 124 – 162.
- Bach, K. 1997. "The Semantics-Pragmatics Distinction: What It Is and Why It Matters." *Linguistische Berichte* 8, *Special Issue on Pragmatics*, 33-50.
- Bach, K. 1999. "The Myth of the Conventional Implicature." *Linguistics and Philosophy* 22: 327 – 366.
- Bach, K. 2001. "You Don't Say." *Synthese* 128: 15 – 44.
- Bach, K. 2002. "Seemingly Semantic Intuitions." in Campbell, J.K. and O'Rourke, M. (eds.). *Meaning and Truth*. New York: Seven Bridges Press. 21 – 33.
- Bach, K. 2004. "Pragmatics and the Philosophy of Language." in Horn and Ward 2004, 463 – 487.

- Bach, K. 2005. "Context *ex Machina*" in Szabó 2005, 15 – 44.
- Bach, K. 2006. "The Excluded Middle: Semantic Minimalism without Minimal Propositions." *Philosophy and Phenomenological Research* 73: 435 – 442.
- Bach, K. 2007. "Review of Emma Borg's *Minimal Semantics*." *Philosophical Review* 116: 303 – 306.
- Bok, S. 1978. *Lying: Moral Choice in Public and Private Life*. Sussex: The Harvester Press.
- Borg, E. 2004. *Minimal Semantics*. Oxford: Oxford University Press.
- Borg, E. 2007. "Minimalism vs. Contextualism in Semantics." in Preyer and Peter 2007, 339 – 359.
- Borg, E. 2012. "Semantics Without Pragmatics?" in Allen, K. and Jaszczolt, K. (eds.). 2012. *The Cambridge Handbook of Pragmatics*. Cambridge, UK: Cambridge University Press. 513 – 528.
- Brandom, R. 1983. "Asserting." *Noûs* 17: 637 – 650.
- Brandom, R. 1994. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge: Harvard University Press.
- Bratman, M. 1992. "Practical Reasoning and Acceptance in a Context." *Mind* 101: 1 – 15.
- Camp, E. 2007. "Prudent Semantics Meets Wanton Speech Act Pluralism." in Preyer and Peter 2007, 194 – 213.
- Cappelen, H. and Lepore, E. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.
- Carnap, R. 1942. *Introduction to Semantics*. Cambridge, MA: Harvard University Press.
- Carson, T. 2006. "The Definition of Lying." *Noûs* 40: 284 – 306.
- Carson, T. 2009. *Lying and Deception*. Oxford: Oxford University Press.
- Carston, R. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Chisholm, R. and Feehan, T. 1977. "The Intent to Deceive." *The Journal of Philosophy* 74: 143 – 159.
- Chomsky, N. 1957. *Syntactic Structures*. The Hague: Mouton & Co.
- Chomsky, N. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Clapp, L. 2012. "Three Challenges for Indexicalism." *Mind and Language* 27: 435 – 465.

- Coen, J. and Coen, E. (Producers and Directors) 2001. *The Man Who Wasn't There*. Los Angeles: Working Title Films.
- Cohen, J. and Rickless, S. 2007. "Binding Arguments and Hidden Variables." *Analysis* 67: 65 – 71.
- Coleman, L. and Kay, P. 1981. "Prototype Semantics: The English Verb *lie*." *Language* 57: 26 – 44.
- Davidson, D. 2001. *Essays on Actions and Events*. Oxford: Clarendon Press.
- DeRose, K. 1996. "Knowledge, Assertion, and Lotteries." *Australasian Journal of Philosophy* 74: 568 – 580.
- DeRose, K. 2002. "Assertion, Knowledge, and Context." *The Philosophical Review* 111: 167 – 203.
- Dummett, M. 1973. *Frege: Philosophy of Language*. London: Duckworth.
- Fallis, D. 2009. "What Is Lying?" *The Journal of Philosophy* 106: 29 – 56.
- Fodor, J. 1975. *The Language of Thought*. New York: Thomas Crowell.
- Fodor, J. 1983. *Modularity of Mind*. Cambridge, MA: MIT Press.
- Frankfurt, H. 2005. *On Bullshit*. Princeton, N.J.: Princeton University Press.
- Gibbs, R. and Moise, J. 1997. "Pragmatics in Understanding What is Said." *Cognition* 62: 51-74.
- Grice, P. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Green, S. 2000. "Deceit and the Classification of Crimes: Federal Rule of Evidence 609(a)(2) and the Origins of 'Crimen Falsi'." *The Journal of Criminal Law and Criminology* 90: 1087-1124.
- Green, S. 2001. "Lying, Misleading, and Falsely Denying: How Moral Concepts Inform the Law of Perjury, Fraud, and False Statements." *Hastings Law Journal* 53: 157 – 212.
- Gregor, M. (ed. and trans.). 1996. *Immanuel Kant, Practical Philosophy*. Cambridge: Cambridge University Press.
- Grotius, H. 2005. *The Rights of War and Peace*. Tuck, R (ed.), from the Edition by Jean Barbeyrac. Indianapolis: Liberty Fund.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Horn, L. 1989. *A Natural History of Negation*. Chicago: University of Chicago Press.
- Horn, L. 2004. "Implicature." in Horn and Ward 2004, 3 – 28.

- Horn, L. 2009. "WJ-40: Implicature, Truth, and Meaning." *International Review of Pragmatics* 1: 3 – 34.
- Horn, L. and Ward, G. (eds.). 2004. *The Handbook of Pragmatics*. Malden, MA: Blackwell.
- Jaszczolt, K. M. 2009. "Cancellability and the Primary/Secondary Meaning Distinction." *Intercultural Pragmatics* 6: 259 – 289.
- Jaszczolt, K. M. 2010(a). "Defaults in Semantics and Pragmatics." *The Stanford Encyclopedia of Philosophy* (Winter 2010 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2010/entries/defaults-semantics-pragmatics/>>.
- Jaszczolt, K.M. 2010(b). "Default Semantics." in Heine, B. and Narrog, H. (eds.). 2010. *The Oxford Handbook of Linguistic Analysis*. Oxford: Oxford University Press. 193 – 222.
- Jaszczolt, K.M. 2012. "Semantics and Pragmatics: The Boundary Issue." in Meinenborn, C., von Stechow, P., and Portner, P. (eds.). 2012. *Semantics: An International Handbook of Natural Language Meaning, Vol. 3*. Berlin: Walter de Gruyter GmbH. 2333 – 2359.
- Kamp, H. and Reyle, U. 1993. *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Dordrecht: Kluwer.
- Kant, I. 1797. "Über ein vermeintes Recht aus Menschenliebe zu lügen." *Berlinische Blätter* September 10, 1797: 301 – 314.
- Kaplan, D. 1989. "Demonstratives." in Almog, J., Wettstein, H., and Perry, J. (eds.). 1989. *Themes from Kaplan*. New York: Oxford University Press. 481 – 536.
- King, J. and Stanley, J. 2005. "Semantics, Pragmatics, and the Role of Semantic Content." in Szabó 2005: 111 – 164.
- Knobe, J. and Nichols, S. 2007. "An Experimental Philosophy Manifesto." in Knobe, J. and Nichols, S. (eds.). 2007. *Experimental Philosophy*. New York: Oxford University Press. 3 – 16.
- Koethe, J. 2009. "Knowledge and the Norms of Assertion." *Australasian Journal of Philosophy* 87: 625 – 638.
- Korta, K. and Perry, J. 2011. *Critical Pragmatics: An Inquiry into Reference and Communication*. Cambridge: Cambridge University Press.
- Lackey, J. 2007. "Norms of Assertion." *Noûs* 41: 594 – 626.
- Levinson, S. 2000. *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Cambridge, MA: The MIT Press.

- Lewis, D. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8: 339 – 359.
- Luntz, F. 2007. *Words That Work: It's Not What You Say, It's What People Hear*. New York: Hyperion.
- MacFarlane, J. 2007. "Semantic Minimalism and Nonindexical Contextualism." in Preyer and Peter 2007, 240 – 250.
- MacFarlane, J. 2009. "Nonindexical Contextualism." *Synthese*, 166: 231-250.
- MacFarlane, J. 2011. "What Is Assertion?" in Brown, J. and Cappelen H. (eds.). 2011. *Assertion: New Philosophical Essays*. Oxford: Oxford University Press. 79 – 96.
- MacIntyre, A. 1994. "Truthfulness, Lies, and Moral Philosophers: What Can We Learn from Mill and Kant?" in *Tanner Lectures on Human Values*, vol. 16. Salt Lake City: University of Utah Press.
- Mahon, J. 2008(a). "Two Definitions of Lying." *International Journal of Applied Philosophy* 22: 211 – 230.
- Mahon, J. 2008(b). "The Definition of Lying and Deception." *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2008/entries/lying-definition/>>.
- Mahon, J. 2009. "The Truth about Kant on Lies." in Martin, C. (ed.). 2009. *The Philosophy of Deception*. Oxford: Oxford University Press.
- Mannison, D.S. "Lying and Lies." *Australasian Journal of Philosophy* 47: 132 – 144.
- Marley, P. 2009. "Gableman Ethics Case Hinges on Defining Lie." *Milwaukee Journal Sentinel*, September 14, 2009.
- Meibauer, J. 2005. "Lying and Falsely Implicating." *Journal of Pragmatics*, 37: 1373 – 99.
- Meibauer, J. 2011. "On Lying: Intentionality, Implicature, and Imprecision." *Intercultural Pragmatics*, 8: 277 – 92.
- Neale, S. 2007. "Heavy Hands, Magic, and Scene-Reading Traps." *European Journal of Analytic Philosophy* 3: 77 – 132.
- Nicolle, S. and Clark, B. 1999. "Experimental Pragmatics and What Is Said: A Response to Gibbs and Moise." *Cognition* 69: 337-354.
- Pagin, P. MS. "Against Normative Accounts of Assertion."
- Pagin, P. 2008. "Assertion." *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2008/entries/assertion/>>.

- Perry, J. 2001. *Reference and Reflexivity*. CSLI Publications.
- Preyer, G. and Peter, G. (eds.). 2007. *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics*. Oxford: Oxford University Press.
- Recanati, F. 2004(a). “Pragmatics and Semantics.” in Ward and Horn 2004, 442 – 462.
- Recanati, F. 2004(b). *Literal Meaning*. New York: Cambridge University Press.
- Recanati, F. 2004(c). “‘What Is Said’ and the Semantics/Pragmatics Distinction.” in Bianchi, C. (ed.). 2004. *The Semantics/Pragmatics Distinction*. Stanford: CSLI Publication. 45 – 64.
- Sadock, J. 1978. “On Testing for Conversational Implicature.” in Cole, P. (ed.). 1978. *Syntax and Semantics, Vol. 9: Pragmatics*. New York: Academic Press. 281 – 297.
- Saul, J. 2012. “Just Go Ahead and Lie.” *Analysis* 72: 3 – 9.
- Saul, J. 2013. *Lying, Misleading, and the Role of What Is Said*. Oxford: Oxford University Press.
- Searle, J. 1969. *Speech Acts: An Essay in the Philosophy of Language*. London: Cambridge University Press.
- Searle, J. 1978. “Literal Meaning.” *Erkenntnis*, 13: 207 – 24.
- Simon, W. 2002. “The Belated Decline of Literalism in Professional Responsibility Doctrine: Soft Deception and the Rule of Law.” *Fordham Law Review* 70: 1881 – 1901.
- Simpson, D. 1992. “Lying, Liars, and Language.” *Philosophy and Phenomenological Research* 52: 623 – 639.
- Soames, S. 2008. “Drawing the Line Between Meaning and Implicature—and Relating Both to Assertion.” *Noûs* 42: 529 – 54.
- Soames, S. 2009(a). *Philosophical Essays, Volume 1: Natural Language: What It Means and How We Use It*. Princeton, NJ: Princeton University Press.
- Soames, S. 2009(b). “The Gap Between Meaning and Assertion: Why What We Literally Say Often Differs from What Our Words Literally Mean.” in Soames 2009(a), 278 – 297.
- Soames, S. 2010. *Philosophy of Language*. Princeton, NJ: Princeton University Press.
- Solan, L. 2002. “The Clinton Scandal: Some Legal Lessons from Linguistics.” in Cotterill, J. (ed.). 2002. *Language in the Legal Process*. Hampshire, UK: Palgrave MacMillan. 180 – 195.
- Sorenson, R. 2007. “Bald-Faced Lies! Lying Without the Intent to Deceive.” *Pacific Philosophical Quarterly* 88: 251 – 264.
- Sperber, D. and Wilson, D. 1986. *Relevance: Communication and Cognition*. Oxford: Blackwell.



- Sperber, D. and Wilson, D. 2004. "Relevance Theory." in Horn and Ward 2004, 607 – 632.
- Stalnaker, R. 1978. "Assertion." in Cole, P. (ed.). *Syntax and Semantics 9: Pragmatics*, 315-322.
- Stainton, R. 2004. "The Pragmatics of Non-Sentences." in Horn and Ward 2004, 266 – 287.
- Stainton, R. 2005. "In Defense of Non-Sentential Assertion." in Szabó 2005, 384 – 457.
- Stanley, J. 2000. "Context and Logical Form." *Linguistics and Philosophy*, 23: 391 – 424.
- Stanley, J. 2002. "Making It Articulated." *Mind and Language* 17: 149 – 168.
- Stanley, J. 2007. *Language in Context*. Oxford: Oxford University Press.
- Stanley, J. and Szabó, Z. 2000. "On Quantifier Domain Restriction." *Mind and Language*, 15: 219 – 61.
- Stokke, A. 2013(a). "Lying and Asserting." *The Journal of Philosophy* 110: 33 – 60.
- Stokke, A. 2013(b). "Lying, Deceiving, and Misleading." *Philosophy Compass* 8: 349 – 359.
- Strawson, P.F. "Intention and Convention in Speech Acts." *The Philosophical Review* 73: 439-460.
- Szabó, Z. G. (ed.). 2005. *Semantics vs. Pragmatics*. Oxford: Oxford University Press.
- Taylor, K. 2001. "Sex, Breakfast, and Descriptus Interruptus." *Synthese* 128: 45 – 61.
- Tiersma, P.M. 1990. "The Language of Perjury: 'Literal Truth', Ambiguity, and the False Statement Requirement." *Southern California Law Review* 63: 373 – 431.
- Travis, C. 1985. "On What Is Strictly Speaking True." *Canadian Journal of Philosophy*, 15: 187 – 229.
- Travis, C. 1997. "Pragmatics." in Hale, B. and Wright, C. (eds.). 1997. *A Companion to the Philosophy of Language*. Oxford: Basil Blackwell. 87 – 107.
- Unger, P. 1975. *Ignorance: A Case for Skepticism*. Oxford: Oxford University Press.
- Weiner, M. 2005. "Must We Know What We Say?" *The Philosophical Review* 114: 227 – 251.
- Williams, B. 2002. *Truth and Truthfulness: An Essay in Genealogy*. Princeton, NJ: Princeton University Press.
- Williamson, T. 1996. "Knowing and Asserting." *The Philosophical Review* 105: 489 – 523. Reprinted in Williamson 2000, Ch. 11 (page references to that book).
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

Wilson, M. 1982. "Predicate Meets Property." *The Philosophical Review* 91: 549 – 589.

Wilson, M. 2006. *Wandering Significance: An Essay on Conceptual Behavior*

Ziff, P. 1972. *Understanding Understanding*. Ithaca, NY: Cornell University Press.