

UCLA

UCLA Electronic Theses and Dissertations

Title

The Logic of Indexicals

Permalink

<https://escholarship.org/uc/item/6211x71b>

Author

Radulescu, Alexandru Viorel

Publication Date

2012

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

The Logic of Indexicals

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy
in Philosophy

by

Alexandru Viorel Radulescu

2012

© Copyright by
Alexandru Viorel Radulescu
2012

ABSTRACT OF THE DISSERTATION

The Logic of Indexicals

by

Alexandru Viorel Radulescu

Doctor of Philosophy in Philosophy

University of California, Los Angeles, 2012

Professor David Kaplan, Chair

The semantics of indexicals, at least in their most basic uses, was convincingly developed in David Kaplan's "Demonstratives", and his views are widely accepted today. Having proposed a formal semantics of indexicals, Kaplan also worked out a logic of indexicals. Most of the features of his logic come from the formal semantics; these I will accept, and even defend when prompted. But one feature goes against the grain of the semantics: arguments are forced to have all their steps within a single context of utterance. My central goal in this dissertation is to argue that this limitation is incorrect, and to provide a logic that does without it. This requires a new definition of validity, to take into account not just the logical form of sentences, but also certain abstract, relational properties of the sequence of contexts that the argument takes place in.

The logic I propose is designed for indexicals; traditionally, these are contrasted with demonstratives. I propose a new way to make this distinction, based on the observation that typical utterances have a speaker and an addressee, irrespective of the words

being uttered, and, most importantly, whether or not the words “I” or “you” are used. By contrast, demonstratives need personal assistance from the speaker: utterances don’t have demonstrata unless the sentences used contain demonstratives. I then argue that the traditional distinction, which is also inherited from “Demonstratives”, along with other criteria proposed in the literature, are mistaken. The main study case is the second person singular pronoun, which has been commonly, and wrongly, thought to be a demonstrative.

Finally, I discuss the ways in which the logic of indexicals differs conceptually both from classical first order logic, and from Kaplan’s logic. By looking not just at sentences and their parts, but also at the contexts, and the relations between them, we discover new and unexpected ways in which we can get valid arguments and logical truths.

The dissertation of Alexandru Viorel Radulescu is approved.

Sam Cumming

Andrew Hsu

Donald A. Martin

Yiannis Moschovakis

Terence Parsons

David Kaplan, Committee Chair

University of California, Los Angeles

2012

to my parents,
Violeta and Iancu Rădulescu

Contents

Introduction	1
1 A Logic of Indexicals with Context Change	4
1.1 Arguments against Any Logic of Indexicals	4
1.1.1 Strawson’s Argument	5
1.1.2 Soames’s Argument	9
1.2 A Simple Reason for a Logic of Indexicals	11
1.3 How to Allow Contexts to Vary, But Not Too Much	15
1.4 Similarity between Conversational Situations	23
1.4.1 Similarity and Conversational Participants	25
1.4.2 Similarity and Days	28
1.4.3 Similarity and Validity	29
1.4.4 The Deduction Theorem	30
1.5 Logical Truth	33
1.6 Why Not a Logic on Contents?	37
2 The Second Person Singular Pronoun and The Difference between Indexicals and Demonstratives	41
2.1 Introduction	41
2.2 The Difference, Traditionally	45
2.3 The Difference	53

2.3.1	Robust Context Features	54
2.3.2	Fragile Context Features	56
2.4	A Bit More about the Role of Intentions	65
2.5	From Utterances to Sentences in a Context	69
2.6	The Tradition and “You”	73
2.7	Other Criteria	78
2.7.1	Perry’s Narrow and Wide Contexts	78
2.7.2	Referent Change within a Context	81
2.7.3	Emptiness	87
2.8	What about Demonstratives?	91
3	Whence Logicality?	95
3.1	Introduction	95
3.2	A Worry to Put Aside: Logical Constants	95
3.3	A Challenge to LD	100
3.3.1	Kaplan Semantics, Quickly	101
3.3.2	The Methodological Argument	103
3.3.3	Logical Truths and Consequence	106
3.3.4	The Empirical Argument	111
3.4	Special Features of LI	114
3.4.1	Logical Truths and Valid Arguments in LI	115
3.4.2	LI and Analyticity	117
3.4.3	LI and Apriority	123
3.5	A Different Version of LI	125
3.6	Two Kinds of Logical Truth, Conversationally Speaking	128
3.6.1	What We Expect, and What We Assume	128
3.6.2	What We Assume about Ourselves as Participants	133

List of Examples	136
Bibliography	138

Acknowledgments

Philosophy dissertations are rarely written without the author incurring many debts of gratitude. I am no exception. Moves from one end of a continent to the opposite end of another are also rarely made without lots and lots of help. I certainly wouldn't have managed it alone. Put the two together and, even with all the overlap, I get a long list of people to thank, many of whom surely escape my memory just now. So I begin with a grateful acknowledgment to all those who ought to get an explicit one.

Thanks to Gabriel Sandoiu, who first taught me philosophy, and to Mircea Dumitru, both because his contribution to the state of the Philosophy Department in the University of Bucharest is so significant, and because he helped me leave when the time came.

Thanks to Laura and Sorin Bangu, Sorin Costreie, Sorin Constantinescu, Ruxandra and Gabriel Preda, who made moving to North America much more pleasant than it would have otherwise been.

Moving closer to UCLA itself, I really ought to thank the whole department, both as an institution with very good traditions, and as a group of faculty, graduate students, and staff. Special thanks are due to Barbara Herman, John Carriero, and Kristin Olson, who made these last two years an ideal time for dissertation writing.

And speaking of institutions, the UCLA Philosophy of Language Workshop deserves more gratitude than any other. It has come to feel like a philosophical home to me, the kind you take with you no matter where you go. Meeting almost every week, sometimes even during the summer, I cannot think of a better way for a student to confront philos-

ophy at its hardest, and at its most fun. For the first couple of years, I didn't see how I could really be a part of it; the discussion was clearly over my head, and I now wish I had a recorder at the time. I kept going, and that was one of my better choices. I still don't have a recorder, and I'm sure I'll come to regret that too. I have presented in the Workshop many of the ideas that ended up becoming my dissertation, along with many would-be parts, which thankfully did not. I thank all those who ever came to a Workshop meeting in the seven years I was at UCLA; I hope it continues for many more.

I have also presented parts of this dissertation at UCLA (in the Workshop, in the job seminar, and at several meetings of the Albritton Society), at the 3rd Semantic Content Workshop, Barcelona, Spain, at the 11th International Workshop on Semantics, Pragmatics, and Rhetoric, in Donostia (or San Sebastian, if you prefer), Spain, at the University of Rochester, at the University of Illinois, Chicago, at the APA – Central 2012, Chicago, and at the APA – Pacific 2012, Seattle. I thank all these audiences, especially my APA commentators, Lynsey Wolter and Isidora Stojanovic. Both this dissertation and I have also been helped by conversations with Geoff Georgi, Eliot Michaelson, and Scott Soames.

Very grateful thanks are due to Sam Cumming, Tony Martin, and Terry Parsons, all members of my committee, and all regulars of the Philosophy of Language Workshop. I learned a great deal from each of them. My dissertation would have been poorer if I hadn't listened to their advice, and much better if I had listened to it more. Thanks also to Andrew Hsu, for his friendship, and for the grace with which he listened to some of my ideas in a very rough form, and polished them to a much nicer form.

Philosophically speaking, I have benefitted immensely from working with my advisor, David Kaplan; certainly more than I could have imagined one could hope for. I learned to read, think, and write from him more than anyone else, and I learned to then do it again, and again, right up to the deadline. And then think some more. We have sometimes met for more than six hours at a time, without a break, and tired though I was by the end, as soon as I left, I wanted to go again. It's not that he makes hard philosophy look easy;

it's that he makes it look joyous. Most sentences in this dissertation were, to some extent, co-written, co-edited, or co-produced by him; think of this sentence as standing for a footnote placed at the end of each paragraph, thanking him for his help. His comments on the many drafts this dissertation went through were sometimes longer than the text itself, and usually better than it. And almost always in red ink. Purely for legibility reasons, he says.

Finally, I would like to thank Radiana Folescu for everything.

I dedicate this dissertation to my parents, and I thank them for their unwavering support, from near and from afar.

Vita

Education

- M.A. Philosophy, University of California, Los Angeles, 2008
- B.A. Philosophy, University of Bucharest, Romania, 2002

Presentations

- “The Logic of Indexicals”
American Philosophical Association – Central, Chicago, Feb 2012
The 11th International Workshop on Semantics, Pragmatics, and Rhetoric,
Spain, Nov 2011
The 3rd Semantic Content Workshop, Barcelona, Spain, Nov 2011
- “The Difference between Indexicals and Demonstratives. A Case Study”
American Philosophical Association – Pacific, Seattle, Apr 2012
Albritton Society, UCLA, May 2011

Fellowships, Honors, and Awards

- UCLA Dissertation Year Fellowship; 2011–2012
- UCLA Dean’s Humanities Fellowship; 2005–2011
- UCLA Chancellor’s Prize Summer Award; 2006, 2007
- Department of Philosophy Valedictorian, University of Bucharest, Romania; 2002

Introduction

Indexicals are words like “I”, “you”, “today”, and “tomorrow”. My dissertation is about their semantics and their logical properties.

The semantics of indexicals, at least in their most basic uses, has been convincingly developed in David Kaplan’s sibling papers, “Demonstratives” and “Afterthoughts”. Having proposed a formal semantics of indexicals, Kaplan also developed a logic of indexicals, somewhat misleadingly called “LD” (as in “logic of demonstratives”). Some of the features of his logic come from the formal semantics; these I will accept, and even defend when prompted. But some features go against the grain of the semantics: arguments are forced to have all their steps within a single context of utterance. My main goals in this dissertation are these: to argue that this last feature is incorrect, to provide a logic that does without it, and then to look at several startling implications of this new logic to our notion of an argument, of validity, and of what an indexical is.

The restriction of arguments to a single context is not without reasons. Ch. 1, therefore, begins with Scott Soames’s claim that this is what must be done in order to construct a proper logic of indexicals. The heart of his argument is that if contexts are allowed to change mid-argument, there would be no guarantee that two occurrences of an indexical will have the same semantic value. This, in turn, would make it that even the most basic logical laws, like Modus Ponens, would turn out to be invalid. On the other hand, I give examples of arguments which we intuitively judge valid, and which not only allow, but depend on each step in the argument having its own, distinct context. In the rest of

the chapter, I develop a logic of indexicals, “LI”, which does allow contexts to change within an argument. Soames’s mistake, shared by the whole tradition that follows LD, is to leave aside the possibility that abstract relations between contexts (such as one being exactly one day later than another, or two contexts having the speaker and the addressee switch roles) can be captured formally, so that validity is defined partly as a function of such features, rather than only on the basis of the sentences used.

LI, as its name suggests, is designed for indexicals. There is another class of context-sensitive expressions included in the vocabulary of LD, namely demonstratives, words like “this” and “that”. Traditionally, the distinction between indexicals and demonstratives has been built on the observation that demonstratives get their semantic value from the context of utterance only with the help of the speaker’s intentions, whereas indexicals can do it automatically (e.g. “I” picks out the speaker of the context, irrespective of the speaker’s intentions to refer to, say, David Hume, which will be thwarted for anyone but David Hume himself). In Ch. 2, I present a new way to distinguish these two kinds of words. The most important case for me is “you”, the second person singular pronoun. It is the most important, first, because “you” is a central part of the vocabulary of LI, which is only designed for indexicals, and second, because it points precisely at the place where the tradition went wrong. I argue that, even though the speaker’s intentions play a role in determining whom she is addressing, “you” is an indexical. The reason is that typical utterances have an addressee, just as much as they have a speaker, irrespective of the words being uttered, and, most importantly, whether or not the words “I” or “you” are used. By contrast, demonstratives need personal assistance from the speaker: utterances don’t have demonstrata unless the sentences used contain demonstratives. The intentions needed by demonstratives, then, are targeted at those particular uses, or tokens, of demonstratives. I find this new distinction fairly easy to grasp, but spelling it out requires a number of delicate distinctions, which take up much of the chapter. Then, I survey a number of other criteria proposed in the literature, with the conclusion that

most of them fail to do the job, while one of them can be developed in a manner friendly to my proposal. In the final section, I make some tentative proposals about a future logic of demonstratives with context changes: it will be closer to a logic of actual utterances than indexicals require.

In Ch. 3, I discuss ways in which LI differs conceptually both from classical first order logic, and from LD, its closest ancestor. The most central difference is the fact that in LI, an argument is made up of sentence-context pairs, not just sentences. In the first part of the chapter, I argue that this feature of LI has its beginnings in a disputed part of LD: “I am here now” is a logical truth, that is, is true in every context, because there is a requirement on all contexts that the speaker of each is at the place, time, and world of the utterance. I first defend this from an objection made by Stefano Predelli, and then draw the lesson that thinking about essential features of typical contexts of utterance should be part of what motivates any logic of indexicals. Then, I point out the ways in which LI is forced to depart from LD, a conceptual distance that explains why Soames thought impossible any logic that allows context changes. The end of the chapter is dedicated to ways in which we can be said to know an argument. This issue is fairly simple for regular arguments, since it seems to involve only linguistic knowledge. Since contexts are part of the argument in LI, different ways in which we are epistemically related to various relations between contexts will generate different intuitions about what to count as valid. This will explain away the feeling of uneasiness that comes with LI declaring some truths to be valid, even though they seem intuitively quite far from it.

Chapter 1

A Logic of Indexicals with Context Change

1.1 Arguments against Any Logic of Indexicals

Here is a valid argument:

- (1) John is happy. Therefore, John is happy.

It is quite easy to see that the conclusion cannot be false if the premise is true: they are the same sentence, and hence they both say the same thing. The more technically minded will prefer to say this in terms of models, or possible worlds, or what not. But intuitively, (1) is obviously valid.^{1, 2}

If we try a bit, we can make things seem less obvious. We can worry about the apparent contextual sensitivity of “happy”: maybe John is melancholic, and in the premise we’re saying that he’s happy by his standards, but in the conclusion we’ve moved on to

¹In my dissertation, I will be talking quite extensively about Kaplan’s work on indexicals. Since the two central papers were published in the same book, but written at least 12 years apart, I have found it useful to refer to them in a more descriptive manner than the usual citation conventions. I will use “[Demons]” for Kaplan (1989b) and “[Aft]” for Kaplan (1989a). We are talking loosely here.

²See [Demons], Remark XIX for a stricter presentation, which is also relevant to our topic.

judging him by more common standards, and then he won't count as really happy. Or we can start to worry about the fact that "John" is a name shared by many people, and for validity we need the name to stay constant throughout the argument. Or, getting closer to indexicals, we could worry about the tenses in the two sentences. As logicians, we usually put these worries aside, and leave them for later treatment. This is not to say that the worries should not be addressed at some point. But these are just topics for further work; they don't seem to present a challenge to the idea that "John" and "is happy" are the kinds of things that a logician can deal with.

Let's look at indexicals now. Consider this argument:

(2) Today is my birthday. Therefore, today is my birthday.

Is this a valid argument? Can we even have a logic of indexicals? A first difficulty is this: if we're just looking at the two sentences, it's hard to tell what it even means to say that they form a *valid* argument. "Today" only gets its value when placed in a context, and so does the first person pronoun. If we focus just on the sentences, we don't get truth values; without truth values, we cannot discuss validity, and hence logic has no grip.

This shows that we must look not only at sentences, but at sentences in a context of utterance, so that, e.g. "today" doesn't just float around, and instead its semantic value is determined by certain facts about the world at the time of utterance, i.e. what day it is. Logic, then, is not about sentences alone, but about sentences in a context.³ Would that be sufficient? Strawson thought not, and it will be instructive to see why.

1.1.1 Strawson's Argument

In order to understand Strawson's position, we need to look at two things: his views on the scope of logic, and his views about arguments which involve indexicals. This aside

³I want to avoid issues about what exactly worlds are, since this will be not important for the main issues I will discuss. Worlds might be thought of as logically possible worlds or even as models. In particular, they are allowed to vary from the actual world both metaphysically and with respect to the intensions of predicates and the denotations of names.

will take some effort, but, apart from the intrinsic historical interest, I hope that the fairly large passages we will look at will yield an important lesson. My ultimate position will be that Strawson was very nearly right, except that he drew the wrong conclusion because of his assumption that logic cannot take contexts into account.⁴

First, the scope of logic, according to Strawson, determines its central role in a theory of meaning:

To know the meaning of a [declarative] sentence [...] is to know under what conditions someone who used it would be making a true statement; to *explain* the meaning is to *say* what these conditions are. One way of giving a partial account of these conditions is to say what some of the entailments of the sentence are. For to say that one sentence entails another is to say that a statement made by the use of the first is true only if the corresponding statement made by the use of the second is true; and to say that one sentence entails and is entailed by another is to say that a statement made by the use of the first is true if, and only if, the corresponding statement made by the use of the second is true. This might make us think that to give the two-way entailments (the logical equivalents) of a statement-making sentence is all that can be done, in the way of *saying*, to give its meaning. [...]

To think this is a mistake, though a common one. Let us return to the point that to explain the meaning of a statement-making sentence is to say under what conditions someone who used it would be making a true statement: and let us call this 'giving the rules of use' of the sentence. We have just noticed the temptation to think that the only kind of rules involved are entailment-rules. I want to show, first, that this view is false, and, second, that the fact of its falsity imposes an unavoidable limitation on the scope and application of formal logic. The limitation in question is not one to be deplored or to be welcomed. It is one to be noticed; for the failure to notice it leads to logical mythology. (Strawson (1960, pp. 211-212); his emphases.)

A brief summary so far: part of the job of a theory of meaning is to say of each sentence what other sentence it entails. To do the latter is to say if the following conditional holds:

⁴Quine (1982, p. 56) agreed with Strawson. He thought that as far as logic was concerned, indexicals ought to be treated as ambiguous words: just fix their referent, and everything is fine. In other words, all we need to do with indexicals is make them into something they are not: very promiscuous proper names. Quine's reasons are not identical with Strawson's, but they are also less interesting for my purposes, so I put them aside.

if the first sentence is used truly, the second one, if used, would also be true.⁵ This is the business of logic. In case this doesn't come out clearly of the quote above, he makes the point explicit just a couple of pages later: "Formal logic is concerned with the meanings of sentences only in so far as these can be given by entailment-rules".⁶ But, and this is Strawson's bigger point, this is only part of what we need from a theory of meaning; logic cannot do the whole job. And his reason concerns indexicals:

The same sentence in different mouths may be used to make one true, and one false, statement ('My cat is dead'); the same sentence in the same mouth at different times may be used to make one true, and one false, statement ('My cat is dead'); and so on. [...] Since this is so, the assertion that a sentence S entails a sentence S' cannot be generally taken to mean that if any statement made by the use of the first is true, then any statement made by the use of the second would be true. It must rather be taken to mean the following: If, at some time, at some place, in the mouth of some speaker, the utterance of S results or would result in a true statement, then the utterance of S' at that time, at that place, in the mouth of that speaker, would result in a true statement. [...] Entailment-rules, as the above schematic formulation shows, abstract from the time and place of the utterance and the identity of the utterer: so they cannot tell the whole story about the conditions under which a sentence is used to make a true statement, unless the sentence is one of which it is true that, if its utterance by anyone, at any time, at any place results in a true statement, then its utterance by anyone else, at any other time, at any other place, results in a true statement.

Entailment rules, then, must be supplemented by rules of another kind. We may call these 'referring rules'. Referring rules take account of what entailment rules abstract from, viz., the time and place of the utterance and the identity of the utterer. [For example:] the word 'I' is correctly used by a speaker to refer to himself. (Strawson (1960, pp. 212-213))

Strawson's observation is surely correct: indexicals have the special ability to get different semantic values when placed in different contexts. "I" picks out the speaker, whoever he or she turns out to be, the present tense in simple sentences picks out the moment

⁵I use a modal expression where Strawson did not; but he clearly did not mean that any two statements with no indexicals in them that actually have the same truth value are logically equivalent. The formulation in the next quote is better in this respect.

⁶Strawson (1960, p. 214).

of time of the context, etc. Indexicals, then, are not like “bucket”, which, in its most straightforward uses, just refers to buckets, or the property of being a bucket.⁷

Strawson’s central point is this: once we notice that indexicals are special, we notice that entailment rules do not capture everything about our intuitive notion of validity, and, therefore are not sufficient for a theory of meaning. To see this, take (2).⁸ Recall that logic is about entailment. Surprisingly, it turns out that (2) is not a valid argument. It is simply false that the truth of the statement made by one use of “Today is my birthday” guarantees the truth of the statement made by another use of that sentence; if the uses are, say, one week apart, and given our being born only once, their truth values will differ. In fact, Strawson argues, logic will only identify as valid those arguments which are made up of sentences whose truth value just doesn’t depend on contextual features. Contra Strawson, some of these will involve indexicals ($p \vee \neg p$ will be true no matter how many indexicals p contains), but most truths involving indexicals will fall outside the domain of logic, as Strawson conceives it.

What about (2)? Well, says Strawson, there is some potential there. As long as we keep certain facts fixed (same speaker, same time and place), we do get a valid argument. But keeping facts fixed is not something logic is about, nor could it even do it: you need referring rules for that, and those are outside the province of logic. Since you cannot keep the context fixed by logical means, logic and the entailment rules it studies are not sufficient to provide a theory of meaning for indexicals. For that, you need both entailment and referring rules.

⁷However, he is still not addressing modality properly. To put it in modern terms, he should add “same possible world” to the other sameness conditions. Just as the value of “I” changes with the speaker, the value of “actually” changes with the possible world. Even worse, possible worlds do double duty semantically, since, besides providing semantic values for indexicals, they also provide the world of the circumstance of evaluation; presumably, Strawson wants it to be kept constant throughout the argument, just like the time, the place, etc.

⁸(2) uses the same sentence twice, whereas Strawson states the more general case in which we use two different sentences, but his point applies just as well.

1.1.2 Soames's Argument

[Demons] showed that Strawson was wrong: we *can* have a logic of indexicals. The fundamental idea was that we need to capture characters (Strawson's referring rules) formally, and this systematic way in which indexicals get their semantic value from the context will allow logic to get a grip on them.⁹

Now we can go back to (2) and account for its validity. Suppose that the premise and the conclusion are placed in the same context, as Strawson thought we must. This guarantees that the semantic value of the two occurrences of "today" and of "my" are the same. This allows us not to worry about the actual truth values of the premise and the conclusion. Instead, we get something very similar to what we had for (1): for any model, for any context, if the premise of (2) is true in that context, so is the conclusion. Quantification over contexts is added, to take care of the special features of indexicals, but otherwise, everything is just as before.¹⁰

I have assumed that the context stays the same for the whole argument, as does [Demons]. Strawson also relied on the idea that keeping the context fixed must be what our intuitive notion of validity presupposes. But we haven't seen an argument to this effect. Can we just assume that this requirement must be met? Do we have to? Scott Soames argues that we do:

In actual conversation, it is, of course, possible for different speakers to utter

⁹As noted in Kaplan (1979, p. 85), Kaplan's characters correspond quite closely to Strawson's referring rules, and so do his contexts of utterance (a label that they share). They do differ, though, on the ultimate truth bearers: Kaplan takes contents of sentences to be propositions, which have a truth value relative to a circumstance of evaluation (which, in [Demons], is a pair of a possible world and a moment in time). Strawson's statements are more closely related to human action and sentences than propositions. His notion of a statement isn't quite the idea of an actual use of a sentence by a person, since he allows talk about the truth value of the statement someone could have made, but did not. Still, at least on Kaplan's account, propositions exist independently of language, whereas Strawson tends to speak of statements as what one makes in using a sentence. These differences are important, and it is not trivial to spell them out. Still, none of this obscures the point (later acknowledged by Strawson, according to Kaplan (p.c.)) that Strawson underestimated the ability of logic to deal with indexicals in exactly the way in which Kaplan did it.

¹⁰Quantification over contexts may be needed anyway, even for a language without indexicals; see [Aft], p. 595. All the better for the logic of indexicals.

different sentences, or even for one speaker to interrupt and finish another's sentence. It is also possible to start a discourse on one day, and finish it on the next. For those reasons, real discourses can contain multiple occurrences of pure indexicals with different referents/contents. This discrepancy between real-life speech situations and Kaplan's system should be regarded as an idealization. [...] [T]he point to notice here is that the idealization is all but forced on him by his goal of developing *a logic* for indexicals. For, it is natural to think, any putative logic in which the transition from A and 'A \supset B' to B sometimes failed to preserve truth, because different occurrences of the same expression in A or B received different interpretations, would scarcely count as a logic at all. (Soames (2010, p. 101-102))

Let's see how this argument works when applied to a simple practical syllogism:

(3) I'm hungry. If I'm hungry, I should eat. Therefore, I should eat.

Soames's point is this: if every step in the argument takes place in the same context, the argument is clearly valid, and a simple instance of Modus Ponens.

Suppose we allow contexts to vary, and, in particular, we allow the speaker to vary. Let me be the speaker of the first premise, and Captain Spock the speaker of both the second premise and the conclusion (who else would bother to go through an explicit argument in favor of eating?). Neither I nor Captain Spock should draw that conclusion from those premises; my hunger shouldn't affect his lunch arrangements many years from now (around 2266, to be precise). The first premise is about me, while the second is about him, and Soames is right in saying that (3) is not a valid argument.¹¹

But (3) sure *looks like* an instance of Modus Ponens. So, Soames seems to think, if we allow contexts to vary within an argument, we'll get exceptions to any and all reasoning patterns. A logic which allows exceptions to rules "would scarcely count as a logic at all", he continues. Hence the conclusion: we cannot allow contexts to vary within an argument.¹²

¹¹Soames is not alone in making this point, though his is the most explicit discussion I have seen. See also Quine (1982, p. 56) and Rumfitt (2010, p. 37–38).

¹²For reasons unrelated to indexicals, Soames in fact believes that there can be no logic of *English*, as

It looks like the choice is either to give up on the idea of a logic of indexicals, or force arguments to take place within a single context. But this would be Strawson's "I told you so!" moment. True, it turns out that we can make referring rules part of a formal system. But, in fact, in conversation we accept as valid more arguments than a logic that requires that the context be fixed seems to be able to account for. And, after all, Strawson never claimed that something resembling English indexicals in some respects could not be captured by a formal language. [Demons] presents an interesting system, he could agree, but the restriction to arguments in a fixed context is not a mere "idealization", as Soames calls it. Rather, it shows that [Demons] is just more formal work, which doesn't apply to the richness of "today" and "you".

1.2 A Simple Reason for a Logic of Indexicals

One underlying motif of Kaplan's work is the search for a characterization of the meaning of natural language terms in a formal framework, which naturally raises the possibility of a search for logical properties, like validity and consistency. This is a bold project, and it may yet fail. Have we found one such failure? Will indexicals ironically show that, because of the restriction to a fixed context, the logical part of the project does not work, as Strawson thought?

I will provide an answer to Soames's simple and powerful argument. But before that, I'd like to appeal to a simple intuition, which seems to be pulling in my direction. Let us look at this quote from Frege:

If a time-indication is conveyed by the present tense one must know when the sentence was uttered in order to grasp the thought correctly. Therefore the time of utterance is part of the expression of the thought. If someone wants

opposed to its formalized shadows, and his skepticism is part of a tradition of worries caused by *demonstratives*. For a brief presentation of Soames's views, see Soames (2010, §7.21). Other important papers about the possibility of a logic of demonstratives are Reimer (1992), Caplan (2003), Salmon (2002), and Braun (1996).

to say today what he expressed yesterday using the word 'today', he will replace this word with 'yesterday'. Although the thought is the same its verbal expression must be different in order that the change of sense which would otherwise be effected by the differing times of utterance may be canceled out. (Frege ([1918] 1997, p. 332))

This passage is puzzling for many reasons. It is unclear how the time of the utterance could be part of the expression of a thought. We usually think of expressing a thought as an activity which only involves signs, not features of the world like times and places. And it isn't clear why Frege thinks that the same thought is expressed when the method he proposes involves an interchange of obviously non-synonymous words ("yesterday" for "today"). All these are interesting and difficult questions both for Frege scholarship, and for any broadly Fregean theory of language and thought.¹³ But I'd like to focus on the simple, intuitive part of Frege's point. Consider these situations:

(4) Josh says "It's raining today". The next day, Josh says "It rained yesterday".

(5) Jerry says to Elaine: "If you want to go to the movies, we'll go to the movies." Elaine replies: "I do want to go to the movies."

(6) James Madison says to Napoleon: "I'm taller than you are". Napoleon replies (correctly) to Madison: "You're not taller than I am".¹⁴

(4) is just the kind of situation Frege is talking about. Josh starts out by saying something about the weather, and uses an indexical. The next day, to emphasize that his opinion about the weather had not changed, he would like to repeat what he had said. But he cannot use the same words; the tense and the indexical need to be changed. We'll put tense aside for now, and such issues as the fact that rain only happens in particular places,

¹³Kripke (2008, p. 204-205) argues that it is impossible to provide a consistent account of Frege's philosophy of language and allow him to claim that the same thought is expressed on the two days. See Evans ([1980] 1996) for a defense of Frege.

¹⁴Napoleon was 5'6.5" (1.69m); Madison was shorter: 5'4" (1.63m). I don't know if they ever met, but Madison was the Secretary of State of the US during the Louisiana Purchase, and Napoleon was the selling party.

so the sentences need to be evaluated relative to the same place as well; just focus on the indexicals. The important thing to notice is this: no matter who said those sentences, and no matter who they were addressing, if we keep fixed the fact that they were said one day apart, in that order, if the first sentence was true in its context, the second sentence *must* be true in its own context. This relation is remarkably like validity. It's a bit more complicated than garden-variety validity, because the argument only works given a certain relation between the contexts. But this is a technical problem; Frege's observation was simpler: Josh must change the indexicals, and if he does, what he says will reiterate what he said yesterday, and therefore Josh is right to do so. Logic shouldn't prevent him from doing that.

As we noted, Frege makes a bolder claim: that Josh expresses the same thought on both occasions. Whether or not that is true, surely he's right about one consequence of what he says. Recall that, for Frege, thoughts are the primary truth bearers; so logic operates on them. Since he claims that the same thought is expressed, (4) must somehow involve a valid argument. This is the part I'd like to focus on: there must be a valid argument lurking somewhere around (4).

(5) only works as a conversation because Elaine can expect her remark to connect with Jerry's in the right way: she does want to go to the movies, so, given what Jerry just said, if all goes well, they'll go to the movies. We have two premises just waiting to be put together into a modus ponens. Why is that? Because Jerry was addressing Elaine when he said "you", and Elaine replied by using "I". As long as the addressee of the first sentence is the speaker of the second, we have all we need to safely draw the conclusion in a valid argument. Since this condition is satisfied in (5), there is no danger in letting the contexts vary. For concreteness, suppose this condition was not satisfied, so that the first step remains as in (5), but instead of Elaine, Newman says to Jerry "I want to go to the movies".¹⁵ The two utterances just don't connect; Jerry was addressing

¹⁵The background, which is not necessary for the logical point, but adds meat to the example, is that, as

Elaine, not Newman, so Newman's reply is logically irrelevant, besides being impolite. No non-trivial conclusion follows from the two utterances, and the fault belongs not to the sentences, but to the fact that the participants did not match correctly: the speaker of the second utterance was not the addressee of the first.

A similar point can be made about (6): if Napoleon wants to contradict Madison, what can he do? He can't just say "I'm not taller than you are", since that would get the comparison mixed up. A natural choice seems to require changing the indexical from "I" to "you" and from "you" to "I". Again, the key is a relation between the two contexts: the speaker and the addressee of the first context are switched in the second. That's what allows Napoleon to *contradict* Madison, and still use indexicals.

This shows that people take full advantage of the logical properties of indexicals: (5) does not just happen to look like an off-beat version of the premises for a *modus ponens*; they work like it, logically speaking. So our intuitions that there are logical relations involved in these context-changing conversations are quite robust: they are about validity and contradictions, and involve the kind of changes that Frege was talking about. And these intuitions are about arguments (more exactly, sequences of utterances) in whose midst the context changes. This is my simple reason for the search for a logic of indexicals: it's not that this would validate a particular view about semantics (though some of us would like that), but rather that we use indexicals as if our uses had logical properties, as if a person in one context can draw inferences from what others have said in other contexts, for example, by affirming or contradicting what others have said, or by supplying additional premises for an obvious inference.

Of course, Soames is right in saying that if we allow any context change, we're in trouble. If Josh misremembers the day when he had originally said it was raining, and says the second sentence two days later, he is no longer just repeating what he had said.

every fan of *Seinfeld* knows, Jerry really doesn't like Newman; he has no interest in Newman's desires, and he certainly doesn't want to go to the movies with him.

The second claim is now independent of the first. And if Napoleon repeats his remark to Thomas Jefferson (who was 6'2", or 1.89m), he is no longer correct, and, more importantly, he is no longer contradicting Madison. The challenge, then, and my goal in this chapter, is to construct a logic that validates my claims about (4), (5), and (6), while invalidating their intuitively invalid counterparts.

This is something that [Demons] cannot handle, since that logic does not allow the context to change within an argument. But the semantics for indexicals that Kaplan proposes seems adequate to a context-changing sequence of utterances: it nicely explains why Josh must change "today" into "yesterday", Elaine must replace "you" for "I", and so on. Kaplan's semantics did not get the logic it deserved. I want to remedy this situation, by offering a Logic of Indexicals (henceforth, "LI"), to be contrasted with Kaplan's Logic of Demonstratives ("LD").¹⁶

1.3 How to Allow Contexts to Vary, But Not Too Much

First, we must answer Soames's argument, which claims to show that no logic of indexicals can account for situations like (5), which require changes in context. Modus Ponens (just like any other logical rule) is obviously safe if we disallow context changes within an argument. All occurrences of "I" and "today", for instance, will corefer, because the speaker and the time will be guaranteed to be the same. So Modus Ponens will have the simple form we are accustomed to: if one premise is A and the other $\lceil A \supset B \rceil$, we can safely infer B , no matter what indexicals they contain, if any. In fact, for logical purposes we don't care about the contexts themselves; we only care that there is the one context, which does its job of supplying values for the indexicals. For discerning validity, we don't look at the context at all; instead, we need only look at the general properties of indexicals,

¹⁶Both are misnomers, since Kaplan's logic deals with both indexicals and demonstratives, and I would like to extend mine to do the same thing, though we will see in Ch. 2 that this is not a trivial matter. Anyway, "LD" is already in use, and since here I will only talk about indexicals, the names should not cause too much confusion.

just as for all other parts of the language.

Soames's error is to assume that once we allow contexts to vary, we cannot put any constraints on their new found freedom. His worry is that indexicals could get any old values, and that no argument with an indexical in it will ever count as valid. This is a mistake. We saw that Elaine concludes that she and Jerry will go to the movies, because in any context like the one that we described, where the speaker and the addressee switched roles in the two premises, it logically follows from those premises. Once we constrain context variation in certain ways, we can eliminate the cases that Soames worries about, and recover a working notion of validity. The technical problem is to figure out how to characterize the constraints.

In fact, we should have seen this all along. An early sign of trouble is failing to pay attention to a central part of Kaplan's answer to Strawson: we shouldn't just look at sentences, but at sentences in a context. Once we follow [Demons], we can formally specify the behavior of indexicals as dependent on contexts. Admittedly, his system does have the requirement that the context not vary. But the idea that there's more to an argument than the sequence of its sentences is a lesson not yet fully exploited.

We haven't said what contexts are yet. The intuitive idea, derived from [Demons], is simple enough: we want to capture the fact that utterances take place embedded in a world, so that they involve a speaker at a given time, at a given location, with a certain addressee. Contexts are thus part of the environment in which an utterance usually takes place, and their role is to fill out the content of whatever indexicals may occur in the sentence uttered.¹⁷

Unlike in [Demons], in a normal conversation, the participants take turns speaking, and the discussion may take a considerable period of time. I like to think that, when doing philosophy, we're part of the same conversation as the one started by Plato, or Thales.

¹⁷In this paragraph, I have in mind a fairly robust notion of possible world: not just as a model, or as a logically possible world, but as a way the actual world could have been. The notion of a possible world as a model I take to be an idealization, made for theoretical purposes, of the robust notion.

But we don't need to get so fanciful. Even though most conversations take considerably less time, they are not instantaneous, and they usually involve more than one speaker. Conversations should be allowed to flow freely, moving from participant to participant, and sometimes staying up all night and into the next day. In fact, it is quite amazing how much work in semantics has been on the semantics of monologues.¹⁸

My proposal is to take seriously this banal feature of conversations: think of an argument as a special type of conversation, in which the last remark is taken as the conclusion and the earlier remarks are taken as premises.¹⁹ In some cases, this conclusion will be a logical consequence of the premises, in some cases not. More formally, think of an argument as a sequence of context-sentence pairs, the last of which is the conclusion.²⁰ For example, (2), putatively an example of the logical rule of Repetition, is of this form: $\langle [c_1, \phi], [c_2, \phi] \rangle$. Like examples (4) – (6), this argument is valid for some context sequences, most obviously where $c_1 = c_2$, and invalid for others. This new notion of argument is captured by this characterization:

- Think of an *argument* as a sequence of the form $\langle [c_1, \phi_1], [c_2, \phi_2], \dots [c_n, \phi_n] \rangle$, where $n \geq 1$, and each pair $[c_i, \phi_i]$ is a pair of a context and a sentence.²¹

For such an argument we can abstract two important notions. One is the older notion of an argument, as just a sequence of sentences, which I will call a “conversational thread”.

¹⁸Which is not to deny that a lot of good work has recently been done on the semantics and pragmatics of conversations; see, for instance, the SDRT tradition, starting from Asher (2000).

¹⁹I will assume that premises are ordered in an argument. This is not unnatural, since most of the time we do utter sentences one after the other. But we also want to allow that several sentences can occur in the same context. So I will let the ordering be arbitrary. Historically, any ordering is unusual: most of the time, premises are allowed to form a simple set, without an ordering. My choice will help my presentation, and it can be proven that no harm comes from it.

²⁰Rumfitt (2010, p. 36–37) also thinks that steps in an argument are context-sentence pairs, but he does not allow contexts to vary.

²¹Note that I'm reserving square brackets for context-sentence pairs. We will impose more constraints later; for instance, I will require that the world feature stay constant throughout an argument.

Definition 1. Given an argument $\langle [c_1, \phi_1], [c_2, \phi_2], \dots [c_n, \phi_n] \rangle$, the sequence of its sentences $\langle \phi_1, \phi_2, \dots \phi_n \rangle$ is its *conversational thread*.

This notion is useful when we look for a notion of validity, which requires that we take the original thread, and look at what happens when we place it in other sequences of contexts. So we need sequences of contexts too, which I call “conversational situations”, to capture the idea that they are like a context for an utterance: they are the environment in which arguments take place.²²

Definition 2. Given an argument $\langle [c_1, \phi_1], [c_2, \phi_2], \dots [c_n, \phi_n] \rangle$, the sequence of its contexts

$\langle c_1, c_2, \dots c_n \rangle$ is its *conversational situation*.

I talk of conversations because that’s where arguments reside. This has also a contrastive purpose: LD, due to its insistence on sameness of context, is more like a logic of instantaneous monologues. I want to allow arguments to be more like normal exchanges: they take place over time, at varying locations, and often contain several speakers. Hence, “conversations”.

We will soon see some technical uses for the distinction between conversational situations and threads. But I hope they are natural notions: conversations can be alike because the same sentences are uttered, even if the speakers are different. In cases like this, I’ll just say that the conversational threads are identical.

A final preparatory definition: the notion of context. For formal purposes, it is best to represent contexts as sequences of parameters, one for each kind of indexical. For ease of exposition, I occasionally slip into talking of contexts as if they just are sequences of parameters. The slip is innocuous, since my goal is to show that we can represent context change with formal means. But behind the scenes, contexts are just what they are: parts

²²“Situation” has another use, stemming from Barwise & Perry (1983). There, it means something like “a small part of a possible world”. I won’t be addressing their theory here, so there should be no danger of ambiguity.

of the world (including possible contexts, which are parts of possible states of the world), not abstract sequences.

In this chapter, I focus on just a few indexicals: “I”, “you”, “today”, “yesterday”, and “tomorrow”. So my contexts are impoverished versions of the contexts of [Demons], which also contained locations. I only need a speaker, an addressee, a day, and, for reasons that we’ll discuss later, a possible state of the world. I also don’t account for the fact that sometimes we address several people at the same time, so we will never have more than one individual playing the role of addressee. In future developments, this lack of pluralities should be dealt with.²³

- Think of a *context* as (S, A, D, W) , where S is the speaker, A is the addressee, D is the day (represented by a positive integer), and W is the world.²⁴

Let me express my response to Soames’s argument in these newly introduced terms. I begin with a more detailed version of (4), using the notations just defined:

(7) $\langle [c_1, \text{“It’s raining today”}], [c_2, \text{“It rained yesterday”}] \rangle$

Let’s assume that the two contexts are alike with one exception: the day of c_2 is exactly one day later than the day of c_1 . Contexts, remember, specify the relevant features of the world that matter for indexicals. So let’s fix on certain specific contexts. Let $c_1 =$ (Josh, Kent, Sep 21, 2011, the actual world) and $c_2 =$ (Josh, Kent, Sep 22, 2011, the actual world).^{25,26} In order to say that this is a valid argument, we need a certain amount of

²³[Demons] does not have an addressee as a feature of the context, since it does not have “you” in the vocabulary. It is not trivial to claim that “you” is an indexical; but I will argue that it is in Ch. 2.

²⁴Note the use of round parentheses. As in the case of steps in an argument, the style of parentheses is relevant only for readability. This is not a definition, because later we will impose more constraints on what counts as a good representation of a context, stemming from constraints on what it takes to be a context.

²⁵Josh Rubenstein is the chief meteorologist and Kent Shockneck is a fellow presenter at KCAL 9, a CBS-owned network in Los Angeles. They both work on the 11 am news.

²⁶This formulation is a bit sloppy, since days are represented in contexts by integers, and Sep 22, 2011 is not an integer. I do this for readability: strictly speaking, the day should be represented, say, by 1871994, the number of days since the latest creation according to the Mayan calendar.

generalization. After all, we cannot simply look at the truth values; in our case, both the premise and the conclusion are false, which leaves open the question whether the conclusion *must* be true if the premise is true.

Our intuitions about (7) are simple. It is valid because:

For any conversational situation $\langle c'_1, c'_2 \rangle$ that is similar to $\langle c_1, c_2 \rangle$ in that they have the same speaker, the same addressee, the same world, and the day of c'_2 is one day later than the day of c'_1 , if the premise is true at c'_1 , then the conclusion is true at c'_2 .

Using the notions defined above, we could put this by saying that all similar conversational situations are truth-preserving for the given conversational thread. You might wonder why we only look at conversational situations that are *similar* in how *all* the context parameters relate rather than only requiring similarity with respect to days. Applied to our example, the worry is this: it would still have rained the previous day even if Josh couldn't come to work on the second day, and someone else had taken his place. All that matters is that the argument take place on successive days. That is correct. I have two answers. One is that this limitation – requiring similarity with respect to the relations between all context parameters – makes things easier from a technical point of view, and the second is that (7) would also be valid if we made $c_2 = (\text{Fred, Marie, Sept 22, 2011, the actual world})$. In looking at conversational situations similar to this, we would allow free variation for the speaker and the addressee parameters.

The most important lesson is that we are looking for conversational situation types in which the contexts stand in certain abstract relations: same speaker, or same date, or a switch of speaker and addressee, etc. To determine whether an argument is valid, we start from its own conversational situation. Then, we look for conversational situations which are *similar* to it in terms of these abstract relations, since similarity is a kind of isomorphism of conversational situations. Once we do this, we can answer Soames's challenge: contexts can vary within an argument, and we can still have something that looks like

logic, with rules and generality. To do this, we must formally characterize the abstract relations between contexts that are intuitively relevant for validity. The question is: what does it mean to say that the relations between the features of an arbitrary conversational situation generate a class of conversational situations? The answer will be given in terms of isomorphism and similarity between conversational situations.

This idea is a weaker version of the constraint in [Demons], that required that the context stay the same, and then generalized over all conversational situations which kept the same context throughout. Since this generalization over sequences of contexts was the same for any argument, it didn't deserve any special attention. This is why LD is a sublogic of LI for instantaneous monologues. Once we try to liberalize the conversational situations we wish to address, the challenge is to characterize a notion of similarity which gets things just right: it must not be too strict to miss intuitively valid arguments, nor too loose, to rule out intuitively invalid arguments. We must not only avoid the problems Soames mentions, but also explain why some arguments with context change just seem valid.

To see how this works, suppose we're interested in a given argument, in my sense, made up of sentence - context pairs. On the one hand, if we impose strict conditions for two conversational situations to count as similar, we get too many validities, because we don't get the generality we need. At one extreme, we could only look at that particular conversational situation. Clearly, this is no notion of validity, since we would have no generality; all we could tell would be whether the premises are in fact false or the conclusion is in fact true. Take (7). It didn't rain on either day in Los Angeles; but that doesn't tell us that the argument is valid.²⁷

²⁷In Georgi (2011), and for different purposes, Geoff Georgi proposes that we take the context we begin with, and look at that very context in different models, i.e. with different interpretations of the non-logical vocabulary. Depending on how that idea is applied to the issues I am concerned with, and how the logic is developed, it might turn out to be equivalent to my proposals. However, keeping the context fixed is conceptually backwards; what makes indexicals special is their systematic reliance on contexts, and that is lost if we just keep the context fixed. There may also turn out to be some special problems, since we would only look at models which contain the objects in the original context; if they have any necessary properties,

In fact, if we insist on keeping any element of the context exactly the same, we can generate unwanted validities. For instance, the agent could be required to stay the same in all similar conversational situations; but what is it about that particular agent that matters for validity? Validity should be independent of the identity of the speaker. Otherwise, we would have “I am Margaret Thatcher” be a logical truth if uttered by Margaret Thatcher. What is then said may be a necessary truth, but it shouldn’t be a logical truth, and it is not in LI. To avoid such unwanted results, we must allow all elements of the context to vary.

On the other hand, if we make it too easy for two conversational situations to count as similar, we don’t get enough validities, and we’re saddled with Soames’s problem. Thus, if we count two pairs of contexts as similar just in case the speakers either are the same in each pair, or are different in each pair, we would be missing out on describing (6) as a contradiction. Recall that Madison says to Napoleon “I’m taller than you”, and Napoleon wants to contradict him, and says “You’re not taller than I am”. What is relevant there is not merely the fact that the speakers are different; what matters is that Madison is addressing Napoleon, and then Napoleon is replying to Madison. If we don’t take into account the addressees of the contexts, we would count as similar the context we talked about earlier, where Napoleon is addressing the much taller Thomas Jefferson. And in that case, both sentences are false, and Napoleon is obviously not contradicting Madison.

If this sounds too technical, here’s an even starker example. Suppose we require for similarity that familial relations between the speaker and the addressee be preserved. This would give us the wrong account of obviously invalid arguments. Here’s one example: “I am older than you. Therefore, I am your uncle”, where the conclusion is uttered truthfully. Then, since it is required that genealogical relations be preserved, the only similar conversational situations would be ones in which the speaker of the conclusion is the

that might generate unwanted validities. The relation between modality and ways of understanding logical validity is complicated, as noted in Kaplan (1986, Appendix E), and it gets even more complicated when indexicals come up. In any case, talk of models is technical, and possibly technically problematic; talk of contexts is intuitive, and, I will try to show, technically unproblematic.

uncle of the addressee. Since the conclusion could not fail to be true in all such ‘similar’ conversational situations, the argument would be valid. This is undesirable both because there is no relation between the conclusion and the premise, and because being someone’s uncle, just like any other such properties, is commonly thought to be a non-logical property.

The moral is that, from all the features that contexts may share, we must carefully pick only the ones which matter for the relevant sense of validity.

1.4 Similarity between Conversational Situations

Intuitively, the validity of an argument with indexicals in it depends on what happens in contexts which are similar to the ones in which the argument takes place. The conversation between Jerry and Elaine, for instance, is a Modus Ponens waiting to happen only because Jerry addressed Elaine, and Elaine replied. In any pair of contexts such that in the first, the speaker addresses the speaker of the second premise, the contents match in the right way. For validity, we don’t really care who did the speaking and when; we only care whether certain relations between contextual features obtain, and then we look at similar context sequences, and see what happens in those contexts with the truth values of the sentences in the argument. If we get the right link between the truth of the premises and the truth of the conclusion, we have a valid argument; if not, we don’t.

The way to formally capture this intuitive notion is by working on isomorphic relations between context sequences. As we saw in the previous section, the challenge is for this isomorphism to capture all and only those relations between the contexts of the original conversational situation which are intuitively relevant for validity.

Since in this chapter I’m only working with just a few indexicals (“I”, “you”, “yesterday”, “today”, and “tomorrow”), the notion of similarity between conversational situations relevant for validity will have to look at relations between the respective compo-

nents of contexts: speakers, addressees, and days.

I didn't mention worlds just now, even though they are parts of contexts. Here I will follow Kaplan's idea, and require that worlds stay constant throughout an argument. The main motivation for this restriction is that I'm trying to capture something about real arguments, and these take place in just one world. Just as there are no trans-world conversations, there also are no trans-world arguments. Our Lewisian counterparts, if there are any, cannot address us, and we cannot address them. What we can do is think about merely possible conversations: remembering a past conversation, I can think that if I had said q instead of p , I would have won the argument. What we cannot do is engage in such conversations. Since this is the logical situation with arguments, it needs to be captured in logic, and the effect is our definition of propriety.

Still following [Demons], I will require that contexts be proper, i.e. that both the speaker and the addressee exist at the world of the context on the day of the context. This last requirement is just a temporary simplification. There are cases where we mean to address someone, and there's nobody there.²⁸ This is similar to the case of empty names, and, as is customary for this latter problem, I will leave it for further work. Let me say this formally:

Definition 3. A conversational situation

$$\langle (S_1, A_1, D_1, W_1), (S_2, A_2, D_2, W_2), \dots (S_n, A_n, D_n, W_n) \rangle$$

is *proper* iff $(\forall i, j \in [1, n]) W_i = W_j \wedge (S_i \text{ and } A_i \text{ exist at } W_i \text{ on } D_i)$.

This is the kind of argument that drove Kaplan to say that only some contexts are proper: the ones where the speaker is at the place of the context, at the time of the context. Our formal representations of contexts are smaller: they don't contain locations, which we don't need because we don't have location indexicals in our language. And the main gist of this chapter is that, unlike in [Demons], we need to take special care of

²⁸We can, for instance, address a note to an unborn child; or unknowingly address a hallucination. So the addressee ideally should not be guaranteed to exist at the time of the context, or, indeed, ever.

relations and transformations among contexts. But the intuition is the same: arguments, and the contexts in which they occur, are worldly things, parts of a particular world, and our initial concerns are just with them. The difference is just that the kind of moving arguments we are focusing on also require a new form of generalization over conversational situations.

Doubts have been expressed about the constraints imposed on contexts in [Demons], and I won't address them here.²⁹ Some concern answering machine messages, which say "I'm not here now" and seem true just in case the speaker is *not* at the location of the utterance, contrary to Kaplan's original properness condition. These worries don't easily apply to my new, much weaker constraint that the world be constant; after all, there are no inter-worldly answering machines. Still, since I preserve much from [Demons], it should come as no surprise that LI also faces the answering machine problem. Think of a text in a testament: "By the time you read this, I will no longer exist". I take it that, for logical purposes, the utterance takes place not when the dying person writes his will; that is just the context of recording. The context in which we judge the truth value of the sentence is the one where the will is read by the lawyer. I have required that the speaker exist at the context of utterance, but here we have two contexts, and the speaker no longer exists in one of them. My intuition is that a special semantics for recordings is needed, which will differ from the semantics of the more typical uses of language. I leave this issue aside in this chapter.

1.4.1 Similarity and Conversational Participants

Now that we have seen that we need to keep the *world* component of the contexts in a conversational situation fixed, let's move to the conversation participants: the speaker and the addressee. Let's begin with *prima facie* cases of Repetition:

²⁹A few early papers are Smith (1989), Predelli (1998), and Corazza et al. (2002).

(8) $\langle [c_1, \text{"I am hungry"}], [c_2, \text{"I am hungry"}] \rangle$

(9) $\langle [c_3, \text{"You are hungry"}], [c_4, \text{"You are hungry"}] \rangle$

Intuitively speaking, (8) is valid only when the speaker stays the same throughout, and (9) only when the addressee stays the same (these conditions are not sufficient, obviously, since the time of utterance matters too; but for now we're just focusing on the conversation participants). This suggests that two conversational situations are similar only if they are alike in the following respect: if there are two contexts in one conversational situation such that they have the same speaker (or addressee), the corresponding pair of contexts in the other conversational situation will have the same speaker (or addressee). Note that this does not require that the four contexts have a person in common; all that is required is a kind of isomorphism between conversational situations. For instance, if a pair of contexts in one conversational situation both have me as the speaker, and the corresponding pair of contexts in the other conversational situations both have Napoleon as the speaker, the two pairs are similar with respect to the speakers. This is how we capture formally the abstract relation between contexts of *having the same speaker*.

I call such relations *logical relations* because they are among the relations that we exploit when we reason with indexicals.

These arguments show only that identities between speakers matter for validity, as do identities between addressees. A very similar argument shows that identities between any participants matter, even if the identity is between the speaker of one sentence and the addressee of another. As one might expect, we just need to play with the indexicals of (8):

(10) $\langle [c_5, \text{"I am hungry"}], [c_6, \text{"You are hungry"}] \rangle$

(10) is valid only if the speaker of the premise is identical to the addressee of the conclusion. Identities between any participants in an argument count as logical relations; so they need to be preserved in any similar argument.

This should not be too controversial. But we need to note something else: if we stop here, non-identities would not count as logical relations, which, for us, means that we would allow similarity to be non-symmetric. The requirement established so far is only that if there are identities between participants in the original conversational situation, that feature be shared by any conversational situation that it is similar to. The converse hasn't been established yet. Here's an argument that we need the converse. Consider this conversational thread, and assume we're talking about full siblings, i.e. siblings who share both biological parents, and that A, B, and C are distinct individuals:

(11) A says to B: "I am your sibling". B says to C: "I am your sibling". As a conclusion, C says to A: "I am your sibling".

Assume, as seems natural, that sibling-hood is symmetric, and that nobody is their own sibling; in fact, you could just add it as a premise:

$$\forall x \forall y ((\text{Sibling}(xy) \leftrightarrow \text{Sibling}(yx)) \wedge \neg \text{Sibling}(xx)).$$

(11) is of the same kind as all the other examples we've been discussing. It is intuitively valid, and the reason is that the speaker of the first premise is the addressee of the conclusion, and so on. More formally, we need to only look at conversational situations in which $A_1 = S_2$, $A_2 = S_3$, and $A_3 = S_1$. So far, nothing new. But note the added premise that nobody is their own sibling, and look at a conversational situation which has $S_3 = A_3$. Here, the conclusion is spoken by one person to himself. Let's assume that we have a model in which the premises are true, i.e. the speaker of the first premise is the sibling of the addressee of that premise. The second premise is trivially true, because of the symmetry of sibling-hood. The conclusion is obviously false, because nobody is his own sibling. By the standards we have in place, this new conversational situation is similar to the conversational situation we started with: each identity which obtained between participants in the latter is preserved in the former. The difference is that the new conversational situation also identifies the speaker of the conclusion with its addressee.

If we were to count such a two-person conversational situation as similar to the one in (11), we would be forced to conclude that the argument in the three-person conversational situation was not valid. But this would be contrary to intuition. So *non*-identities are also among the relations we exploit when we reason with indexicals.

The point should be clear by now. We are looking for criteria for similarity, which in turn will tell us which conversational situations are relevant for ascertaining the validity of a given argument. Since we want (11) to come out valid, we must not allow the above conversational situation to count as similar to our original conversational situation. This is easily accomplished: we need to require that any identity and difference between participants in a conversational situation be preserved in any similar conversational situation. We thus need the biconditional of which we previously had only one direction, and this will restore symmetry to our relation of similarity.

1.4.2 Similarity and Days

Identities and non-identities also matter for days. This was already obvious from earlier examples:

(2) Today is my birthday. Therefore, today is my birthday.

To get this argument to be valid, we are obviously interested only in conversational situations in which both occurrences of “today” pick out the same day. But days are more structured than participants: they have a metric, which is why we can represent them by positive integers. So there is more to worry about: the relations between the days of the contexts. Recall Frege’s argument:

(7) $\langle [c_1, \text{“It’s raining today”}], [c_2, \text{“It rained yesterday”}] \rangle$

(7) should count as valid just in case the day of c_1 is exactly one day before the day of c_2 . So as far as days are concerned, similarity requires that the *relations* between the days

stay the same, where we include not only identities and non-identities, but also relations like “is exactly one day before”, and so on. Note that it’s not enough to look at just one-day differences. Here’s an argument which requires us to care about two-day differences:

(12) $\langle [c_3, \text{“It will rain tomorrow”}], [c_4, \text{“It rained yesterday”}] \rangle$,

where the day of c_3 is April 2, 2010, and the day of c_4 is April 4, 2010.

Our language is quite impoverished in terms of time-related vocabulary. It is an open question whether we need to make it a part of the definition of similarity that differences of more than two days between contexts need to be preserved. But English contains many more day-operators, like “a week ago”, and so on, so we might as well require that any difference between days be preserved. One thing is clear, though: we don’t need to look at n-tuples larger than pairs of contexts, since days are all ordered on the same scale. More carefully, we will require that, for two conversational situations to be similar, it must be the case that for any two contexts in one conversational situation, if the difference between the days of those contexts is n , the difference between the corresponding contexts in the other conversational situation must also be n .

1.4.3 Similarity and Validity

We are now ready to sum up our results, and gather them together into a characterization of similarity, and then of validity. I will assume we are dealing only with proper conversational situations, so we won’t worry about changes in the worlds of contexts.

For any context c_i , let D_i be the day of c_i , and let P_i be either the speaker or the addressee in c_i , i.e. a participant in c_i . Then two conversational situations will be similar just in case corresponding identities and non-identities among participants are preserved, and so are differences between days. More carefully:

Definition 4. Two proper conversational situations $\langle c_1, c_2, \dots, c_m \rangle$ and $\langle c'_1, c'_2, \dots, c'_n \rangle$ are *similar* iff

- (a) $m = n$
- (b) $(\forall i, j \in [1, n]) A_i = A_j \leftrightarrow A'_i = A'_j$
- (c) $(\forall i, j \in [1, n]) S_i = S_j \leftrightarrow S'_i = S'_j$
- (d) $(\forall i, j \in [1, n]) A_i = S_j \leftrightarrow A'_i = S'_j$ and
- (e) $(\forall i, j \in [1, n]) D_i - D_j = D'_i - D'_j$

Note that conditions (b)–(e) have the effect that similarity of conversational situations is an equivalence relation.

We can now finally present what we were looking for: a definition of validity. The fundamental idea is that an argument is valid just in case for any conversational situation similar to its situation, if the premises are true in their respective contexts, so is the conclusion. More carefully:

Definition 5. An argument $\langle [c_1, \phi_1], [c_2, \phi_2], \dots [c_n, \phi_n] \rangle$, where we think of $[c_n, \phi_n]$ as the conclusion, is *valid* iff for any conversational situation $\langle c'_1, c'_2, \dots c'_n \rangle$ similar to its conversational situation, if $(\forall i \in [1, n - 1]) \phi_i$ is true in c'_i , then ϕ_n is true in c'_n .

Some of these requirements may seem unfortunate. Intuitively, the validity of (7) does not depend on the non-identity of the speaker and the addressee, nor on the fact that the speaker of the premise is the same as the speaker of the conclusion. Once we set up the formal system (which I leave for future work), I conjecture that it will be provable that whenever there are no indexicals sensitive to a certain part of a context, any restriction of the similarity class due to that feature of the context does not matter for validity. So the result is at most cumbersome, but it makes no real difference.

1.4.4 The Deduction Theorem

We should also note that the Deduction Theorem fails, for two reasons: as commonly expressed, it does not apply to LD, and, second, when re-written to suit it, it is false. A

classical formulation of the theorem is this:

Classical Deduction Theorem: $(\forall n \geq 0)$ an argument $\phi_1, \phi_2, \dots, \phi_n, \psi \therefore \chi$ is valid iff $\phi_1, \phi_2, \dots, \phi_n \therefore \psi \rightarrow \chi$ is valid.

The Deduction Theorem allows us to move freely between discussions of validity and discussions of logical truth. The reason is that it guarantees the equivalence of an argument with any finite number of premises with an argument with no premises, with the former premises turned into antecedents in the conclusion.

As given, the Theorem doesn't apply to LI, since we think of the steps of an argument as context-sentence pairs. Note that we cannot have something like $[c_n, \psi] \rightarrow [c_{n+1}, \chi]$, because that is not well formed. The reason is that we have a stark contrast between the conditional, which is a sentential connective, and \therefore , which is a relation between steps in an argument, which might consist of pairs of a context and a sentence.

The most straightforward adaptation of the Classical Deduction Theorem to LI would be this:

Bad Indexical Deduction Theorem: $(\forall n \geq 0)(\forall i)$ an argument

$\langle [c_1, \phi_1], [c_2, \phi_2], \dots, [c_n, \phi_n], [c_{n+1}, \psi], [c_{n+2}, \chi] \rangle$ is valid iff $\langle [c_1, \phi_1], [c_2, \phi_2], \dots, [c_n, \phi_n], [c_{n+1}, (\psi \rightarrow \chi)] \rangle$ is valid.

That the Bad Theorem fails is clear: on the left hand side we have ψ in c_{n+1} and χ in c_{n+2} , whereas on the right hand side we have $(\psi \rightarrow \chi)$ in just one context, c_{n+1} . Moving χ to the same context as ψ can obviously make all the difference between validity and invalidity.

Besides, we allow changes in context between steps in an argument, but not within each step in the argument. So we cannot just choose one of the contexts on the left hand side of the Deduction Theorem and have it be the context of the conclusion of the right hand side. The proof is not complicated. Take the following argument: $\langle [c_1, \phi], [c_2, \psi] \rangle$.

If we try to turn it into a single conditional, we are forced to have something of the form $\langle [c_3, (\phi \rightarrow \psi)] \rangle$, where we must choose a single context c_3 . The problem is that it doesn't matter if we choose c_3 such that $c_3 = c_1$, or $c_3 = c_2$, or a completely different context; no matter what we do, we cannot guarantee the equivalence.

Here's an example: $\langle [(S_1, A_1, D_1, W_1), \text{"It's raining today"}], [(S_1, A_1, D_2, W_1), \text{"It's raining today"}] \rangle$, where we suppose that D_2 is the day right after D_1 . This is not a valid argument, since the days of the two contexts are not the same, and rain one day does not logically guarantee rain the next day. Now if we try to apply the Indexical Deduction Theorem, and choose the first context to work with, we'd get $\langle [(S_1, A_1, D_1, W_1), \text{"(It's raining today} \rightarrow \text{It's raining today)"}] \rangle$. But this *is* a valid argument, since we just repeated the same sentence, in the same context. And it doesn't matter which context we choose; in fact, the sentence "If it's raining today, then it's raining today", when put in any context, will turn out to be a logical truth, as, I think, it should.

Does this failure of the Deduction Theorem mean that we don't have a logic? Such worries may be dispelled when we notice that a narrower version *does* hold. If we restrict the theorem so that the conclusion and the last premise in the left hand side are in the same context, we preclude the problem we had earlier:

Indexical Deduction Theorem: $(\forall n \geq 0)(\forall i)$ an argument

$\langle [c_1, \phi_1], [c_2, \phi_2], \dots [c_n, \phi_n], [c_{n+1}, \psi], [c_{n+1}, \chi] \rangle$ is valid iff

$\langle [c_1, \phi_1], [c_2, \phi_2], \dots [c_n, \phi_n], [c_{n+1}, \psi \rightarrow \chi] \rangle$ is valid.

In fact, the choice of the last premise is arbitrary; any premise that has the same context as the conclusion can be moved to the right hand side as an antecedent. This is just one of the consequences of the fact that the logic does not make any use of the fact that the premises are ordered in an argument, as opposed to a simple set, as the usual definition has it.

All this is a consequence of our decision to allow the context to change between steps in an argument, but not within any single step (i.e. we cannot have two contexts per

sentence). There is something intuitive behind this decision, even if in real life, each step of an actually made argument takes time, and may well even span two days (think of a premise stated beginning just before midnight, and ending immediately afterwards. Which is the day of the context?). As argued in [Demons], p. 522, we must abstract away some features of utterances, lest we end up with issues like the fact that some arguments are just too long to be uttered in the life-span of the Universe, and so, they would never be uttered truthfully, and so any argument containing them would seem to be valid by default.

1.5 Logical Truth

This is a natural place to introduce the notion of logical truth. The topic deserves its own section because logical truth in LI will turn out to be more interesting than the usual notion of logical truth: on the technical side, its connection to validity is more complicated, and on the philosophical side, some initially problematic examples will hopefully lose their discomfoting nature by teaching us something about the relation between semantic knowledge and logic. One clear example is the fact that all logical truths of LI are also logical truths of LD, *except* for the ones introduced by the addition of “you”. As we will see, this indexical brings with it new kinds of logical truth, which did not come up in LD; the logic of “you” and “I” is more interesting than the logic of “I” alone.

A single sentence is usually considered a logical truth when its truth is guaranteed, irrespective of the premises of any argument of which it is a conclusion. This allows us to include arguments without any premises. We need to adapt this definition to the current framework, where we don’t just look at sentences, but sentences in contexts:

Definition 6. $[c_1, \phi]$ is a *logical truth* iff for any $\langle c_2 \rangle$ which is similar to $\langle c_1 \rangle$, $[c_2, \phi]$ is true.³⁰

³⁰Note that we talk of the sequence $\langle c_1 \rangle$, because we have defined similarity in terms of sequences. For ease of presentation, I will occasionally slip into talking of similarity between contexts when we only have sequences with just one context in them.

This definition of logical truth is a natural extension of the usual definition. In LD, a sentence is said to be a logical truth just in case it is true in all contexts. According to the new definition, a sentence *in a context* is a logical truth just in case it is true in all contexts which are similar to its context.

The reason why the new definition doesn't collapse into the old one is that because we introduce "you" in the vocabulary, we now must include an addressee in our contexts, something contemplated, but not elaborated in [Demons]. Not all contexts are similar, because the identities and non-identities between the speaker and the addressee need to be preserved. All contexts can be divided into those in which we have a self-addresser, and those in which we do not.

Technically, the definition of logical truth given above stems easily from our framework. But once we apply it to actual conversational situations, some strange things happen: we get examples of logical truths which do not feel like logical truths at first blush. I think that this seeming problem is, in fact, an advantage, because it will prove that some very simple logical truths are not always epistemically immediately available.

Before I present the problem technically, I'd like to relate it to an older problem about self-knowledge. Put yourself in Mach's shoes, who didn't realize that he was addressing himself in the mirror when he said "what a shabby pedagogue you are". Perry (2001, §5.4) used this example to show that there is something special about first-person beliefs, since Mach did have a belief about himself in that conversational situation, but not a *first-person* belief about himself.³¹

Perry's example is very persuasive, but it does not include an assertion, so let me consider a simplified, though somewhat artificial, version:

(13) $\langle[(\text{Ernst Mach, Ernst Mach, D, W}), \text{"I am identical to you"}]\rangle$

This is a soliloquy: the speaker is identical with the addressee. By claim (d) in our defi-

³¹In fact, Perry's example used "what a shabby pedagogue that is". I changed the demonstrative to an indexical, since I have not included demonstratives in LI.

nition of similarity, if the speaker of any context in the conversational situation is identical to the addressee of any context, that identity must hold in all similar conversational situations. This definition includes cases in which the speaker of a context is identical to the addressee of that very context, and this is the conversational situation in (13). So the argument is a logical truth, since there is no conversational situation similar to the initial one in which the speaker is not identical to the addressee.

But this seems unsettling. If it does not, here's how one may be unsettled: begin by granting that it is a necessary truth; after all, true identities with directly referential terms are always necessary. If you're in a conciliatory mood, you may admit that most of the time, the speaker knows when she is addressing herself. This is surely correct; but doesn't it also just seem wrong to say that Mach didn't know a *logical* truth? And yet this is exactly what my theory claims. So, should Mach's sad realization be blamed on lack of logical acumen?

One might think that the moral is that we should exclude intra-contextual identities from our definition of similarity. I think not. Think of this situation: speaking to myself, I say "If you stop drinking three milkshakes a day, you will lose weight." I continue: "I will stop drinking three milkshakes a day". I conclude: "I will lose weight". This argument is fine when lecturing others who accede, why shouldn't it be good when lecturing myself? The issue is complicated, since it will require a discussion about the relation between LI and live conversations, and about the proper scope of logic. But this will have to wait until Ch. 3. For now, I just want to show that any reason to put aside intra-contextual identities would generalize to an argument against any definition of similarity like the one I've proposed. Similarity is an all or nothing affair.

Suppose we change condition (d) in our definition of similarity of arguments to (d*), so that we only require that identities between speakers and addressees be preserved between distinct contexts, but not within them:

$$(d^*) (\forall i, j \in [1, n], i \neq j) A_i = S_j \leftrightarrow A'_i = S'_j$$

This way, we make (13) logically contingent, since there are conversational situations now deemed similar to ours where the speaker is not identical to the addressee, and in those, “I am identical to you” is false. Appropriately, we also change the status of “I am not identical to you”: in LI, in the context of (13), it was a logical falsehood; now, it is logically contingent.

It might look as if we have succeeded in making Mach logically blameless. But he would remain just as blameless in situations much like the original one. The reason is that changing (d) to (d*) does not affect the importance of cross-contextual identities to similarity. These we cannot just excise from our logic; they were the main instrument in answering Soames’s challenge. Therefore the following argument, which is like having (13) for a conclusion, but adds some premises, remains valid:

(14) $\langle [(\text{William James, Ernst Mach, } D', W), \text{“You are a genius”}], [(\text{Ernst Mach, Ernst Mach, } D, W), \text{“I am identical to you”}] \rangle$

Note that what is said in the premise is irrelevant, as is the speaker and the day; the example is good only because it triangulates the identity involved in the conclusion with the help of the premise. Thus, (14) comes out valid even with (d*) because cross-contextual identities are sufficient to ensure that every conversational situation similar to the original one has the following property: the addressee of the premise is identical with the speaker of the conclusion, and with the addressee of the conclusion. Since identity is an equivalence relation, that is sufficient to guarantee that the conclusion comes out true.

Now I can make my point: is (14) any different from (13)? What did Mach not know just before his realization? That he was looking at himself, in Perry’s example; if he had addressed that person, he would not have known that he was addressing himself. According to the original definition of similarity, that makes him ignorant of two facts: that (13) was a logical truth, and that (14) was a valid argument. According to the version with (d*), he would have not known a logically contingent truth, namely that (13) is true, but he would still have not known that (14) was a valid argument. The latter seems just as

insulting to Mach's intelligence as the former; namely, not at all. These valid arguments are valid not merely in virtue of the sentences uttered, but in virtue of facts about the contexts. Logical acumen doesn't guarantee self-recognition, or tracking of days.

1.6 Why Not a Logic on Contents?

One of the distinctive features of LD was that it claimed that some logical truths were not necessary truths. The distinction between the two kinds of truths is already a difficult one, but Kripke (1980) had already presented a few kinds of truths which he argued were necessary, but a posteriori, such as "Water is H₂O". Since logical truths are generally thought to be a priori, and since this example doesn't really look like a logical truth, this already shows that necessity and logical validity are not coextensive.

One surprising consequence of [Demons], which goes beyond Kripke here, is not only that there are some necessary truths which are not logical truths, but the entailment doesn't work in the other direction either: some logical truths are not necessary. And this is surprising: how can a logical truth not be necessary?

To make it easier to keep track of this issue, here are some examples: "I am here now", "I exist", "I am rich iff I am actually rich". These sentences are logical truths in [Demons], but not in FOL, because of certain special features of indexicals. Let me focus on "I am here now" and explain briefly why this sentence is a logical truth, without being a necessary one as well. This explanation applies equally well to the other examples considered above.

Any sentence containing indexicals is evaluated for its truth value at the context in which it is uttered. This is meant to capture the idea that, say, the truth value of an utterance of "I am happy" should depend on the happiness of the speaker at the time and in the world in which he utters the sentence, not on someone else's happiness. As we saw earlier, [Demons] only allows *proper* contexts, i.e. contexts which have the speaker being

at the location of the utterance, at the time of the utterance, in the world of the utterance. This says nothing about the speaker's happiness, so "I am happy" is not a logical truth. But "I am here now" cannot fail to be true at any context, since it expresses just what it is for a context to be proper.

Necessity, according to [Demons], applies not to sentences, but to what sentences say, to propositions. For instance, if I say "I am here now" in Austin, TX on Dec 30th, 2010, I express the proposition that I, the author of this chapter, am in that city, on that day. But I could have been somewhere else on that day, obviously, and I could even have died before then. In general, no proposition expressed by "I am here now" is necessary; after all, any speaker could be somewhere else at that time, or, if unlucky, could be nowhere at all at that time.

Kaplan got these nice results because LD was a logic of sentences, not propositions. But LI works on pairs of a sentence and a context, and, once given a context, every sentence expresses a proposition. So a natural objection to my system is this: why not a logic on those propositions, on contents? Why have a complicated logic, on pairs of things, instead of a simple one, on just propositions? What is the difference? And are there any advantages either way?³²

The answer is that there probably are differences. The hedge is needed because there is no such thing as "the logic on contents", partly because there is no consensus on what these contents are. Here are some live options: they could be sets of possible worlds; or structured Russellian propositions, in which case views differ on the structure, and on the constituents; or senses (possibly structured); and many other permutations are also available. They all deserve a longer discussion, but for now, let me make two simple points, which suffice to show that there may well be some advantages to my view.

³²There are also "mixed" systems, which put into the proposition, for instance, both names (typical sentential denizens) and their referents (good candidates for propositional constituents). They will likely be as fine-grained as my system, if not more, so they are a more direct challenge to my system. Larson & Ludlow (1993) is a good example of such a system. They thus abandon the idea that people speaking different languages ever say the same thing.

The first difference is that some arguments come out valid in my system, and probably invalid in a logic on contents. Just as in LD, the following is a valid argument: $\langle [c_1, @\phi], [c_1, \phi] \rangle$. Assume that ϕ is a non-modal sentence (i.e. it contains no modal operators). Then, the argument made up of the propositions expressed by those sentences in those contexts is not valid in a system obtained by suitably adapting LD. Here's a simple outline of the proof. Assume that the premise is true when set in the context c_1 . The proposition expressed says that ϕ is true in a particular world, the actual world of c_1 . This proposition, if true at all, is true when evaluated at any other world. The conclusion, on the other hand, expresses a proposition which says nothing explicitly about any particular world. It thus expresses a proposition which is true when evaluated at a particular world just in case it ϕ is true in that particular world. So the proposition expressed by the conclusion is false when evaluated at worlds in which ϕ is false, and at those worlds the proposition expressed by the premise is still true (since it is still about the actual world of c_1 , no matter what world we evaluate it at). So the argument is not valid according to this logic on contents.³³

There are also differences in the other direction: there are truths which come out to be logical truths in a logic on contents, but not in my system. For instance, if logic is on contents, then the repetition of any proposition is a logical truth.³⁴ But, according to many conceptions of propositions, the following sentences express the same proposition: "Hesperus is a planet", and "Phosphorus is a planet". So whenever someone argues from the former to the latter, they make a valid argument, as valid as just repeating the first sentence. This is usually taken to be a bad result.³⁵

³³This argument should remind readers of Kaplan's argument that some logical truths are not necessarily true. We should also note that it all depends on how the logic on propositions is set up; one could, for instance, claim that there are no non-modal sentences, just as one can claim that there are no tense-free sentences. My argument, thus, only works given the assumptions I made, and, reasonable as I think they are, they can be avoided, even if at some cost.

³⁴This needs to be qualified: as stated, repetition is valid only if propositions have their truth values by themselves. If they have truth values relative to some indices, such as possible worlds and times, then repetition is only valid relative to the same indices.

³⁵If you're not a Millian, pick any two expressions your favorite theory counts as synonymous, and run

Here's another example, using an indexical: "I am Alex Radulescu", in the context (Alex Radulescu, David Kaplan, D' , W), expresses a proposition of the same form as above, $x = x$, and so it must be valid in a logic on propositions. But this sentence in that context, though true, is – as it should be the case – *not* valid in LI (or in LD, for that matter).

The point is that for logical purposes, we care not only about referents, but also about the linguistic items. For names, having some referent or other is important, but the identity of the referent is irrelevant. For indexicals, I have argued that we care about structural relations within and between contexts; so we care about a particular aspect of reference, though not quite down to the level of actual reference. In general, going directly to contents leads to a loss of logically relevant information.

the same argument. If your favorite theory is so strict that it allows no synonymous expressions, it is wrong, and shouldn't be your favorite to begin with.

Chapter 2

The Second Person Singular Pronoun and The Difference between Indexicals and Demonstratives

2.1 Introduction

Words which clearly depend on context for their semantic value are not hard to spot. As Cappelen & Lepore (2005) remarked, once you begin the list, anyone can continue it. Just mention “I”, “today”, and “here”, and others will easily spring to mind: “you”, “she”, “yesterday”, “there”, “local”. “Actual” may not come to mind immediately, but some training in the philosopher’s modern usage of the word would be sufficient to add it to the list. Many other kinds of expressions are also thought to be contextually sensitive: modals, epistemic modals, comparative adjectives, tenses, etc. I will focus on the short list above.

The most influential analysis of these context-dependent words is in [Demons] and its companion, [Aft]. In this chapter, I will be concerned with Kaplan’s claim that this class of words should be split into two: pure indexicals and true demonstratives. The

distinction is not quite as easy to see as the one between context-sensitive words and the rest, and I will spend a fair amount of time teasing it out, but a simple characterization will do for now. Pure indexicals are words like “today” and “I”, which get their semantic value directly from the context of utterance, but take no account of the speaker’s referential intentions.¹ True demonstratives are words like “this” and “that” (and their complex offspring, like “this book”, and “that tree”), which do not work as straightforwardly as pure indexicals. Theories about the mechanism by which true demonstratives get their referent abound, but let’s just say that true demonstratives can only get a referent by taking account of the speaker’s intentions, or actions, or some other speaker-related feature of the situation.

I find it hard to remember which of them are pure, and which are true, so I will forgo these adjectives in this chapter. The downside is that Kaplan used the simple label “indexicals” for the whole class, so I am forced to use another expression: I will use “deictics” for both indexicals and demonstratives, and keep “indexicals” for “pure indexicals” and “demonstratives” for “true demonstratives”.

I have several goals in this chapter. First, I propose a new way to draw the distinction between indexicals and demonstratives. Traditionally, the distinction has hinged on the semantic relevance of the speaker’s intentions. I argue that this is not the right way to go. My proposal depends on distinguishing between kinds of contextual features and kinds of roles that intention plays in uses of deictics. An interesting case will be the second person pronoun, “you”. The tradition has not focused too much on it, but it provides me with a good example of the way in which my proposal is different, and, hopefully, preferable to the alternatives. In fact, I will argue that “you” is an indexical, even though traditionally it has tended to be categorized as a demonstrative. Second, I discuss some variants of the classical distinction, and I argue that some fail to distinguish even be-

¹I’m talking about the simplest use of the word “today”. When a curmudgeon begins a sentence with “the problem with young people today is...”, he is not complaining about something that happened that very day. I leave such uses aside.

tween paradigmatic indexicals and demonstratives, and others could be developed into a variant of my proposal. Finally, in §2.8 I give some reasons why I think the distinction matters, beyond the claim that it captures an important feature of context sensitivity.

Here's the plan, in a bit more detail. In §2.2, I present the difference between indexicals and demonstratives as it was first given by David Kaplan, in [Demons], where the terminology was introduced. This is the most influential discussion of our issue, influencing the rest of the literature. The basic idea is that the speaker's intentions are needed to secure a referent for occurrences of demonstratives, but not for indexicals. This idea is embedded in a general picture about deictics, which are context-sensitive words. Therefore, it will turn out that how we think about contexts will influence how we think about deictics. Towards the end of the section, I present David Lewis's notion of a context, and argue that, in its naturalistic bent, it is a very useful starting point.

In §2.3, I present my own take on the difference between indexicals and demonstratives. The central point is fairly straightforward: indexicals are those words which depend on those contextual features which belong to utterances in general, irrespective of the expressions used. For instance, utterances are typically made by a single speaker; "I" gets its referent from this feature, but utterances have this feature irrespective of whether they contain the first person pronoun. The contrast is that demonstratives need a helping intention that is specific for them; an utterance has that kind of feature – one or more helping intentions – only for the sake of the demonstratives it contains. No demonstratives, no such contextual features. The point is quite intuitive, I think, but spelling it out will require a fairly careful distinction between this kind of contextual features, and the various roles that intentions can play in our linguistic theory. One good consequence of this discussion is exploited in §2.4, in which I attempt to show where and why the tradition went astray.

Most of my arguments draw their power from reflections on typical utterances. But the literature is split on the question whether these are the right vehicles for logic, and for

formal semantics. The other candidates, sentences in a context, are obtained by abstracting away certain features of real utterances, such as the fact that utterances take time to make, and therefore extremely long ones may be physically, or even metaphysically, impossible. In §2.5, I show how one can accommodate my proposals within the strictures of the more abstract sentences in a context.

Having made my position clear, in §2.6 I come back to the topic of the second person pronoun, and show how the tradition differs from my proposals in the way “you” has been categorized. This section is, on the one hand, an application of the more general discussion that precedes it, and, on the other hand, an opportunity for me to show how naturally it divides the kinds of deictics, and, in particular, why its consequence that “you” is an indexical is a correct.

What I have been calling “the tradition” is mostly based on certain seminal passages from [Demons]; and the label is appropriate, because all other discussions, to some extent, are elaborations of those ideas. But my proposal has other rivals as well, which I present and discuss in §2.7. Many of them, I will argue, are inadequate for the purpose, since they do not in fact achieve their goals, even before we bring in my claims about “you”. For instance, in §2.7.3, I discuss Eros Corazza’s idea that indexicals cannot fail to have a referent, whereas demonstratives can. I argue that indexicals can be empty too; “you” will be a clear case, but there are others as well. The observation may hold for a distinguished class of indexicals, and, if so, it would be interesting to figure out the consequences of this fact for logic. But I will be content to note that this is not the way to separate indexicals from demonstratives. This section is not purely negative, however, since not all criteria in the literature miss the mark in this way. In §2.7.1, I discuss Perry’s two separate distinctions within the field of deictics, and I argue that one of them can be developed in a way that is friendly to my own proposals.

In the course of my dissertation as a whole, and in particular in this chapter, most of my time is spent on indexicals. In the final section, §2.8, I talk directly about demon-

stratives. I do not propose a semantics for them, nor a logic. But some lessons can be drawn from the fact that demonstratives need the assistance of made-to-order intentions: for demonstratives, we need full-on utterances, not mere sentences evaluated relative to a context. The reason is that each use of a demonstrative needs to be accompanied by an actual intention. And these intentions don't occur in logical space; they belong to real speakers, when they use that very token in a particular sentence. This provides a programmatic answer to the question whether we need utterances or just sentences in a context: for indexicals, we can have (and, in fact, need) the latter; for demonstratives, we need the former. I say "programmatic" because this still leaves the work of actually incorporating demonstratives in LI, which I leave for further work.

2.2 The Difference, Traditionally

The simplest way to make the distinction is the way I did it at the beginning: by example. "Today" and "I" are indexicals, and "this" is a demonstrative. And then we can begin to wonder if some other words are more like paradigmatic indexicals, or more like demonstratives. This method does not depend on any particular semantic theory, which is a virtue if there really are two semantic kinds here. And I do think there is an intuitive difference between the two kinds of words. One problem is that once we work out a theory for them, it may turn out that there's no semantically interesting difference. And if we find more than one interesting difference, it may not be easy to figure out which one is criterial, and which one simply depends on our first choice of paradigms.

The tradition has many tributaries, since there is quite a vast literature on the semantics of demonstratives, and there doesn't seem to be any widely accepted view just yet.² But here is a sketch that should, broadly speaking, be acceptable to most: "today", in

²Focusing just on bare demonstratives, see Bach (1992), Reimer (1992), Braun (1996), Salmon (2002), and Elbourne (2008). For complex demonstratives, see King (2001) for a very influential account, and Wolter (2009) for a good brief overview of the literature.

virtue of its meaning, picks out the current day, the day of the context of utterance. This linguistic rule is sufficient to secure it a referent in any context. By contrast, “this” has no such simple rule associated with it; it doesn’t pick out a particular feature of the context. Rather, the speaker must intend to refer to a particular thing with that use of the demonstrative, and, unless something precludes it, like the speaker hallucinating the intended object, “this” will pick out the intended object.³ It may be a part of the conventional meaning of “this”, best seen in contrast with “that”, that the former suggests proximity, the latter remoteness. But this is not enough to give the word a referent in every context; after all, there always are many things near the speaker, and only one of them can be the referent, and the choice is left open by this sort of linguistic meaning of the word.

Here is a representative quote from [Demons], pp. 490–491, where the distinction was first made and the terminology was proposed:

“[Demonstratives] require, in order to determine their referents, an associated demonstration: typically, though not invariably, a (visual) presentation of a local object discriminated by a pointing. These [deictics] are the true demonstratives, and ‘that’ is their paradigm. The demonstrative (an expression) refers to that which the *demonstration* demonstrates. [...] For [indexicals], no associated demonstration is required, and any demonstration supplied is either for emphasis or is irrelevant. Among the pure indexicals are ‘I’, ‘now’, ‘here’ (in one sense), ‘tomorrow’, and others. The linguistic rules which govern *their* use fully determine the referent for each context. The speaker refers to himself when he uses ‘I’, and no pointing to another or believing that he is another or

³The intentions involved are probably more complicated than this. According to Recanati (2007, p. 215), they are much more complicated: the referent of a demonstrative is “the entity such that the speaker makes manifest to the hearer his or her intention to bring it to the hearer’s attention by means of the hearer’s recognition of this intention”. I doubt that this is correct. First, it seems to require too much of the speaker: the hearer should not be able to prevent the speaker from referring to something just by refusing to pay attention, and therefore making the speaker fail to make his intentions manifest to him. Second, this kind of Gricean intention seems too rich and too complicated for most uses of demonstratives. Still, the structure of the intentions is not relevant to my goals, so I will continue to talk in the simplest terms, of an intention to refer to something.

intending to refer to another can defeat this reference.”

Kaplan talks here more about demonstrations than about intentions. He switches completely to talk of intentions in a paper, [Aft], written 15 years later. This is symptomatic of a large debate about the mechanism of reference for demonstratives: are intentions sufficient to secure a referent? This is an important question, but it need not be answered here.

What is common to most views is that the speaker’s intentions play *some role* in determining the referent of demonstratives, and none for indexicals.^{4, 5}

Intentions do, indeed, look like the distinguishing mark between indexicals and demonstratives. Here is a just-so story about how one could come to have this view.⁶ Start out from looking at a couple of sentences: “I am happy”, and “This is interesting”. For simplicity, put aside the tense; it is relevant, since few of us are happy all the time, so the first sentence is true at some times, and false at others. It is clear that neither sentence has a truth value *tout court*, since, by themselves, the singular terms lack a referent. They both depend on the context of utterance in a fairly systematic way. “I am happy” is true in a context if the speaker of the sentence is happy at that time. So in order to get a truth value for a sentence containing an indexical, we need to evaluate the sentence relative to a context with a speaker in it. Note that I didn’t mention the speaker’s intentions. “I” does its job automatically, without looking into the mind of any participant; it just picks

⁴This is the majority view, but, as always, there is a minority. Wettstein (1984, pp. 72–73) argues that the line stretching from the speaker’s finger determines the referent of demonstratives accompanied by a pointing. Quine (1968, p. 194) is often cited as a precursor, but Quine was not talking about demonstratives in that passage; Ryle (1949, p. 188) is a better, though still imperfect, fit. These difficulties have led to a mini-tradition of intentionally ignoring the issue, and just continuing to talk in terms of demonstrations as whatever it is that demonstratives need to get a referent. See Salmon (2002) and Kaplan (2003).

⁵Of course, the speaker’s intentions play a role in selecting which word to utter. They determine that on a particular occasion we have a use of the first person pronoun, not an expression of surprise, nor a mention of the ninth letter of the English alphabet. The remarks in the main text only apply once these pre-semantic issues are resolved.

⁶It is a just-so story because I do not claim that anyone ever actually thought along these very lines; but I do find it compelling, more than Kipling’s just-so stories. It should also be obvious that I am focusing only on the issue of the mechanism of reference for deictics; truth in a context is a much more complicated affair than I make it sound, but the complications are not relevant here.

out the speaker. In fact, we don't even need to look at a real situation involving a person really speaking. We can say: suppose Barack Obama was speaking; the sentence would be true just in case he was happy at the time. "I" still does its job: it picks out the speaker from this (imagined) context.

"This is interesting" is different. Suppose we fix a context of utterance; it involves Obama again, speaking to Joe Biden in the Oval Office on March 9, 2012. Our sentence is still up in the air; we don't yet know enough to evaluate it at that context. Intuitively, various questions are appropriate: what was Obama talking about? Was he pointing at anything? Or, the even more inclusive question: what did he mean? All of these questions can only be answered by looking inside the mind of the speaker, so to speak, or at the outward manifestations of his intentions, such as pointings, nods, etc.

This is not to say that intentions play no role in uses of indexicals. When Obama says "I am happy", he intends to talk about himself just as much as when he says "This is interesting" he intends to talk about something in particular. Typically, when the speaker uses a singular term, he intends to talk about the referent of that singular term. We can think of possible counterexamples, of course; it is not unreasonable to think that when a parent says "Santa Claus brought you a gift", she doesn't intend to talk about a non-existent person.⁷ But our two sentences are not problematic; they both involve intentions to talk about something, to refer to something. What matters is the role of the intentions in fixing the referent. When Obama says "I am happy", his intention to talk about himself is fulfilled. This intention however is irrelevant to the way "I" gets its referent, which is just in virtue of the meaning of the word and the fact that Obama is the speaker. "This", on the other hand, depends both on the meaning of the word and also on the intention-laden fact that Obama intended to refer to something in particular. This *further* intention is the crucial part.

⁷Not unreasonable, but also not undeniable: Kripke (2011) and Salmon (1998) claim that "Santa Claus" does refer to an existent abstract object.

A composite criterion, which I offer as an explication of the quote reproduced above from [Demons], emerges from these observations:⁸

- **Linguistic Rules** by themselves never fully determine the referent of a demonstrative in a context, but they always fully determine the referent of any indexical in any context.
- **Intentions** are required for demonstratives to gain a referent, but not for indexicals.

I have presented these ideas as neutrally as I could. Still, they do conflict with a number of views. For instance, mutually recognized salience has been proposed as a mechanism of reference for both indexicals and demonstratives.⁹ On this view, the need for intentions cannot help us distinguish between the two kinds of deictics, and, indeed, we shouldn't even attempt to distinguish them.

The part about linguistic rules is a bit less controversial, but it also depends on some assumptions about what counts as a linguistic rule, and what such a rule is allowed to work on in a context. For instance, it depends on the idea that each word is governed by a single linguistic rule, which is then counted on by the speaker to do its job in any context. Some philosophers have been arguing against this view, claiming that the stable meaning of most words vastly underdetermines their semantic value in a Kaplan-style context.¹⁰ From now on I will assume that these worries can be addressed, and will not mention them again.

In the remainder of this section, I discuss a preliminary worry: from the beginning,

⁸Note that I am only claiming to present the ideas in that particular quote from [Demons], not the overall view from the whole paper. In fact, there are several views about contexts, both in [Demons] and in [Aft], which make it unlikely that what I present was, at any point in time, Kaplan's considered view. One example: if contexts are chunks of the world which may (but don't have to) contain and surround an utterance, as is clearly stated in [Demons], p. 494, then they may well contain the speaker's intentions as well; if demonstratives have a kind of character which is sensitive to the speaker's intentions, it will be able to find these intentions in the context, just like the character rule for "I" finds the speaker in the context. What happens if there are no intentions? It turns out that [Demons] avoids this question, by analyzing away demonstratives in terms of the rigidifier *dthat* and definite descriptions; but the issue comes up later in [Demons], as well it should, conceptually speaking. That the speaker's intentions are part of the context is clear in token-reflexive accounts of deictics; see §2.7.1 for a presentation and discussion of Perry (2001).

⁹See Mount (2008).

¹⁰See, for instance, Travis (1996) and Bezuidenhout (1996).

talk of deictics has been tied to talk of contexts of utterance. One might think that this is where the distinction between indexicals and demonstratives could be made most easily. The reason is this: if linguistic rules fully determine the referent of indexicals in each context, then all we need to do is make an inventory of words with a similar behavior, and what is left will be the demonstratives.

It's not that easy: why are the speaker's intentions not part of the context? After all, when speakers use a demonstrative, they normally do have a corresponding intention. If we think of contexts in this richer manner, and add that occurrences of "this" pick out the intended object (plus whatever other conditions we place on it), then demonstratives and indexicals seem to function in exactly the same way. The difference would then be just that some features of contexts are about, or contain, the speaker's intentions, and some not.

To some extent, this dispute about contexts may look like a decision about terminology; but it is also a decision about our goals. We could simply think of a context as a device designed for semantical theorizing, which gives values to indexicals, and whose connection to real world utterances is left open for research.¹¹ Then, plainly, the difference between indexicals and demonstratives will depend not on their semantical representation, but on the way in which we decide which are the relevant features to be represented in our contexts. Still, the traditional distinction between indexicals and demonstratives looks right, since intentions are a plausible candidate for a central part in the story of demonstratives, but not in that of indexicals.

Another option is to think of contexts more robustly, as parts of a possible world. Here is David Lewis:

Whenever a sentence is said, it is said at some particular time, place, and world. The production of a token is located, both in physical space-time and

¹¹Simplifying, this is the position of Predelli (2005). See also [Aft], p. 591, where this idea is mentioned, though not adopted.

in logical space. I call such a location a *context*.

That is not to say that the only features of context are time, place, and world. There are countless other features, but they do not vary independently. They are given by the intrinsic and relational character of the time, place, and world in question. The speaker of the context is the one who is speaking at that time, at that place, at that world. (There may be none; not every context is a context of utterance. I here ignore the possibility that more than one speaker might be speaking at the same time, place, and world.) The audience, the standards of precision, the salience relations, the presupposition... of the context are given less directly. They are determined, so far as they are determined at all, by such things as the previous course of the conversation that is still going on at the context, the states of mind of the participants, and the conspicuous aspects of their surroundings.¹²

According to Lewis, contexts are simply locations in a possible world. These locations may contain a speaker, or they may not. The speaker, if there is one, may be intending to refer to something, or not. What makes this conception quite robust is that the semanticist does not dictate what is to be found in a context. The world does that, and, presumably, the world also tells us if a particular sentence is being uttered in a context or not. All semantics can do is look at the situation it is being given, and build a theory to fit it.

In short, Lewis starts by looking at a sentence as uttered; when we have that, we have everything we need to give a semantic value to any deictic, indexical or demonstrative. If we follow Lewis, the difference between indexicals and demonstratives is not that the former get a referent in a context, unaided by intentions. Since users normally have demonstrative intentions when they use a demonstrative, the intentions will be right there for the semantics to see; the context contains those intentions, and presumably their object, so there is no technical difference between indexicals and demonstratives. Rather,

¹²Lewis (1981, p. 85–86); his emphasis.

the difference is made by the fact that demonstratives get their referent “less directly”, which I take to mean that they work, in part at least, off the speaker’s intentions, whereas indexicals get their referent “directly”, i.e. just in virtue of the location of the utterance, and the identity of the speaker.

How we think of contexts will influence the structure of our semantical theory; we will come back to these issues in Ch. 3. But, and this is the lesson I draw for now, none of it matters for the issue at hand: intentions, wherever they belong in the final analysis, are the place to look for the distinction between indexicals and demonstratives.¹³

This much seems hard to dispute. To see why, let me quickly pick an admittedly straw man-ish theory¹⁴: suppose the speaker’s intentions were irrelevant, and all that mattered was what was being pointed at by the speaker. Imagine my arm is acting up again, and it’s pointing at random things. Suppose it’s pointing at Venus, but I don’t pay any attention to my arm; I really want to talk about the cup of coffee in front of me. I say: “This could use more cream”. I take it to be very implausible that I’m talking about Venus, even if some observer, unaware of my ailment, may erroneously take me to be doing so. All this is not to say that the pointing theorist must categorically be wrong. Suppose someone thinks that some way of making the intended referent salient is required for demonstrative use. Then I could, for instance, observe my arm more carefully, wait until it’s pointing at the coffee, and utter my sentence then. The pointing would not be an intentional act, but I’d be making use of it nonetheless. It could then be claimed that the pointing was necessary, but insufficient without the accompanying intention.

¹³There is one question I will not address: which intentions are relevant for demonstratives? Suppose I really want to talk about Venus. I spot a light up in the night sky, and I say “That is a planet”. Suppose the light comes from an airplane, and Venus is not visible from where I’m standing. Prima facie, I had two intentions: to talk about Venus, and to talk about that light. Which one wins? And why? Important questions, to be sure, about which intentions are relevant for what, but however we answer them, most answers share the element I need: some intentions are relevant for demonstratives. For dissenting views, see Bach (2005) and Gauker (2008).

¹⁴It is not very far from Wettstein (1984), though far enough not to make Wettstein a straw man. He talks of being responsible for the clues one gives to the audience; in my case, the speaker is not responsible for his gestures.

2.3 The Difference

In this section, I argue that the tradition does not have it quite right; that simple talk of intentions is not sufficient to distinguish between indexicals and demonstratives. In particular, I will claim that we need to look at the kinds of features of contexts of real utterances that deictics depend on, and that some indexicals depend on intentions.

My central argument is quite complicated, and it will require a fairly long discussion. But there is a simple way to see the attractive nugget at its center. Focus on typical utterances. When an utterance is made, it is made by someone, the speaker of that utterance. This fact is independent of what sentence is uttered, and it is a general fact about all utterances. To see this, note that most utterances are not about the speaker; we say many things, and most are not about ourselves. But *we* say them, and we are the speakers. Similarly, an utterance is made at a time, whether it concerns issues related to a location in time, or features of natural numbers. Indexicals are the words which take advantage of these facts; if “I” is used in an utterance, it will pick out the speaker of that utterance, but the property of having a speaker is shared by all utterances, whether they contain a use of the first person pronoun, or not.

Demonstratives differ from indexicals, in that not all utterances have demonstrata. Consider an utterance of “two times two equals four”. It’s an utterance, so it has a speaker, it is made somewhere, at some time, by someone, to someone. But it has no demonstratum. In fact, the question “what is the demonstratum of that utterance” is inappropriate; it is not that kind of utterance. What kind of utterance has demonstrata? Typically, utterances with demonstratives in them. If I say “that’s interesting”, you can ask: what is the demonstratum? And you can ask it, because typically when we use a demonstrative, it has a demonstratum. Demonstratives, then, are words which depend on contextual features that occur only when the sentence uttered contains demonstratives. These features are not common to the contexts of all sentences. In contrast, the contextual features that provide the referents for the indexicals are features of all con-

texts of utterance, and are independent of whether the relevant indexicals occur in the utterance.

I find this distinction intuitively compelling, but in order to make it clear, and try to convince the unconvinced, I will have to leave our discussion aside for a bit, and begin elsewhere: the metaphysics of real world utterances.

2.3.1 Robust Context Features

Kaplan tried to solve the following problem: take a sentence with a deictic in it; how does it get a truth value? As we saw earlier, Strawson and Quine claimed that the meaning of these kinds of words cannot be captured formally, since they keep changing their referent, without changing their meaning. Kaplan's reply was that deictics get their referents in a disciplined way from the environment, which, as we also saw earlier, had already been noticed by Frege in 1918 (e.g. "I" refers to the speaker), and this allows them to be treated formally. All we need is a formal representation of this environment, a context of utterance, and we can deal with the problematic expressions. This allows us to think of evaluating sentences in any context, whether or not the person designated as the speaker really was uttering that sentence, or even speaking, at that time and place, etc. Once we have the notion of a sentence evaluated in a context under control, we can start thinking about real world utterances, since they are what we represent formally as sentences in contexts.

I would like to start backwards, from Lewis's notion of a context: think first of utterances, i.e. sentences being used by actual speakers, with actual intentions, audiences, goals, etc. In particular, I will focus on declarative statements, such as Obama's saying "I am happy". Usually, these have exactly one speaker, one (or more) addressees, happen in just one place, at one particular time, etc. None of these features is essential; none is a criterion establishing what counts as an utterance. Some couples finish each other's sentences, so there is more than one speaker per utterance. Sometimes we begin saying

something in one place, and mid-performance we move to another place, and this may matter for what we say. Sometimes it takes us quite some time to say something (imagine saying a long sentence slowly). Sometimes we begin a sentence by addressing one person, and end it by addressing another. Even for the notion of a goal there is variation: we can begin to say something with one communicative goal in mind, decide it's not feasible, and switch to a different goal, all while still working on uttering that one sentence. There are messages in a bottle, where we want an audience, but have no idea whether we will succeed in having one or more people read our message. These are not otherworldly, *recherché* examples. They happen often, and we expect them to happen; sometimes, we depend on it. Even the message in a bottle is not exceptional; every philosophy paper ever written is like it. So any theory of communication needs to deal with each of them at some point; if it cannot, that counts as a *modus tollens*.

Still, the right way to start is with a simple utterance: one speaker, one addressee, one time, one place, etc. They are the prototypes, they count as typical, and, I will claim, they shape the structure of our utterances. To see this, note that there are two kinds of features of contexts: robust and fragile.

Robust features are the ones which an utterance has just in virtue of being an utterance.¹⁵ For instance, an utterance has a speaker; not because of what is said, but because that's how utterances generally work. Again, not always: there may be more speakers for one utterance; it's harder to think of utterances without a speaker, but perhaps with some imagination one could do it. But it always makes sense to ask: who said that? Even if the answer is sometimes complicated.

This generalizes to the other robust features mentioned above: it makes sense to ask where was that said? When? To whom? The latter shows that having one (or possibly several) addressee is as robust a feature as having a speaker: typically, we say things *to someone*, even if only to ourselves. And all this is independent of *which* sentence is

¹⁵I borrow the tag "robust features" from Crimmins (1995, p. 393).

being said. These are facts about the utterance *qua* speech act, since we can also ask corresponding questions about promises, orders, requests, thanks, questions, etc.¹⁶

This is a claim about metaphysics, and not about semantics. There is nothing in the semantics of “here” that forces each utterance of a whole sentence to happen in a particular place. The utterance may happen in several places, as when the speaker is moving about. And perhaps before God created space, he could have used “here”, as a kind of trial run. He would have known, of course, that it is empty in that usage, but so what? Perhaps, for once, the goal was not to speak the truth. Semantics doesn’t preclude any of this.

Or perhaps metaphysics does prevent at least emptiness; perhaps my scenario is impossible, and God could not have used “here” if there was no space to use it in. And perhaps angels, when they lack a location, cannot use “here” as an indexical. These metaphysical claims are stronger than the one I am making; I’m just talking about what is typical for utterances. To be sure, what is not typical needs to be addressed as well. But, to make a methodological point, it seems to me likely that in order to explore the semantics of an expression, one starts out with its typical uses, where the expression is at its most comfortable.

2.3.2 Fragile Context Features

I want to contrast robust features with fragile ones. In order to do so, I first need to make one more clarification, regarding three uses of the notion of intention in the theory of communication.¹⁷ First, when someone uses a singular term, it is typically accompanied by the intention to refer to something. This cuts across all lexical categories: it’s true

¹⁶I don’t claim that they can be asked of any kind of speech act. Namings of ships, for instance, don’t typically have an addressee: “I name this ship XYZ”. In some traditions, namings of people do have an addressee, though: “I baptize you...”.

¹⁷Henceforth, I will only talk about the speaker’s intentions. If other participants’ intentions are also relevant, similar points could be made about them.

of proper names, demonstratives, indexicals, etc. These intentions as to the topic of the forthcoming discourse are pre-semantically relevant, in the sense that they inform the speaker's choice of words, and their object is what is about to be said, but they do not necessarily provide any input to the semantics or the truth value of what the speaker ends up saying. The pre-semantic use of intentions in a theory of communication is to be contrasted with a second use of intentions, the semantic use of intentions, those intentions that play a role in the mechanism of reference of a certain expression. Once the word is used, only the semantic intentions are relevant for matters of reference and truth value. Finally, a third use of intentions: some intentions are only post-semantically relevant, in the sense of Grice ([1975] 1989b): the speaker may intend to communicate more than what is expressed by his utterance. For instance, when we say "Could you pass me the salt?", we are directly asking if the addressee is able to perform a certain action, but the question is normally taken as a request for a passing of the salt. These intentions also inform the choice of words in uttering a sentence, just like pre-semantic intentions. I separate them because Gricean intentions are not related to the semantics of the sentences. Most pre-semantic intentions will not contribute to the semantic values of the expressions used; but these intentions concern the semantics nonetheless. In any case, our focus will be on pre-semantic and semantic intentions, so the third category will be largely left aside. For the sake of brevity, I will speak of pre-semantic, semantic, and post-semantic intentions, which makes it sound as if there are three mutually exclusive classes of intentions. I think it is not implausible that the same intention can play several roles, but the issue of individuating intentions is too complicated to try to settle it here.

I take it as obvious that there are pre-semantic intentions which provide no semantic input. If I want to talk about the planet Mercury, but all I say is "Always recognize that human individuals are ends, and do not use them as means to your end", unless I am talking in code, I cannot be talking about Mercury, except in the trivial sense that I intend Mercurians to obey that law as well. I also take it as obvious that there are post-

semantic uses of intention. If Grice is right, my goal-related communicative intentions are particularly relevant to implicatures.

It is not, however, obvious that there are semantically relevant intentions. Bach (2005), for instance, claims that the boundary between semantics and pragmatics is drawn by the presence of intentions: when they come into the picture, we're in the midst of pragmatics. So let me give an example of a dispute which can be put exactly in these terms.

Donnellan (1966) and Kripke (1977) disagree precisely about the role of intentions in the use of definite descriptions.¹⁸ Here is a typical case, adapted from Donnellan (1966): two people are just getting to a party, and one of them, Paul, sees someone in the distance, Oliver, who looks a bit drunk. In an attempt to make some small talk, Paul says "the guy drinking a martini is tipsy". He chooses that description because pointing is rude, and because he sees that person drinking from a martini glass, half filled with a transparent liquid. Suppose now that Oliver, the particular person Paul was focused on, was drinking rum out of a martini glass, but was indeed tipsy; suppose further that at the party there was only one person at that time drinking a martini, Pip, who was not visible to Paul (he was in another room), and was sober. To recap: Paul is the speaker, Oliver is tipsy and drinking rum out of a martini glass, and Pip is drinking a martini, but is not tipsy. Here's the issue now: was what Paul said true, or false? The choice depends on whether you think that "the guy drinking a martini" picks out Oliver or Pip, since Oliver is tipsy, and Pip is not. And there are reasons for both choices: on the one hand, Pip is the only person who satisfied the description; on the other hand, Paul clearly meant to talk about Oliver, would have had no reason to say anything about Pip, and the addressee wouldn't have thought of Pip either.

¹⁸Kripke makes a similar case about proper names. Donnellan doesn't ever clearly discuss proper names. His Aston-Martin example in Donnellan (1970) comes closest, but his discussion is complicated, and the moral he wants to draw from it is not related in very simple ways to the issues we're discussing. Proper names would have been more handy for me, since they are undisputedly singular terms, just like deictics. If this turns out to matter, the discussion could be switched to proper names, without attributing to Donnellan the view that Kripke is attacking. This last move is made perhaps a bit easier by the fact that Almog (2012) argues that Donnellan ought to have said the same things about proper names and definite descriptions.

Simplifying, Donnellan (1966) argues that Paul said something true about Oliver, in spite of Oliver not satisfying the description he used, and thus some uses of definite descriptions pick out referents by some other mechanism than satisfaction of the description, and Kripke (1977) argues that Paul said something false about Pip, because Pip satisfied the description.¹⁹

Donnellan and Kripke agree about Paul's pre-semantic intentions: he wants to talk about that guy over there, and, since he thinks that guy is drinking a martini, he also wants to talk about the martini drinker.

They disagree about Paul's semantic intentions. Donnellan claims that Paul's use of "the guy with the martini" refers to Oliver, because he intends to talk about a particular person in his field of vision, whatever he's drinking, and this intention is semantically relevant because it trumps the fact that the description picks out Pip. So what Paul said is true, since Oliver was tipsy. If Donnellan is right, Paul's intention to talk about X influences the semantics of definite descriptions, at least some of the time.

Kripke, on the other hand, claims that "the guy drinking martini" is about Pip, not Oliver, as a matter of semantics. He admits that, besides wanting to talk about the martini drinker, (which is why he uses that description), Paul also wishes to talk about that guy over there, who is in fact Oliver. The addressee may also figure it out, and respond as if Paul's remark was about Oliver. Still, Kripke claims, Paul's Oliver-intentions (which he calls "speaker intentions") play no semantical role; what he said is false, since Pip is not tipsy. In fact, Paul's Pip-related intentions also play no role in Kripke's account; he uses the description, which does what it does, i.e. it picks out Pip, whatever Paul wants from it. So his pre-semantic intentions play no semantic role; they are just part of a story about how he came to use that particular description. According to Kripke, only the speaker's

¹⁹I use "about" loosely, since Kripke does not claim that definite descriptions are singular terms, and it is compatible with Kripke's views that Russell's theory of descriptions is correct, so that the proposition expressed does not contain the satisfier of the description. I mean "about X" in the sense that the truth value of what was said depended, in some way to be specified, on facts about X. This notion of aboutness ought to be neutral enough to be acceptable to both sides.

general intentions to use the words with their usual semantics matter semantically, and these intentions clearly play no role in the semantics of those words.

The disagreement between Donnellan and Kripke is delicate. Donnellan would agree with Kripke about the claim that Paul intends to use the words in their usual semantic capacities, not as part of some special code. They part ways with respect to the question whether referential uses of definite descriptions are normal, usual, customary, semantically respectable uses.

For my purposes, it doesn't matter which view is right about definite descriptions. But these two roles of intentions are relevant to deictics. Start with indexicals. When Obama says "I am happy", he pre-semantically intends to talk about himself. He could have done it in various ways; he could have said "Obama is happy", or "The President of the USA is happy", or even "This guy is happy", pointing at himself. The choice matters semantically, and besides, the implicatures will be quite different. But each of these sentences would be true, and would in fact be about himself. It should also be clear that his saying something is an intentional act. You count as the speaker if you are the agent of that utterance, if you're talking in your own name. But once he says "I", all these intentions are irrelevant. Even if he thinks he is David Hume, and he pre-semantically intends to speak of David Hume, he is not referring to David Hume; "I" picks out the speaker, and that is that.²⁰ On indexicals, then, I stand on Kripke's side: we pre-semantically intend to use indexicals with their customary meaning, and once used, they refer to whatever they refer, without a care for our intentions. Here is my generalization: intentions play no semantic role for indexicals. As we saw above, most philosophers agree that intentions do play a semantic role in uses of demonstratives.

What does this have to do with fragile features of contexts? Here's the idea: indexicals work off the robust features of contexts. Some robust features are intention-independent; for instance, utterances happen in a certain place, at a certain time. Our pre-semantic

²⁰I borrow the argument, and the moral, from Perry (2001, p. 67).

intentions are relevant here. Speech acts are intentional acts, and part of that intentionality is our planning when, where, and to whom to speak. I can wait until we are alone to give you the good news that I know you'll want to keep from envious ears. During a Q&A, some audience members stand up, so others can see who is speaking, and hear them better. But, again, once the utterance is made, those plans are semantically irrelevant. I speak when I in fact speak, and that's what gives the semantic value of "now", whether I'm right about being alone with you or not; the questioner is the person who asks the question, and that is what gives "I" its referent, whether others hear her or notice her or not.

Other robust features depend on pre-semantic intentions, like wanting to address someone. But once in place, indexicals just get their referent automatically from the context. The relevant intentions, if there are any, are not *semantically* relevant. Indexicals are those words which take advantage of robust features of contexts.

The case of "you" is a special one, because it is hard to distinguish:

- having an intention to refer to X by using "you", and
- having an intention to address X, and then using "you" to refer to X.

Hard, because planning to address X resembles planning to talk to X over there, where nobody will overhear, and I said that the last kind of intention is not semantically relevant. Hard, because my intention to address X does not by itself determine the referent of any indexical; I must also succeed in addressing X, and one might confuse this fact with the claim that the intention is semantically inert. Hard, but not impossible.

To see how special "you" is, let's look at demonstratives. Here, we see that some intentions do play a semantic role, and a sentence with a demonstrative in it needs these features, these intentions, to exist in order to get a truth value. They are fragile features, because they are instituted by intentions that do semantic work. Obama's intention to speak is not relevant to an account about how his use of "I" gets its referent; but his

intention to pick out something in particular when he says “This is interesting” plays a semantic role because the demonstrative works directly off the intention. This directness is one way to see what sets “you” apart from demonstratives. An occurrence of “you” picks out its referent in virtue of that referent being addressed, not just in virtue of its being the intended referent.

Here’s an example. Suppose I am talking to X, but wish to refer to Y and then Z. I can address X without ever wanting to refer to her, but if I do wish to do so, “you” is ready for action, just as “I” can be used to refer to me. While talking to X, I cannot use “you” to refer to Y or Z. For that, I would need to be addressing them, in turn. While talking to X, I have many other options: I can use the demonstrative “he” to refer to Y, and I can use Z’s name to refer to him. But I cannot use “you”. That is why the intention to refer to X by using “you” is not relevant. I normally address someone not in order to talk *about* them, but to talk *to* them. If I want to do both, I can, using “you”. What I cannot do is talk about X using “you” without addressing them. And failures are caused by addressing the wrong person, not by intending to refer to the wrong person.²¹

“You” latches on to the addressee, and it needs no help to do it. Demonstratives are more malleable: I can refer to Y by using “he”, whether or not Y is part of our conversation, or anywhere near it. But this malleability comes at a cost: they need the help from the speaker’s intentions to get a referent. Furthermore, these intentions are token-specific: I want this particular use of “he” to refer to Y, and this other to refer to Z. I can do this at will, since the intentions are token-specific for demonstrative uses of pronouns. I can also address several people in turn, and use “you” to refer to them in turn; but I cannot do the latter without doing the former.

This issue comes up, for instance, in one of Kaplan’s discussions of Donnellan; in a footnote, he asks rhetorically:

²¹Addressing can be more or less ‘in person’, as when I say “Oh Kant, why didn’t you use shorter sentences?”. I count this as a case of me addressing Kant, even if he doesn’t know it.

I may have you in mind, and believing that it is you whom I see hiding under the bed, begin berating you. Even if it was not you under the bed, might it not still be you whom I criticized? ([Aft], p. 583, fn. 36)

From the main text, it is clear that this is intended as an intuition pump, and that we are meant to agree that it would be you who is being criticized. But this is wrong: you're not there to be addressed; I tried to address you, and failed. You are not the referent of my use of "you", because you're not there for me to address you. If someone else is there, I am addressing them. If there's no one under the bed, I am not addressing anyone. In neither case am I addressing the absent you.

Not noticing these facts made Kaplan count "you" as a demonstrative. He was looking at the following example:

(15) You, you, you, and you can leave, but you stay.²²

The surrounding discussion is of the fact that different occurrences of the same demonstrative in the same sentence may not corefer, and the question is how to fit that into his theory. Here's his proposal:

Such cases seem to me to involve an exotic kind of ambiguity, perhaps unique to demonstratives [...]. Where different intentions are associated with different syntactic occurrences of a true demonstrative, we would want to use distinct symbols in our formal language in order to avoid equivocation. [...] In [cases like (15)] in which there is simultaneous perception of all addressees, I think it correct to say that there are several distinct, simultaneous, directing intentions, indexed to distinct intended utterances of the demonstrative "you" (which are then voiced one at a time). ([Aft], p. 587)

²²[Aft], p. 586.

We will come back to the referent change issue in §2.7.2. For now, I will focus on Kaplan's idea that "you" gets to have this variability because each occurrence has the benefit of its own intention, and thus the possibility is open for each occurrence to have its own referent. This is exactly what I deny: it's not our intentions that make this possible; it's the fact that we can change whom we're addressing several times while uttering a single sentence. The intentions that are relevant for "you" do not provide a referent for the word directly; they merely provide an addressee.

As I said, it is hard, because one might be confused by the fact that at least some uses of demonstratives also have other conditions for success than merely the existence of an appropriate speaker's intention. When the speaker hallucinates, her attempts to pick out things that do not exist will fail. Her intentions are thwarted by reality. In the case of "you", the speaker will also be unable to use it to pick out the hallucinated rabbit in the corner. But the failure happens via the failure to address the rabbit. This intermediary step does not exist with demonstratives.

A word about nomenclature. I call some features "robust" because they set up the scene, just in case something is about to be uttered. They are like theater stages, which are not normally purpose-built for a particular play. The other features are "fragile" in the sense that they are only needed for particular kinds of sentences. You don't need a demonstrative intention if your sentence doesn't need it. But you need to make your utterance in a particular place, because that's how utterances are made. And you cannot felicitously use a demonstrative in the absence of a guiding intention.

I should also note that I don't mean that only utterances with demonstratives in them require demonstrative-type intentions. As Stojanovic (2007, p. 14) noted, if Obama had said "Impressive!" while looking at a painting and intending to make a comment about it, what he said is intuitively as much about that painting as if he had said "That is impressive". How exactly to handle such types of sentences is beyond the scope of this

chapter.²³

2.4 A Bit More about the Role of Intentions

I said that I wouldn't propose a view about the mechanism of reference for demonstratives, and I am resolute on the matter. But if the picture proposed above is correct, we can already draw some conclusions about the nature of the relevant intentions. To explain what I mean, I will sketch a very simple picture of the way in which we use deictics. One can think of this as a fragment of the picture provided in Korta & Perry (2011), but their goals are broader than mine: they talk about all aspects of communication, while I talk merely of reference. I will explain the differences shortly. But I need to mention immediately that I don't mean to imply that the two projects should be treated independently. It's just that for my narrow interests, I think it is possible to do so.

Some utterances are purely general; they aim to be about no thing in particular, and they succeed. I want to focus on the ones which aim to communicate singular thoughts. They all share a certain kind of intention: to refer to some particular thing.²⁴

Different kinds of linguistic tactics can be used to achieve that goal. One thing we can do is to use a proper name, if the intended object has one. If Kripke is right, once we select a proper name which we believe to be borne by that object, our job is done; the name takes over, and it refers to its bearer, with no need for further guidance, and sometimes against our wishes, if the intended object does not, in fact, bear the name.

We can also use an indexical. The speaker's reasons to do so are varied: the object in question may not have a name (uses of "here" in unremarkable places); or it would be strange to use a name (most uses of "I" and "you"); or we don't know its name (if calendar dates are names of days, then some uses of "today" fit the bill; that is the point

²³See, for instance, Stanley (2000) and Recanati (2002), for two opposite views on the matter.

²⁴Perhaps even this is a bit too strong; recall my remarks about Santa Claus.

of the question “what day is it today?”). In spite of this variety, I argue that these words all have something in common, namely they depend exclusively on robust features of the context to pick out their semantic value, i.e. features which an utterance has *qua* utterance. For some, the speaker’s intentions play some role (the addressee is addressed, at least in part, because the speaker wants to address her), for others, it does not (the current day is current all by itself). But none are established for the sake of the interpretation of some expression, and, typically, an interpreter is supposed to know them, again, not in virtue of parsing the sentence correctly, nor by attending to what is said, but merely by observing the speech act.²⁵

Perry and Korta are fellow travelers on this road, because they also start from utterances, and talk about roles that things have in making an utterance. But on this stretch, they travel backwards. Of a sentence like “I am Obama”, they seem to claim, as a basic semantically relevant fact, that the speaker’s use of the indexical is what gets the hearer to think of him first as the speaker, of whom it is said that he is identical to Obama, thought of under the Obama-guise.²⁶ Suppose the sentence is true. Their story would nicely explain certain kinds of Frege-type puzzles. The difference between “I am Obama” and “Obama is Obama” would be that the first sentence would occasion two ways of thinking of Obama, whereas the second would be trivial, in Frege’s sense.

But Perry and Korta get the semantic facts wrong. The hearer already thinks of Obama as the speaker of the utterance, because he *is* the speaker of the utterance; the use of “I” is just piggybacking on that independent fact. It may be true that the job of “I” in the sentence is to activate this way of thinking about Obama, in order to supply the hearer with an informative identity fact. But the job depends on facts which pre-existed the use of the word, and the use of the word depends on them. Their mistake, in my terminology,

²⁵I say “typically” because there are exceptions. The world does not guarantee that the interpreter will be able to tell by observation who is being addressed, and similarly for who, in a crowd, is speaking. This claim has been disputed for demonstratives by people who claim that the referent is that which an ideal interpreter would take to be the referent. See, for instance, Wettstein (1984) and King (2012).

²⁶I am talking about Korta & Perry (2011, §6.3); an example like mine is discussed on p. 66.

is to miss the importance of robust features of contexts in the workings of indexicals.

This difference in persistence between the two kinds of features would be well served by calling fragile ones “fleeting”. These features are not fragile in the sense that they can be broken off from the use of the word. On the contrary, fragile, or fleeting, features are tied to the use; the expression doesn’t get a referent without them, and the intentions lack an object if the use doesn’t get realized. They are fleeting, in the sense that they are momentary; they appear for that particular moment, and have no power outside of it.²⁷ Robust features are independent of the words being used, and even of whether a given language contains any particular indexical; their role is related to the general functions of utterances, not to the contents thereof, and therefore their role in the semantic story for indexicals, though important, is not essential for their accomplishing their more general role in structuring a context.

Here is another way to tell the story of indexicals. Suppose the speaker wants to refer to himself, and uses the word “I”. I find it undeniable that he intends to refer to himself by that use. But that is not the semantically basic intention. We first have the intention to speak, and then the act of speaking. That makes Obama the speaker. Since Obama is the speaker, he can be referred to by the speaker by the use of the word “I”. Which he does. So the intentions which matter semantically are the intention to speak, and the intention to refer to the speaker, realized in the use of the word “I” (and, if you wish, the intention to use the English pronoun “I”, not the English pronunciation of the symbol for iodine). They are all 2-place intentions.

Demonstratives require 3-place intentions: the speaker intends to refer to Y by a particular use, a particular token, of a particular demonstrative.²⁸ And these 3-place intentions are the semantically relevant ones. The hearer thinks of the demonstratum in a

²⁷We can use anaphora to pick up their referents later on in the conversation. But by then, the fleeting intention has done its work of securing a referent, and the anaphora merely makes use of that achievement.

²⁸Unless there is no semantic difference between “this” and “that”, in which case, for English speakers at least, one needs only intend to use a particular demonstrative, the choice of which being irrelevant. More about this in §2.8.

certain manner because it is the demonstratum, i.e. because it is referred to by that use of a demonstrative. And the speaker's intentions to communicate do not work on a way of thinking about something already available to the hearer qua participant in the conversation. The demonstratum might already be thought of in a similar manner by the speaker and the hearer, as when two car buffs see a Ford GT at a stop light, and one says "That's the first blue GT I've seen". But that is just a fact about their interests, their positions, etc. It is not intrinsic to the utterance itself; in fact, in my example, it precedes the utterance.

A final note about demonstratives: ever since Kaplan introduced the notion of character, there has been a debate whether demonstratives have them.²⁹ Character is the rule by which, for instance, indexicals get their referent from the context. Clearly, demonstratives as word types do not have a rich enough stable meaning to pick out a referent from an arbitrary context. Demonstratives, of course, are not rule-independent. We use them knowing how they work, knowing that they are sensitive to the speaker's intentions, even though philosophers have been struggling to figure out what exactly the rules are. We could say that the character of demonstratives as types is such as to pair them with either the relevant intention, or directly with the demonstratum. Or we could claim that demonstratives have no character, a notion more at home within the realm of indexicals. I mention this debate only to put it aside. The outcome matters, but it doesn't serve as a more fundamental feature of demonstratives than the ones we've been talking about. Fundamentally, this is a debate about the proper place of intentions in semantics, and, even more generally, about the goals of semantics and how to distinguish it from pragmatics. Applied to deictics, it is a debate about the proper place of indexicals and demonstratives. I have talked as if demonstratives must be treated semantically; I think that's right, but an independent issue. If I am right, I have differentiated indexicals from demonstratives properly. Where the difference belongs in our overall theory is another

²⁹Kaplan, for instance, changed his view on the matter. In [Demos], he thought that demonstrations give a character to each occurrence of a demonstrative; in [Aft], he became unsure how to talk of characters in the case of demonstratives. For discussion, see [Aft], §II.

matter.

2.5 From Utterances to Sentences in a Context

So far, I have talked about utterances, their metaphysics, and the place of context dependent words in the semantic scheme of things. But there are two paradigms of semantics in the literature, and only one of them matches my proposals so far. On this conception, semantics is fundamentally about sentences when used; until a sentence is used, talk of reference, truth, or meaning is confused. This tradition began with Strawson and Austin³⁰, and is still alive and well.³¹ Using an overly broad brush, and paraphrasing Austin, this is a tradition that aims to explain what people do with words. As such, its purview includes many things besides what words people use, such as their intentions, goals, beliefs, etc. Language is a tool with many uses, and consequently any such theory needs to explain many, if not all of them: it is used to assert, to question, to communicate, to imply, to refer, to lie, to misguide, etc. Perhaps not all of these belong to semantics proper; perhaps some things are relegated to other disciplines, like pragmatics, or perhaps sociology. But the basic idea is that language is a tool for communication. So any theory is tested by its accord with intuitions about what passes between participants in a conversation.

The other tradition I have in mind has just as many contemporary co-promoters, all of whom had complicated views, and therefore, upon inspection, none of whom may

³⁰See, for instance, Austin (1950) and Strawson (1950). Their views were in fact more complicated than I make them sound. For instance, both Austin and Strawson thought that truth bearers are statements, but were open to assigning truth to sentences on occasions in which they were not actually used. This makes it quite complicated to figure out what exactly is the relation between their notion of a statement and the more recent distinctions between sentences, sentences in a context, and propositions. I leave this issue aside here.

³¹The idea of utterances as necessary for deictics to get a semantic value is obvious in token-reflexive theories, as proposed first in Reichenbach (1947, pp. 284ff.) and, independently, by Burks (1949). This point is clearly made in Perry (2001, Ch. 3), who also defends a token-reflexive theory.

quite fit the mold: Frege, Russell, Kripke.³² This tradition is much narrower in focus. Semantics is seen as an analysis of language, seen as a tool for representing the world. Which is not to deny that language is used for other purposes; but the idea is that all those can, or, indeed, must be separated from this special, and perhaps primary, function. One can then go on to offer an account of communication: by representing the world in a certain way, the speaker hopes that the hearer will understand the representation, and, with some aid from a theory of cooperation, the hearer may take up what was said. But all that comes down the road.

That some idealization is needed is accepted by all parties.³³ Speaking just of indexicals: suppose it is a metaphysically unavoidable fact that no speaker can pronounce 100 sentences at the exact same time. If we really only looked for real utterances, then we could find no instances of an argument with 100 premises. This is a fact that logic ought to be blind to: such an argument ought not to come out valid merely because there is never an occasion to utter those premises truly, even if we could infer a contradiction. Still, the two places that semantics may begin from will lead to different concerns and theoretical desiderata.

In this section, I want to show that my criterion can be made to accord with the strictures of the second tradition. I will formulate the issue using the contrast I mentioned above between Lewis and Kaplan. Lewis, if you recall, thought of contexts as surroundings of real utterances. Kaplan wanted a more abstract notion: of a sentence in a context, where the context may not contain any utterances of that, or any other sentence. The goal is to explain the intuition that we can talk not only about what is said by someone at

³²I include Frege more because of his influence than his actual views. It would be more accurate to talk of him as thinking that thought represents the world, and language represents thought. If representation is transitive, his inclusion may not be so bad. See also Locke, *Essay* III. i. 2, p. 402 “Besides articulate Sounds therefore, it was farther necessary, that he should be able to use these Sounds, as Signs of internal Conceptions; and to make them stand as marks for the Ideas within his own Mind, whereby they might be made known to others, and the Thoughts of Men’s Minds be conveyed from one to another.”, quoted from Locke (1975). I owe the citation to Lewis Powell.

³³García-Carpintero (1998) correctly argues that this fact can be accommodated by a token-indexical view, properly construed.

some time and place, but also about cases in which nobody says anything. For example, we think of the sentence “I am awake” as false in a context which has me as the agent, but I am sleeping. Obviously, I am not uttering anything; and nobody else could utter that sentence for me, since they would be talking about themselves, not me.³⁴

I started from utterances because that is where language lives. But I don’t want to make this a deep point about the proper home of semantics. Rather, it is a methodological point: words have the meanings they have because we use them in certain ways. Every semantic account of English will be tested by its fit with the way English is actually used. This point has never been denied. But indexicals bring out an unexpected twist: they depend on what I call robust features of utterances. So we start out by thinking about utterance contexts, and then we see what gets stuck to words such that we can abstract our way to sentences in a context. In order to get to the latter, we need to find a replacement for the direct link to robust features, since, obviously, it is not true that sentences in just any old context usually have an addressee, and even though contexts always have agents, there is no requirement that the agent be an utterer. Applying this to “you”, we need to separate the way in which X gets to be the addressee (a matter of actual intentions, which accompany actual utterances) from the way “you” gets its semantic value in a context where X is the addressee. This move from utterances to sentences in a context is necessary in order to avoid having “I am asleep” never come out true, which it would if we were only looking at utterances: presumably nobody utters anything while asleep, and talking in your sleep does not count as an utterance in our sense.

The key is the distinction between the two types of intentions that are semantically relevant for deictics. Indexicals work off robust features of contexts, some of which depend on the speaker’s intentions, e.g. to address a particular person (recall that we are in the realm of the metaphysics of typical utterances; exceptions do not prove the rule). But the words do their work by themselves; and this allows us to move from the more grounded

³⁴This is a simplified version of the arguments in [Aft], p. 584.

notion of an utterance to the more ethereal, sentence in a context. So, for instance, just as we can think of me as the agent of the context that has me asleep (and thus, not really an agent of an utterance), and thus make “I am asleep” true, and, more importantly, “I am awake” not a logical truth, so we can think of “you are asleep” as true in a context having my neighbor as the speaker and his daughter as the addressee, even when they are both asleep. In the last context, my neighbor is not the agent of an utterance, and his daughter is not being addressed by anyone.

Thus, we can separate the role of intentions in setting up robust contextual features from the functioning of the words (i.e. word types) relative to a context. This allows us to have a notion of an indexical relative to a context, and bracket away the issue of the mechanism by which things get to have their roles in that context (e.g. how my neighbor gets to be the speaker, and his daughter the addressee).

Demonstratives are different. In their case, the intention directly connects the token occurrence of the word with its referent. Therefore, we cannot repeat our indexical strategy, and talk of things as being demonstrata in contexts in which no utterance is made, and therefore no intentions are active, since we cannot get a grip on the token of the word and connect it to a referent unless the context contains an intention to do the job. But that’s exactly what’s missing, absent a real utterance. Perhaps this is too strong. Imagine that I want to say “This is mine”, but abstain at the last moment, thinking it impolite to say that. An intention is formed, and maybe there is room for talking about the truth value of that sentence in that context, even if no utterance took place. And it might not even have to be that very sentence; maybe we can also evaluate the sentence “this is blue” in that very context, since my intention with regard to “this” doesn’t have anything to do with how the sentence continues. This would fit nicely with the idea of having one context per word, not per sentence.

So there is some wiggle room between real world utterances and what is needed for demonstratives to get a semantic value. But in all these cases we had real contexts, with

real intentions, and we didn't get anywhere near contexts where the speaker is asleep or doesn't speak the language. So I find it much harder to get a satisfactory notion of a sentence in a context for demonstratives. This is not a comfortable conclusion. Imagine trying to construct a logic of demonstratives, where you were only including real contexts, with a speaker actually intending to refer to things. Kaplan's original worries would come back to haunt us, because we are within striking distance of unwanted logical truths, such as "I am thinking", or "I am awake", or "I speak English".

The bottom line is that according to my distinction, we can easily get a notion of sentences in a context for indexicals, but I don't know if that is possible for demonstratives. We could add sequences of demonstrata to the context, and pair them with the corresponding sequence of demonstratives. But that is an artificial solution; demonstratives do not work like that. One possible conclusion might be that utterances really are the only level of analysis appropriate for demonstratives. Another option might be to take demonstratives as some kind of super-Millian proper names. *Super-Millian* because for proper names we are used to bypassing the issue of the fact that, e.g. there are many people whose name is "John", so that sameness of proper name in an argument guarantees sameness of referent. But we have fewer demonstratives than proper names; otherwise, they wouldn't need to be context dependent. All we are left with are the referents; so our logic for demonstratives would get closer to a logic on propositions than on sentences. Both options have clear disadvantages, but nobody said truth was going to be easy. I will come back to this in §2.8.

2.6 The Tradition and "You"

The time has come to argue for against competing criteria for distinguishing indexicals from demonstratives. "You" is in some ways like prototypical indexicals, and in some like prototypical demonstratives, which may explain why it sometimes gets counted as one,

and sometimes as another, even by people who hold the traditional conception. In this section, I will show that my picture differs from the tradition not just conceptually, but also extensionally: it counts the singular second person pronoun “you” as an indexical, not as a demonstrative.

It is easy to see that “you” is a problematic case: it is counted among demonstratives in [Aft], pp. 585–590,³⁵ Braun (1996, p 171, fn. 12), and García-Carpintero (1998, p. 532), and among indexicals in Salmon (2002, p. 519) and Caplan (2003, p. 208, fn. 1). Even more baffling is the fact that most of the time, when it appears on a list, it is included there without any arguments, as an obvious addition to the canonical list. But it can’t be obvious that it is both an indexical and a demonstrative, especially since an important reason to make these disjoint lists is to capture the contrast between the two types of words.

The traditional view is that demonstratives are the deictics which work off the speaker’s intentions. “You” does: it is up to the speaker to decide who the addressee is, whereas it is not up to the speaker to choose the possible world where he makes the utterance (except in the trivial sense that he may choose not to utter anything). Of course, the speaker is not all-powerful. If he is hallucinating a big white rabbit, and addressing it, he is in fact addressing nothing and nobody. But normally, one addresses the person one intends to address.

This contrast, however, is not as clear cut as it seems. After all, the speaker also decides when and where to speak, and yet nobody has been inclined to think that “now” and “here” are demonstratives for that reason. The point is that the speaker’s referential

³⁵There is little talk about “you” in [Demons], though a tantalizing remark on p. 552 seems to indicate that Kaplan might have contemplated counting it as an indexical: “Aspects of the contexts other than [agent, place, time, and world] would be used if new demonstratives (e.g., pointings, You, etc.) were added to the language.” His usage of “demonstrative” in the original paper is not very tightly regimented, as it is sometimes used synonymously with “indexical”. What is more important is that something like the addressee could be considered part of the context. Then again, it is hard to figure out why pointings would count as on a par with “you”, as expressions of the language, so we should not put too much pressure on this passage.

intentions do not matter for “now”, but they do play a special role for “you” and “this”. Here’s another way to look at it: if you ask the speaker whom she is addressing, she would be in a particularly good position to answer. One can also ask the speaker when she is (or was) speaking, but in this matter, the speaker is not in a better position to answer than anyone else around who has a good watch. Onlookers may also opine about who is being addressed; but they need to guess something about the speaker’s intentions, which is why she is in a special position.

So if one is to hold on to the traditional view, I think that it is a natural choice to count “you” as a demonstrative. But the traditional criterion doesn’t quite force this choice. As we saw, the idea is that the stable word meaning of demonstratives just isn’t sufficient to secure a referent; it is not rich enough. “You”, however, is quite similar to “I” in this respect: the former picks out the addressee, the latter picks out the speaker. Their riches seem equal. So the feeling that some supplementation is needed for uses of “this” to help out its meaning is lacking for “you”. This is why some philosophers have counted it as an indexical: the traditional distinction just doesn’t cut finely enough when it comes to “you”. In the next section, I will in fact discuss a few other criteria that have been proposed, even if the authors didn’t realize they were different from the traditional conception.

Before we look at alternative views, here is what I claim, and what my theory entails: “you” is an indexical. The addressee is a robust feature of contexts: just as utterances are made at a certain time in a certain place, they typically are made to communicate something to someone.³⁶

To illustrate, note the difference between these sentences:

(16) You owe a dollar to him and to her.

(17) He owes a dollar to you and you.

³⁶Well, someone or ones. Plural “you” brings in problems of its own, which I leave for another time.

All the third person pronouns in these examples are demonstratives; the speaker must intend to refer to certain people, who need not be participating in the discussion in any way. But for each occurrence of “you” to get a referent, particular people must be addressed, with all the accompanying requirements: they have to be available for being addressed, and they will count as participants in the conversation.³⁷

Just as before, I’m not talking about essential features, because there are utterances which lack an addressee. Indeed, there are at least two ways that can happen. One is the hallucination case: the speaker intended to address someone, but failed. This is similar to the use of empty names by speakers who are unaware of their lack of a referent. The other case is even more drastic: there are kinds of utterances which lack an addressee by their nature. Talking to oneself, for instance, can take the form of addressing oneself (“You shouldn’t have drunk the whole bottle”, said by a rueful speaker), but it may not, in what is often called “thinking out loud”. Or think of someone in a wistful mood, using Stevie Wonder’s song to express his feelings: “I wish those days could come back once more”. This person is not saying it *to* someone else, nor to himself; there is no addressee; he is just expressing himself, thinking out loud. Besides, any semantics of indexicals worth its salt ought to allow for the use of indexicals in thinking, where typically there is no addressing going on.

Once we shine a light on “you”, other hard cases emerge: what should we say about “we”, or the second person plural pronoun, “you”? I leave these questions open here. But let’s look at a clearer case, because it involves a postulated word. “Fthat” is a term introduced by Gillian Russell in Russell (2010) as an indexical thus: “[it] directly refers to whatever object the speaker’s favorite thing is in the world of the context at the time of the context”.³⁸ “Fthat” is an indexical by traditional lights, since nothing in its rule

³⁷More needs to be said about what it takes to be available for being addressed. There are exotic cases, just as there are with all other indexicals. One can address a letter to one’s unborn children, or to the first person to visit Mars, which suggests that acquaintance may not be a requirement. But what *does* it take to address someone? And, as always, phone conversations bring challenges of their own.

³⁸I should note that it is not all that important for Russell that it is an indexical, not a demonstrative; nor

mentions the speaker's intentions. But this is ad hoc. Why would we not put preferences on the same footing as intentions? They are similarly speaker-dependent, and they are both occasionally hard to figure out by the audience. My inclination, if I was working with the traditional distinction, would be to count "fthat" among demonstratives.

By my lights, it is not a typical indexical, because there is nothing about the metaphysics of typical utterances that forces speakers to have favorite things. In fact, I suspect most of us do not have a single favorite thing in the world at a time; and even if we did, nothing could prevent a tribe of chronically undecided Martians to make utterances, in spite of their lack of favorite things. But it is also not quite a demonstrative, since it works on a contextual feature that is not dependent on a particular utterance. If the speaker does indeed have a favorite thing, then the term just picks it out; the speaker's job is done by the time he has ordered his preferences, with no concern for uses of "fthat". This point is particularly easy to see, since in English such a word is sadly missing, but English speakers do sometimes have favorite things.

This last point inclines me to agree with Russell: "fthat" would, indeed, be an indexical. I admit that nothing in the metaphysics of utterances themselves tells us much about favoring objects. But when we think of utterances, we think of the utterances we make, just as we are, in a world, at a time, with favorite things, people to talk to, etc. Just as there may be beings outside of space or time, so there may be beings incapable of favoring things; perhaps Francis of Assisi really did love everything equally. But a typical, normal utterance is located in space and time, it typically has an addressee, and perhaps is made by someone with well ordered preferences. And "fthat" bears another resemblance with typical indexicals: the speaker either has a favorite object or not, but that fact is not fleeting, like demonstrative-linked intentions. We don't have favorite things in order to refer to them by the use of "fthat", but if we introduced it into the language, we could do the reverse: refer to our favorite thing by using it. If this is correct, indexicals draw their

is the example crucial to her arguments.

power not solely from the structure of conversations, like the changing roles of speaker and addressee; but rather from the wider structure of a context of utterance, with all its conversation-independent richness.

The lesson I want to draw from our discussion of “you” is that the traditional focus on intentions is too crude, because intentions play two distinct roles in the semantics of deictics. In this chapter, I made two distinctions: among intentions, between pre-semantic ones and semantic ones, and among features of contexts, between fragile and robust ones. Using this new terminology, the problem is this: it has generally been assumed that semantic intentions are only connected with fragile features of contexts. “You” proves that idea wrong: the intentions that set up an addressee are pre-semantic, but the addressee feature is a robust one. The hearer may often need to figure out the speaker’s intentions in order to know who is being addressed. But these intentions do not interact directly with the word meanings; it’s only their effects on the context that matter. The other is to set up fragile features of contexts; word meanings operate directly on these intentions.³⁹

In fact, the tradition has not been totally blind to these issues. One symptom is that several other criteria have been proposed for the distinction between indexicals and demonstratives, though it has not often been noticed that they are not all extensionally equivalent. That is my task in the next section: I will present several alternative criteria, show that they differ with respect to “you”, and argue that my proposal is better.

2.7 Other Criteria

2.7.1 Perry’s Narrow and Wide Contexts

In Perry (2001, §4.4), Perry divides deictics into four types, generated by two distinctions. I will deal with them in turn.

³⁹I leave it open that demonstrative meanings may also operate on other contextual features, depending on your view about the semantics of demonstratives.

First, he distinguishes two kinds of context, or, as I would prefer to put it, two kinds of contextual features:

What I call *narrow context* consists of facts about which things occupy the essential roles involved in the utterance, which I will take to be the agent, time, and position. These roles are filled with every utterance. The clearest case of [a deictic] that relies only on the narrow context is ‘I’, whose designation depends on the agent and nothing else. The *wide context* consists of those facts, plus anything else that might be relevant, according to the workings of a particular [deictic].⁴⁰

Perry’s examples of deictics dependent only on narrow contexts are “I”, “now”, “here”, and “tomorrow”. Among the rest he counts all classical examples of demonstratives, such as “that”, “there”, “he”, and two more special deictics, “yea” (as in “The dog is yea big”) and “dthat”. More about the last two in a moment. As far as the rest are concerned, Perry’s distinction separates exactly indexicals from demonstratives. Furthermore, “you” would need to count as a demonstrative, since it depends on a context feature different from the agent, the time, and the location.

I agree with Perry’s broader claim that indexicals are those words which depend on contextual features that normally come along with any utterance. He says “every utterance”, which is stronger than my talk of “typical utterances”, but I don’t think that is a big disagreement; after all, there are utterances that take quite a bit of time, maybe even discontinuous stretches of it. But he gives no reason to include only the agent, the time, and the location on the privileged list. At least the traditional distinction had a principled reason: intentions were the reason for banishment. But once we put worries about intentions aside, as Perry does, why the short list? Why not an addressee? Don’t utterances

⁴⁰Perry (2001, p. 59). Perry uses “indexical” the same way I use “deictic”, so I replaced his term with mine, to prevent confusion. I assume that “position” includes, or maybe simply means, position in possible world space, that is, a possible world. It cannot mean “position in space”, since that would be redundant.

typically have those? For these reasons, I take it that the first distinction will not do as a way to separate indexicals from demonstratives.

Perry's other distinction focuses on the role of intentions, and divides deictics into automatic and discretionary ones:

An utterance of "yesterday" designates the day before the utterance occurs, no matter what the speaker intends. Given the meaning and context, the designation is automatic. No further intention, than that of using words with their ordinary meaning, is relevant. The designation of an utterance of 'that man', however, is not automatic. The speaker's intention is relevant. There may be several men standing across the street when I say, 'That man stole my wallet'. Which of them I refer to depends on my intention. I have some choice or discretion in the matter.⁴¹

Examples of automatic deictics are "I", "tomorrow", "yea", and "dthat". Discretionary ones are the typical demonstratives like "this" and "there", plus most uses of "now" and "here", when used to pick out not just the very instant and spot of the utterance (those uses are automatic ones), but rather regions which contain them. For example, when we say "7 million people live here", we do not mean the very spot we're in, but some region around it. Since intentions determine the extent of the region, these uses of "now" and "here" count as discretionary.

This distinction comes closer to mine. I would count "yea" as a demonstrative, because the gesture that must accompany it is designed for that particular word; we could think of it as the externalized version of the kind of intention that provides a referent for demonstratives. Nor is Perry's reason to count it as automatic quite clear: on what grounds do we discount the speaker's discretion in making the relevant gesture?

⁴¹Perry (2001, p. 60). On p. 62, Perry makes it clear that the speaker's intentions are not sufficient to establish the referent; he settles on salience as the best candidate for that role.

Putting aside “yea”, there need be no conflict between Perry’s distinction between automatic and discretionary deictics and my distinction between indexicals and demonstratives. His analysis is much more brief, and therefore it could be filled out in various ways. In any case, he needs to distinguish between intentions that are relevant for semantics from those that are not.⁴² I offer my semantic intentions, and the distinction between robust and fragile context features, as a way to do just that.

2.7.2 Referent Change within a Context

The main question addressed in Braun (1996) is how to deal with the fact that two occurrences of the same demonstrative in a single context may not corefer. For instance, the following sentence can be truthfully uttered while pointing first to a ship, and then to another:

(18) That is bigger than that.

Braun’s theory deserves a longer discussion. But here, I am only interested in one advantage he claims to have over Kaplan’s theory of demonstratives: that he succeeds in drawing a sharp distinction between indexicals and demonstratives. And that is true: if his theory is correct, then the two are semantically quite different. All this presupposes that indexicals really don’t behave like demonstratives, i.e. that they don’t change their referent within a context. And this is not a hard claim to make; note that it’s hard to imagine how the following sentence could be truthfully uttered:

(19) I am bigger than I am.

This seems to show that indexicals keep the same referent throughout an utterance. So we have another possible criterion, though Braun himself doesn’t propose it as such:

⁴²A fact readily admitted by Perry: Perry (2001, p. 60), where he appeals to David Kaplan’s notion of a directing intention.

- **Referent Change:** two tokens of the same demonstrative can have different referents within a single context. All tokens of each indexical must have the same referent within a single context.

This criterion, or one much like it, also underlies the only sustained discussion I know of that uses “you” as a paradigmatic demonstrative: [Aft], pp. 585–590. The main reason for counting it as a demonstrative can be found on p. 587: “It is no part of the *meaning* of “today” that multiple syntactic occurrences must be associated with different contexts. In contrast, the meaning of a demonstrative requires that each syntactic occurrence be associated with a [its own] directing intention.” I think there is something right in this: there are differences between indexicals and demonstratives in the structure of the semantically relevant intentions. But I will argue that the difference does not lie in the possibility or necessity of having only one referent per context.

Kaplan’s system was well placed to deal with indexicals. But if Braun is right, it cannot deal with demonstratives, without some artificial decisions. All this stems from a problem posed, but not solved, in [Demons], §§ XV and XVI, and [Aft], § II, which contain a few options for dealing with sentences like (18). In a nutshell, the problem is this: if demonstratives are exactly like indexicals, given certain central features of Kaplan’s account, any two occurrences of the same demonstrative are semantically guaranteed to corefer in any context. Since it has always been assumed that there can be no more than one context per sentence, (18) is a clear counterexample, since the two occurrences of “that” do not corefer. So demonstratives are not exactly like indexicals.

So far, it looks like Referent Change is a good criterion, in that its application is quite clear, and it reveals an important and interesting semantic difference between indexicals and demonstratives. And if we apply it to “you”, it comes out as a demonstrative. To illustrate, we can use an example from earlier:

(15) You, you, you, and you can leave, but you stay.

The traditional criterion, which looked at the role of intentions for guidance, does not entail the Referent Change criterion. However, they can be connected by making some natural assumptions. One could easily think: since demonstratives require intentions, whereas indexicals do not, we can blame intentions for the fickleness of demonstratives, and say that indexicals are stable within a context just because their meanings rigidly dictate their behaviour.

Some assumptions are well hidden under this train of thought, but they are there nonetheless. If we allowed one context per deictic occurrence, it would trivially follow that no word changes its referent within a context, since it would be impossible for two words to occur in the same context. This assumption is well motivated, of course, but not unchallengeable.⁴³ In other words, we run again into an old problem: it all depends on what you want to put in a context, and what the nature of a context is.

In fact, I don't even think that Referent Change correctly characterizes indexicals. The following sentence can be true only if we assume that time passes between the utterances of the indexicals:

(20) Now you see it; now you don't.

There are true utterances of this sentence, and that fact can only be modeled if we assume that the two occurrences of the indexical "now" get different values. Admittedly, one could claim that it's really two different sentences that are being uttered, and each sentence gets its own context, so we still haven't proven that indexicals change value within the utterance of one sentence.⁴⁴

⁴³Crimmins (1995, p. 398–399, fn. 12) speculates that the idea of having several sub-contexts for just such a situation might work, but also expresses certain worries about it.

⁴⁴And anyway, even if there was an explicit conjunction between the two sentences, there are theories which treat "and" as a very special connective, not so much as a creator of new sentences, but more as an instruction about how to deal with two consecutive sentences. I have in mind the dynamic logic tradition, starting from Groenendijk & Stokhof (1991). Besides, we can easily construct similar sentences that are not molecular: "The object that you now see has now vanished".

So let me give another example, a disjunction. Suppose you are conducting an experiment, and you cannot see what the test subject sees. You do know the setup, however, so you know that there are only two possibilities: the subject is either seeing a meadow or a desert. You also have a button which switches between the two images. To make sure that the setup is working properly, you ask the subject if she agrees with the following statement:

(21) You're either seeing a dominantly green image now, [you flip the switch] or you're seeing a dominantly green image now.

I think it's clear that this sentence can only be modeled correctly if the values of the two occurrences of "now" are allowed to differ within the disjunction. But you might still think that we have a molecular sentence, made up of two atomic sentences, and one could still cling on to the one sentence - one context model. So consider this final situation: I'm a detective, and I'm searching a room for hiding places beneath the flooring. I have hard shoes, so I just try jumping up and down in various places in the room. I discover that it sounds more hollow in one corner than anywhere else. I say:

(22) [Jumping in one corner] It sounds more hollow here than [now I jump over towards the middle of the room] here.

I admit that this is a rather exotic example, and if it is, it would be very interesting to figure out why. But I hope it's natural enough that I can make this point: there's nothing in the nature of indexicals that stops them from changing their value within the scope of one sentence. Indeed, even "I" is not immune to this, since there are couples who finish each other's sentences, and then "I am taller than I am" can be uttered truly, as long as the switch between the speakers happens between the two occurrences of the indexical. So all deictics can have non-coreferential tokens within the same sentence. We are tempted to describe this situation in terms of mid-sentence context change. But then, by parity of reasoning, why not mid-sentence context change for multiple occurrences of "that"?

In fact, this is also an answer to the main argument for counting “you” as a demonstrative, which relies precisely on the Referent Change criterion. The core of the argument is presupposed in the aforementioned [Aft], § II. My exposition will be more detailed than the original text, but I hope that it captures its spirit correctly.

Technically, it is possible to treat demonstratives like indexicals. Indeed, this is one of the simplest ways to the problem of providing a logic for both kinds of terms. [Demons], p. 528, discusses, but does not adopt, the view he aptly calls the “indexical view of demonstratives”. Recall that indexicals get their referent from a particular feature of the context (“I” picks out the speaker, “today” the day, etc.). According to this view, we can just treat demonstratives as if they were indexicals: they pick out the demonstratum of the context.⁴⁵

So far, so simple. But there is an immediate problem: there can be more than one demonstrative per sentence. This is the point of the Referent Change criterion: there can be two non-coreferring occurrences of the same demonstrative in the same sentence, in the same context. Since this is the *indexical* theory of demonstratives, it needs to prevent demonstratives from changing referent. The quick fix made in [Demons] is to say that there really is a sequence of demonstratives in each sentence, and a sequence of corresponding demonstrata. Thus, each occurrence of a demonstrative is associated with an index, perhaps, but not necessarily, indicating its place in the left to right ordering of occurrences of that demonstrative in the given sentence. So the underlying real form of “That is taller than that” is something like “That₁ is taller than that₂”. And now we don’t have the same *indexed* demonstrative change referent mid-sentence; rather, we have two separate underlying words, governed by different semantic rules: “that₁” picks out the first demonstratum, “that₂” the second, and so on. Furthermore, something like this solution is inescapable, since we may want to say that our sentence, plus the sentence

⁴⁵The view does not lack supporters; it was perhaps first advanced in Scott (1970), later adopted by Braun (1994), attacked and renamed “the bare bones theory” for its simplicity by Salmon (2002, p. 511), and then defended again in Caplan (2003).

“($\forall x$)($\forall y$) x is taller than y , then y is taller than x ”, entail another sentence, “That₂ is shorter than that₁”. In other words, we must keep track of indexed demonstratives.

Even its proponents admit that this solution is artificial. As Salmon (2002) points out, this introduces into our formalism words which do not exist in English; there really aren't infinitely many words beginning with “that”, and having numbers as subscripts. Nor is it the same case as with proper names, like “David”, such that several individuals may bear that name. For proper names, few feel any problem with letting logic work on a properly disambiguated version of English. For demonstratives, as for indexicals, there is just one word, with a constant meaning, and our logic should deal with this fact, or die trying. Caplan (2003, §4) replies to this objection by saying that the notion of context we use in our formal account is just as formal as our account. We need contexts in order to provide us with values for deictics; if we need a sequence, then we need the context to provide us with a sequence, and that is exactly what contexts are for.

There are two ways to reply to Caplan's argument. One is to insist, along with Kaplan, that contexts are not our playthings. They just are *represented as* sequences which serve to give the logician values for (some, possibly all) deictics, but what we aim to represent are parts of the world, which have certain features, like the speaker-feature, and, as I'm arguing, the addressee-feature. So there is a real metaphysical question whether intended referents are like those or not. I endorse this argument, but if you are not convinced, recall an earlier point: indexicals can vary in reference within a sentence just like demonstratives. It's just that the custom is to idealize away the possibility of referent change for indexicals, and not for demonstratives.

In Ch. 1, I argued that context change within an argument should not be idealized away, and that it doesn't need to be. That solution does not apply to the current problem, since what we would need now is to allow contexts to change mid-sentence, rather than just mid-argument. I don't know how to do this, so I won't propose a solution here. One idea is to allow one context per word; this seems to me the most promising option, but

there are technical difficulties which demand further attention.⁴⁶

For demonstratives, the problem may be slightly easier, since we can make use of the different token-specific intentions that can be connected to an utterance of a single sentence to account for each of their referents. But for indexicals, in spite of the difficulties, the difficult path outlined above is the kind of path we will need to take. In the case of “you”, for instance, since it depends on a robust contextual feature, there is no other way to change the addressee mid-sentence than by changing the context mid-sentence.

2.7.3 Emptiness

Another way to make the distinction was proposed in Corazza (2004, p. 140), and it also counts “you” as a demonstrative, and it also fails to really distinguish indexicals from demonstratives:⁴⁷

- **Emptiness:** indexicals cannot be vacuous; demonstratives can.

Here’s the idea: if I say “I’m tired today”, I cannot fail to refer to myself and to the current day. But if I say “He’s funny”, and I am in fact hallucinating the clown I intend to talk about, I fail to refer to anyone. By this criterion, “you” is clearly a demonstrative, since I might as well have been addressing the clown, saying “You’re funny”, and I would have failed to refer, just as before.

Corazza didn’t actually intend this to be a definition of indexicality, of course. After all, some people think that “God” cannot fail to refer, and that doesn’t make the name an indexical. But is it a feature which distinguishes the two kinds of deictics? Well, there are uses of indexicals where the connection to the referent is not as straightforward as our examples suggested. Consider the use of “I” in wills; when read aloud by the lawyer

⁴⁶For instance, truth is normally defined as truth in a context; if we have several contexts per sentence, with respect to which of those should we evaluate the sentence? Surely not all, since the facts may change from one context to another; but then, which one?

⁴⁷See also Corazza (2011, p. 135), which is a good introduction to the subject.

after the appropriate death, the referent no longer exists.⁴⁸ Of course, this doesn't show that the indexical doesn't refer; just that it doesn't refer to anything existing at that time. But "Aristotle" is in the same boat; the owner of the name is long dead, but the name still refers to him. It just shows that indexicals are more complicated than a few examples may suggest.

This is a well known problem. There are uses of indexicals that do not naturally fit in Kaplan's theory, and a lot of work is being done on them, which shows that the dust hasn't quite settled on the issue. For instance, I can leave a note on my office door saying "I'm not here", whose truth depends on me not being at the place denoted by "here".⁴⁹ I want to rely on a simpler counterexample. Suppose that the world ends in 1000 years. Take that to mean what you want: that the Universe will contract; or that the gods decide to put an end to this pesky time business. So here's a sentence which fails to refer to any date: "In 5000 years, chess will still be played". Semantically, "in 5000 years" is very similar to "tomorrow". It just happens to be an expression composed of 3 words, rather than 1, but its semantic rules must be very similar to "tomorrow". And yet, it fails to refer to anything, because there neither is nor will be anything for "in 5000 years" to pick out. And besides, we could make a similar argument for "tomorrow", if we were willing to entertain the depressing possibility that the world might end before midnight.

Corazza (2004, p. 140, fn. 10) recognizes this problem, and proposes a distinction between *basic* indexicals (introduced by means of a list: "I", "now", "here", "today"), and *derivative* indexicals ("tomorrow", "next Wednesday", "my car" etc.), and says that only the former cannot be vacuous, but the latter can be. The idea is that one could use basic indexicals to define the meaning of derivative indexicals, e.g. "tomorrow" picks out the day after today, and "today" picks out the day of... now.⁵⁰ So the Emptiness criterion at

⁴⁸These uses were already noted in [Demons], p. 491, fn. 12.

⁴⁹See, for instance, Smith (1989), Predelli (1998), and Corazza et al. (2002).

⁵⁰Technically, we will need a rigidifier like "dthat" to change the descriptions into singular terms, as appropriate.

most serves to distinguish basic indexicals from all other deictics.

“You” would clearly count as a derivative indexical, like “tomorrow”, since its character could be given thus: “the person I am addressing”, which is similar structurally to the proposal for the character of “tomorrow” in the previous paragraph.⁵¹ And, just like “tomorrow”, “you” may fail to have a referent in some contexts.

Three questions arise at this point. First, is Corazza right about basic indexicals? And if he is, is it a semantically important distinction? And if it is, should it replace the distinction between indexicals and demonstratives?

I am not sure about the first question. Note that his claim is universal: basic indexicals simply cannot be vacuous. Is that true? I will assume that the claim is meant to cover only simple cases, of straightforward assertions, not cases of recorded messages. The latter cases are quite complicated, which is why I haven’t focused on them either. Just to give an example: in Seattle, there is a train connecting the Sea-Tac airport with downtown. The train has no driver on board. Along the way, there are a few stops; before each, a recording of an announcer is played, saying “The doors will open on my right” (or “left”, as the case may be). There clearly is a message creator, and a message announcer, but neither is likely to be on each train, so “my” cannot refer to them. The message conveyed is quite clear: the doors on the right, facing in the direction of travel of the train. I am not sure how exactly to interpret the utterance. And I am quite open to the suggestion that it wasn’t a very good message; it sounded strange to me when I first heard it.

But, as I said, it is best to leave these issues for later. So instead, consider this medieval-sounding question: could God use “today” before time was created? Well, maybe not, if “use” entails an occurrence in time. On the other hand, perhaps omnipotence overrides

⁵¹There is a difference: “the person I am addressing” is not rigid, whereas “you” is. There are various ways to rigidify descriptions, and I am willing to grant that one of them may fix this particular issue well enough. Still, using Fregean language, they do express different thoughts, so it would not be very plausible to claim that “you” *means* “the person I am (now) addressing”. In fact, Frege ([1918] 1997, p. 333) made a claim in the opposite direction: that we sometimes use “I” in the sense of “he who is speaking to you at this moment”.

such issues. But this is a metaphysical issue, not a semantic one. And this brings me to the second question: it is not a semantic fact that basic indexicals cannot fail to refer. Even if they always have a referent, that fact would depend on the nature of contexts, and that is a matter of metaphysics.

This may sound like a strange complaint, since I also claim that the distinction between indexicals and demonstratives is grounded in metaphysics. But there is a difference. I argue that the semantics of indexicals rides on certain kinds of contextual features, which are metaphysically distinguishable from others. And I think this is true of both basic and derivative indexicals.⁵² The rules “pick out the speaker” and “pick out the addressee” seem to me to be on a semantically identical footing, even if metaphysics may make it the case that the former always provides a referent, while the latter may fail. There is nothing in the meaning of “I” that is sufficient to guarantee that it always have a referent. The more general point is this: some parts of the metaphysics of utterances are reflected in the meanings of English words, and some are not. All utterances, for example, are made at a certain temperature; yet we have no single indexical in English which would pick out the current temperature.⁵³ But the structure is there to be exploited, if the need arises.⁵⁴

You may now have guessed my answer to the third question: one distinction should not replace the other. I don’t claim to have shown that basic indexicals lack any special properties with semantically important consequences. Perhaps, if metaphysics entails that basic indexicals always have a referent, this should be reflected in what we require of contexts, so that we would make sure that there are no empty occurrences of “I”, etc. This is one kind of reason behind Kaplan’s claim that we should only talk of contexts with

⁵²“My car” is different, since it is not a robust feature of contexts that the agent normally has a car. But this is to be expected; once we include indexical phrases, anything goes.

⁵³The description in “what’s the temperature?” comes close, but it is not rigid.

⁵⁴Thanks to Walter Edelberg for bringing up temperature. He may have meant it as a counterexample; I like to think he did not. Besides, this does raise the following independently interesting question: are there other robust features of contexts? What are they? Does the speaker’s having a head count? Having a heart? Having thoughts? Having parents? A principled criterion is needed here.

a speaker, a place, a time, such that the speaker exists at that time, in that place; that, in other words, there are no improper contexts.⁵⁵ And perhaps no such requirements should be made for derivative indexicals; for instance, we could allow contexts with a limited number of days, so that some occurrences of “tomorrow” lack a referent. I will not take a stand on these issues. My claim is simply that indexicals cannot be distinguished from demonstratives by pointing out that indexicals are never empty.

2.8 What about Demonstratives?

The logic of indexicals depends on the fact that they have a rich character. For instance, we get that “It’s raining if and only if it’s raining now” is a logical truth because the left hand side of the biconditional must be evaluated at the same moment as the right hand side, and that is because “now” always picks out the moment of utterance, the moment at which the biconditional claim is asserted to be true. If we didn’t notice this special feature of indexicals, the sentence would be equivalent to “It’s raining if and only if it’s raining on September 3, 2012, at 10:20am”, and that is clearly not a logical truth. The bigger point is that we need to distinguish logical truth, i.e. truth in all contexts, from necessary truth, i.e. truth in all circumstances of evaluation.

I’ve also argued that the semantics of indexicals allows us even more. Based on the relations between the characters of indexicals, we can look at relations between contexts, and have a logic that allows the following to be a valid argument:

(4) A says “It’s raining today”. The next day, A says “It rained yesterday”.

This argument is valid because of two relations. One between the indexicals: the first sentence contains “today”, and the second “yesterday”. Their characters guarantee that if they are prefixed to the same sentence, and the first one is said exactly one day

⁵⁵[Demons], p. 509.

earlier than the second one, if the first sentence is true, the second is true as well. In other words, the relation between the indexicals guarantees that the argument is valid if a particular relation holds between the contexts. And, crucially, that relation is dictated by the characters of the indexicals.

The logic of demonstratives cannot be exactly like the logic of indexicals. This is not a new claim, but the argument is new: it's not because different occurrences of demonstratives within the same sentence or argument may fail to corefer. I argued that that is true of indexicals just as much as demonstratives. Rather, the reason is that demonstratives depend on utterances, not just on sentences in a context.

I see two options for a logic of demonstratives. One is to claim that demonstratives draw something like a character *from* a context, perhaps because of the intention they are associated with, or the demonstration, or the salience properties, etc. We don't, then, have the same semantics for all occurrences of any demonstrative, but for each there is a method of assigning a character in every context. This is one way to interpret Kaplan's original solution from [Demons].

In order to move from LD closer to LI, contexts would then play two roles in the logic. First, they would give a character to demonstrative occurrences. Then, the relations between the characters of any deictic, plus the meanings of the other sentence parts, would dictate the relations between the contexts of each premise for the argument to be valid.

Some of the problems this solution faces depend on the choice of the character-giver. If we choose intentions, we have two structures to worry about: what the speaker intends, and what he succeeds in doing. Suppose that the speaker intends two occurrences of "this" to corefer, but they do not (because he mistakenly believed he was looking at one snake, when there were two in front of him). Should the characters of those demonstrative occurrences incorporate the speaker's unrealized wishes? And there's also the reverse problem: suppose the speaker does not want two occurrences to corefer, but they do; what should the characters reflect? Or if the speaker has no opinion on the matter?

In §2.3.2, I argued that my intention to refer to David Hume by using “I” will inevitably be defeated. Its relevance to logic is thus null. “I am David Hume” has its logical form irrespective of whether I believe that I am David Hume and irrespective of my referential intentions connected to “I”. Can we adapt that argument for demonstratives?

If we try the same argument structure, we would look at a sentence like “That is David Hume”. But our puzzle is about connections between occurrences of demonstratives, not between one occurrence of a demonstrative and a proper name. So we need something a bit more complicated, like the following: I look at what appears to me to be a dog partly hidden by a tree, and I say: “that [fixing my attention on the head] is that [fixing my attention on the tail-end]”.⁵⁶ Suppose there really are two dogs. Intuitively, I would say that the two occurrences of demonstratives do not corefer; they each refer to one of the dogs.⁵⁷ So the statement is not true, and my intending to refer to a single dog is just as inefficient as my earlier intention to refer to David Hume by using “I”. The structure of the speaker’s intentions might be relevant for cognitive significance.⁵⁸ But it is not relevant for logic, nor for what is said, since the speaker’s intentions to employ two coreferential demonstratives are not sufficient to guarantee coreference.

This argument might not suffice. Starr (2010, §2.3) argues that intentions only matter when they succeed; that is, for logical purposes, we might want to separate true uses of “that is identical to that” into those in which the speaker intended to refer to the same thing twice, and those in which he did not. But we gain little by making this move. By making success part of the condition, we lose the link to cognitive significance, since “from the inside”, the speaker cannot tell if what she is saying is true or not, and therefore she cannot know the logical form of what she is saying just by looking at her intentions. Why then worry about what the speaker intended about the relation between her uses of

⁵⁶The example is from Perry (2001, §4.5).

⁵⁷And this is how it comes out on all semantical theories of demonstratives that I know of.

⁵⁸This is what Perry uses the example to show.

demonstratives? Isn't it enough that we are looking at the intentions that (help) determine the referent?

I said that there are two options. The first one was to take demonstratives as close to indexicals as they can get. The other option I have in mind is to put them in a class of their own: words for which only the referent matters, i.e. words like proper names, except that the identity of the name doesn't leave any trace in the logical form. This is what I called earlier "super-Millianism", but we can now be a bit more specific. So I propose to call it "logical Millianism". "Millianism" because we disregard all other features of the word except its referent; "logical" because we're disregarding them for logical purposes, not for semantical purposes. On one version, it might be even more strict than classical Millianism: if all demonstratives are alike, then it couldn't matter whether one uses "this" or "that", or "he", etc. So even word identity would be irrelevant for logical purposes. Or maybe we could wish to note the association between "he" and one of the genders, so some demonstratives would be partially descriptive.

I am drawn to logical Millianism, because it keeps logic within the bounds of stable word meaning. Once we allow intentions to play freely with reference, it seems to me that we are outside of logic, and within the bounds of semantics proper. But I'm not sure that logic really needs to be that strict, and I can see why one might care about the structure of the intentions of the speaker. Which is why I remain undecided on the matter.

Chapter 3

Whence Logicality?

3.1 Introduction

The logic of indexicals I have proposed differs from first order quantified logic, the paradigm of logicality, not just in expressive power, but also in more fundamental and unexpected ways. Some of these it shares with Kaplan's LD, which also incorporates context sensitive expressions, but some are specific to it. In this paper, I present some of the differences between these three logics. Then, I argue that the differences do not amount to a reason to reject LI's status as a logic. The two parts of the paper are thus independent: you may agree that I captured the differences correctly, but think that LI, or even LD, is not, properly speaking, a logic. To some, this may seem a terminological issue, of little importance. I disagree: however we decide the matter, we stand to learn something about what we expect from logic and the formal approach to natural language semantics.

3.2 A Worry to Put Aside: Logical Constants

One of the simplest ways to reject the logicality of LD and LI is to claim that indexicals are not logical constants. In general, how principled one's choice must be depends on what

purposes that particular logic was developed for.

That a logic needs to separate logical constants from other expressions is obvious from the fact that logic is formal, i.e. it is about the form of sentences, and the form is given by the distribution of logical expressions. The Medieval distinction between categorematic and syncategorematic words was made with this goal in mind. In more modern times, here is Tarski making this point quite clearly, and then commenting on the fact that, at that time, no criterion had been found and widely accepted:

Underlying our whole construction [of a definition of logical consequence] is the division of all terms of the language discussed into logical and extra-logical. This division is certainly not quite arbitrary. If, for example, we were to include among the extra-logical signs the implication sign, or the universal quantifier, then our definition of the concept of consequence would lead to results which obviously contradict ordinary usage. On the other hand, no objective grounds are known to me which permit us to draw a sharp boundary between the two groups of terms. It seems to be possible to include among logical terms some which are usually regarded by logicians as extra-logical without running into consequences which stand in sharp contrast to ordinary usage. In the extreme case we could regard all terms of the language as logical. The concept of *formal* consequence would then coincide with that of *material* consequence. The sentence X would in this case follow from the class K of sentences if either X were true or at least one sentence of the class K were false.¹

This only shows that logic works by choosing a privileged class of expressions; but we still don't know how to do it. Tarski was not sure that we could make the choice once and for all, on the basis of a single criterion. And, indeed, several attitudes are possible: on one extreme, we could think that there is just one, true logic, which entails that there is a fixed set of logical constants, and we just need to figure out which they are. One kind of motivation here would be the hope that mathematics will turn out to be in some sense subsumed to logic, or to logic plus some conceptual truths. Obviously, one presupposition here is not that we can subsume mathematics to something or other, but

¹Tarski ([1936] 1983, pp. 418–419). See Gómez-Torrente (1996) for a good presentation of Tarski's views on this issue, and a defense thereof.

that there is a specific benefit to subsuming it under the true logic. On the other extreme, we could have a very relaxed attitude, according to which the logician starts with one of many equally good goals in mind, and chooses a set of logical constants accordingly; thus, there is no one criterion for logicity, and no deep general question about which are the real logical constants.²

Whatever the reasons, it is true that indexicals are quite unlike the logical constants of first order logic, and this might be sufficient to count them out. For instance, as far as I know, nobody has yet proposed a proof theory for a logic of context sensitive expressions. But logic is often thought to be essentially tied to proofs.³ Some authors even require a particular kind of proof theory for something to count as a logic:

[S]uppose that an alleged logical constant is [...] offered for which introduction or elimination rules were unavailable. All would then, quite naturally, reject the idea that the expression is a logical constant at all. Logicians would think we are dealing with nothing of that sort.⁴

It would be nice to have a proof system for LI. And I am an optimist; I think it can be done. But such a system does not exist, and it might be uninteresting, or, worse, impossible.

But I do not place great importance on the relevance of this issue to the question whether LI is a logic or not. Indexicals can be captured formally; that was a central goal of [Demons], and I think it succeeded. So we can formally define truth for indexical languages, and thus validity. If that is done correctly, it is to be expected that we would find special sentences of the formal language which turn out to be true in all contexts (that this is the right notion of truth needs to be argued, of course). These are, admittedly, purely semantic, non-proof theoretic, notions; but logic is, in part, about semantics: it is about preservation of truth. As long as the special truths come out of a correct semantics, we

²In this paragraph, I'm following MacFarlane (2009), which is a good survey for this issue.

³See Hacking (1979) for an influential defense of the centrality of natural deduction systems to logic.

⁴Bueno & Shalkowski (2012, p. 13).

should welcome them. The relation goes the other way too, of course, since the logic can be a test for the semantics: if it turns out, for instance, that everything follows from a sentence we deem to be a logical truth, we need to revise our semantics. We are at the three way crossroad between formal semantics, philosophy of language, and logic. Each approach will inform the other, and preconceived ideas about any of them can be overturned. For these reasons, it would be wrong to move from talk of logical truth to talk of semantic truth. There is no distinction worth making here; the logic arises out of the semantics, and it would be wrong to obscure this fact.⁵

Still, I do feel the pull of this worry for some expressions of LD. A much discussed example of an alleged unexpected logical truth is the sentence “I am here”, which is formalized in LD as *Located*(I, Here). It is crucial to this example that “x is located at y” is treated as a logical constant, since if we symbolized the sentence as *F*(I, Here), it would obviously not be a logical truth (think of the interpretation which assigns “is 5 yards away from” to *F*). But what makes “x is located at y” a logical relation, unlike other, seemingly similar ones, like “x likes living in y”, understood as a relation, as in “David Kaplan likes living in Los Angeles”? A better way to ask these questions, given the last paragraph: is there anything of interest semantically in this relation, or is it just like any other relation, like being older than someone?

I don’t find it wildly implausible to think of “x is located at y” as semantically special. If you need a two typed language, with one kind of term for places, and another kind for things that might be located in places, then “x is located at y” is a notion that we would naturally want to capture logically, that is, whose “logic” we would want to capture. After all, presumably the language is typed because we found an important difference between the two kinds of terms, and the *location at* relation is one of the basic aspects of the relation between places and their inhabitants. One might worry about where we should stop; but I am not sure that we need an answer to that question. This is an old debate, which began

⁵For similar, though not identical, considerations about the scope of logic, see Kaplan (2005).

at least as early as 1936, with Tarski's work, and continues to this day.⁶ My attitude is similar to the one expressed in Warmbröd (1999): let there be a central core of logicity, as bequeathed to us by Tarski, and a surrounding cloud of languages, which can be as rich as we need them for the purpose at hand.⁷

One way to bypass issues about what makes something worthy of logical attention is to switch to talk of analyticity. This is what Predelli (2011) does, and it is not unreasonable. It is generally reckoned that all logical truths are analytic truths, since their truth depends only on the meanings of their parts, and their mode of combination. In fact, their truth depends only on the meanings of their *logical* parts, which makes logical truths be a proper subclass of analytic truths. An example ought to make this clear: "All bachelors are unmarried" is a prototypical example of an analytic truth, but few people would want to count "unmarried" a logical constant. This suggests the easy way out we wanted: since we're not sure if "I", or "today", or "located", should count as logical constants, just leave that problem aside, and talk of analytic truths instead. And there is something right about this. When, in proposing an epistemic logic, it is said that "it is known that p" entails that p is true, that is an observation about the meaning of "know".

Still, this strategy would not solve all our problems. If we found something to be a synthetic truth, that would prove it was not a logical truth. But we would still want a way to figure out which truths are merely analytic and non-logical. Still, for many of the critical discussions of LD, talk of analyticity has been sufficient. It is natural to ask, with Predelli (2011), if "I am here" is an analytic truth. In fact, as we will see, his point is even simpler: there are uses of that sentence such that it turns out false; so it cannot be an analytic truth, so it is not a logical truth.

⁶See Tarski ([1936] 1983), and, more explicitly, Tarski (1966). For a modern exposition, see Sher (1991).

⁷Warmbröd's main reason is that science seems to be doing just fine with first order logic ("FOL" henceforth), and the rest are needed only for other purposes. I cannot comment on his claim as far as physics and chemistry go. But formal semantics, which I regard as a science, needs a much richer language if it is to have any hope of characterizing natural languages. I agree with the idea that inasmuch as we want to ring off a part of logic as the central one, FOL is a natural candidate. But I don't agree with his larger claim that FOL is sufficient for scientific purposes.

Useful though it may be, replacing talk of logicality with talk of analyticity brings its own problems. Most obviously, since Quine (1953), analyticity has been seen as a dubious notion. More recently, there have been attempts to fight back: Boghossian (1996), Russell (2008), Juhl & Loomis (2010). All of these authors concede some points to Quine, and consequently defend only some parts of the classical notion of analyticity. I think it fair to say that none of their proposals has quite prevailed yet, and trying to choose among them is only tangential to my project.⁸ Therefore, I will only talk of analyticity in the broadest terms, and only inasmuch as it bears on the question whether LI ought to count as a logic.

3.3 A Challenge to LD

Before we look directly at LI, I would like to discuss first a challenge to LD. I have two reasons. On the one hand, it could easily be adapted to work against LI; so if I answer it now, I will have less to worry about later. Second, a close study of it will help me draw a lesson that I will use later: logical truths are not true only in virtue of the meanings of the logical constants; they also depend on what truth is, and, for both LD and LI, what contexts are.

One of the most surprising results of [Demons] is that the following are logical, analytic truths, even though they are not necessarily true (read @ as “actually”):

(23) [For any sentence ϕ] $\phi \leftrightarrow @\phi$

(24) I exist.

(25) I am here.⁹

⁸Besides, new arguments are still being proposed against any notion of analyticity; see, for instance, Williamson (2008, Chs. 3 & 4).

⁹In fact, the classical example is “I am here now”. I am removing the last indexical for the sake of simplicity; none of the points under discussion here depends on it. I should also note that, unlike the others, (23) is a schema, not a sentence. My imprecision should not generate any misunderstanding.

Kaplan was interested in the semantics and logic of indexicals, so he was most interested in the logical properties of these sentences. More recently, Gillian Russell has made sentences like these central to her account of analyticity.¹⁰ For indexicals, her theory comes down to the claim that “truth in virtue of meaning” should be clarified as “truth in virtue of character”. Since the character of indexicals is formally captured in their semantics, any indexical-involving analytic truths are likely to turn into logical truths. This type of analyticity promises to be interesting. For instance, its connection to other kinds of analyticity is not very simple. After all, the kinds of fact that allegedly make “I am here” special seem quite different from the kinds of fact that make “Bachelors are unmarried” special.¹¹ Even putting aside analyticity, the claim that some logical truths are contingent is a startling one. But before looking at these fascinating issues, the challenge needs to be stated and faced.

Stefano Predelli (2005) gives two arguments – a methodological one and an empirical one – to show that only (23) ought to come out a logical, analytic truth, and that the others are neither semantically nor logically interesting. I will show that his methodological argument does not work. That leaves his other, empirical argument. The latter is on the right track: empirical adequacy is where the action ought to be.

3.3.1 Kaplan Semantics, Quickly

A quick recap of Kaplan’s semantics is needed before we move to our topic. I will leave out many of the details, and even some of the central parts of the account, since they do not affect the current discussion.

Note first that indexicals do not have a value *simpliciter*. Typically, we think of the

¹⁰See Russell (2008, 2010). Kaplan does call a sentence of this kind “analytic” ([Demons], p. 529), and later says that “all logical truths are analytic” ([Demons], p. 540), so Russell is not departing from Kaplan on this issue.

¹¹This is a very delicate matter. For some very interesting, but brief, remarks on the issue, see Higginbotham (1988, p. 33) and Kaplan (2005).

proper name “Barack Obama” as naming Obama, and that’s that.¹² But “I am right handed” has no chance to get a truth value unless a context gives a referent to “I”. The first role of contexts, then, is to provide a semantic value to indexicals. Given the sentences we are working with, we need a context to supply at least an agent (for “I”), a place (for “here”), and a possible world (for “actually”). Once we introduce indexicals into a language, contexts are indispensable.

The second role of contexts is more complicated. The basic semantic value of a sentence ϕ is its truth value *in a context relative to a circumstance of evaluation*. For our purposes, it suffices to think of a circumstance as a possible world.¹³ Suppose that Barack Obama says “I am the president of the US”. What he says is true in the actual world because in that context “I” refers to Obama, and he actually is the president of the US. But we can take what he said and evaluate it at another possible world, where McCain won the 2008 elections. At that world, what Obama said is false. So the context provided the content of “I”, we then computed the content of the whole sentence at that context, and that content can be evaluated at any possible world.

Having defined truth at a context relative to a circumstance of evaluation, we now need a notion of truth *simpliciter*. Kaplan’s solution is to have the context dictate the circumstance of evaluation, so that a sentence ϕ is true at a context c just in case it is true at the context c relative to the world of c . More formally: ϕ is true at a context c iff $\llbracket \phi \rrbracket_{c, c_w} = T$, where c_w is the world of c .

¹²There are dissenters, of course. Pelczar & Rainsbury (1998) argue that proper names are indexicals. A bit more common is the view that they are dependent on context in some other way; see, for instance, Burge (1973), Recanati (1993) and Perry (2001, Ch. 6).

¹³Circumstances of evaluation may contain other elements than a possible world. Lewis (1981) included standards of precision. In the original [Demons], §XVIII, they also contained a time, to deal with the intuition that even if what Obama said is true at the time I’m writing this paper, it was not true, say, at any time in 1990. But I will leave aside time and tenses for the sake of simplicity.

3.3.2 The Methodological Argument

Predelli (2005, ch. 2) gives two arguments purporting to show that (23) is an analytic truth, but (24) and (25) are not, and therefore they are not logical truths either. Briefly, one argument, which I will call “the empirical argument”, points out that (24) and (25) can be used so that they turn out to be false, whereas (23) cannot. But a sentence that is false in some contexts is neither an analytic, nor a logical truth. I believe that this argument is where the action is, and I’ll come back to it in §3.3.4.

The other argument, which I will call “the methodological argument”, aims to show that even if there could not be any false utterances of any of these sentences, only (23) ought to come out an analytic truth, because it is the only one which is true in virtue of the meanings of its parts.

Let’s begin with (23). As we saw above, ϕ is true at a context just in case it is true at that context, evaluated at the world of the context. The clause for @ is this: $[[@ \phi]]_{c,w} = T$ iff $[[\phi]]_{c,c_w} = T$. So @ forces ϕ to be evaluated at the circumstance of the context of utterance, even when we want to evaluate $@\phi$ at other possible worlds. The result for truth *simpliciter* is that ϕ is guaranteed to have the same truth value as $@\phi$, as long as they are evaluated at the original context. In other words, we get that $\phi \leftrightarrow @\phi$ is a logical, analytic truth.

We get this result due to the semantics of the “actually” operator, and the definition of truth in a context. Since the world feature of contexts plays a role in both definitions, we get a guarantee of the truth of (23). So far, Kaplan and Predelli agree.

But so far, (24) and (25) are not analytic truths. After all, it might be that the context assigns me as the agent, but assigns Vancouver as the place, and I’ve never been to Vancouver, so “I am here” comes out false. [Demons], p. 509 argues that proper *contexts of utterance* are contexts in which the agent of the context exists at the world of the context, and is located at the place of the context. His claim was not intended as a postulation, or as a claim about the needs of semantics. It was a claim in metaphysics: that is how

contexts *are*, however else we may wish them to be; they are contexts for real utterances, parts of the world, not mathematical objects. The properness condition guarantees that “I am here” and “I exist” are true at all contexts: by narrowing the domain of contexts to what Kaplan calls the *proper* ones, we get certain special truths out of the structure of the contexts plus the semantic rules for the indexicals.

In fact, Kaplan’s terminology is a bit misleading. He doesn’t just think that there are two kinds of context, and we shouldn’t worry about the improper ones. Rather, his thesis is that all contexts are proper.¹⁴ Here’s Kaplan’s reason: improper contexts “are like impossible worlds; no such contexts could exist and thus there is no interest in evaluating the extensions of expressions with respect to them”.¹⁵ Talk of improper contexts, in other words, would serve no purpose. Indexicals don’t grow on trees; people use indexicals, and this fact imposes certain conditions on the environments we want our semantics to take into account. Furthermore, he claims, this is a good result: “Intuitively, [(25)] is deeply, and in some sense, which we will shortly make precise, universally, true. One need only understand the meaning of [(25)] to know that it cannot be uttered falsely”.¹⁶

Predelli disagrees:

“Is it the case that ‘one need only understand the meaning’ of the relevant expressions, in order to understand this example’s presumed privileged status? I do not think so. [...] In this case, an important premise must be accepted, in addition to the hypotheses regarding the meanings of ‘I’, ‘am’, and ‘here’[...]: the indexes which the system may take into consideration must be proper indexes. It seems clear that such an assumption is not warranted by

¹⁴Lewis (1981) thinks of contexts as temporal and spacial locations in a possible world. So he allows more proper contexts, namely ones in which there is no utterer, but resembles Kaplan in thinking that contexts are not just any ordered tuple of features.

¹⁵[Demons], p. 509.

¹⁶[Demons], p. 509. I changed the numbering to match my examples. Kaplan was talking about “I am here now”, but, as I said, I am putting aside temporal issues.

the conventional meaning of any expression.”¹⁷

You may have noticed that Predelli talks about “indexes”. The reason is that one of the central points of his book is that it is important to think about what goes on when we move from a real world context of utterance to its formal representation, which he calls an “index”. The point of the methodological argument is that even if all contexts of utterance were such that the speaker of the utterance is located at the place of the utterance, this fact should not be a constraint on the indexes. And the main reason is that it would make “I am here” analytically true, when it is clearly not true solely on the basis of the meanings of the parts, but also because of some non-semantic fact about contexts of utterance (and the role of the empirical argument is to show that this fact does not obtain).

Predelli’s point is very clear with “Something exists”. First, it clearly comes out a logical truth in LD, since it follows by existential generalization from “I exist”, and, since the latter is a logical truth, so is the former. And, indeed, that is the result we get from the model theory. Recall that every context must be such that the speaker exists at that world, in that place. So any time we evaluate “Something exists” at a proper context, it will be guaranteed to be true by the speaker’s guaranteed existence. But, Predelli rightly argues, “Something exists” cannot be true in virtue of the meanings of indexicals in it: there are none.¹⁸ But how could thinking about indexicals lead us to discover new analytic truths about indexical-free sentences?

It may be convenient to define indexes as some ordered tuples, so that semantics be able to do its job cleanly. That never happened in [Demons], where in fact contexts are not so defined, and we are not even given identity conditions for them, and few philosophers seemed to mind it. But we could do it, if we wanted to. And yet, this is not sufficient to make Predelli’s argument work. The meanings of indexicals that characters are supposed to capture are not related to any tuples. The attractiveness of saying that “I” refers to the

¹⁷Predelli (2005, pp. 62, 64). As before, I left aside “now”.

¹⁸Recall that we put aside tenses. If we had not, it would be very hard to find an indexical-free sentence of English with a verb in it.

speaker of that context of utterance comes in large part from the fact that we are talking about *contexts of utterance*. So part of the spelling out of the meaning of “I” makes use of the notion of Kaplan-style contexts, not the notion of Predelli-style indexes. This shows that, contrary to Predelli’s claims, if we figure out an important property of contexts of utterance, that ought to affect our account of the meaning of indexicals. If Kaplan is right about our concept of a context of utterance being limited to the contexts he deems proper, then those are the contexts that we use in the semantics of indexicals, and it is a discovery about the meaning of “I” that “I exist” is a special truth.

Nor does Kaplan have to be exactly right about contexts of utterance for Predelli to be proven wrong. Suppose, in the style of Reichenbach and Perry, that we think that a context of utterance for ϕ must be a context in which ϕ really is uttered. This is not part of Kaplan’s notion of a proper context, nor of Lewis’s, and it adds new analytic truths. For instance, “I am speaking” will always be true, since the speaker in any context of utterance of that sentence will be speaking in that context. Even on this conception, “Something exists” remains true in all contexts, as does “I am here”. And these truths are declared special by thinking about the notion of a context, and by allowing these reflections to seep into our semantics.

3.3.3 Logical Truths and Consequence

Suppose that Predelli were right in saying that the meanings of indexicals *by themselves* do not make (24) and (25) logical, analytic truths. He would still be wrong to conclude that they are not logical truths, and he would still be wrong about analyticity. To show this, I offer two arguments.

The first argument is a kind of *tu quoque*: if Predelli is right about these sentences, he is wrong about (23). $\phi \leftrightarrow @\phi$ comes out an analytic truth not just because “actually” means what it does. Part of the work is done by having a single context throughout the sentence, and having that context play the double role of supplying both a semantic value

for @, and a world as a circumstance of evaluation for the whole sentence. The latter fact does not follow from the meaning of any indexical. It is a claim about the relation between truth in a context relative to a circumstance, and truth in a context *simpliciter*. That the two are independent can be seen from the fact that much of the recent literature on relative truth concerns the question whether truth in a context depends just on features of the original context of utterance, as opposed to, say, features at a context of assessment.¹⁹ But Kaplan's semantics of indexicals is not challenged in that literature.

Predelli might answer by pointing out that, true or not, claims about truth in a context are *bona fide* semantic claims, even if they are not about indexicals per se. The idea would then be to draw a line between these acceptable sources of analyticity, and ones which concern the nature and structure of contexts. One way to argue for this is to claim that metaphysics is the proper venue for the latter issues, and that logic should stay as metaphysically neutral as possible. This gets us to the heart of the matter: I want to argue in favour of mixing metaphysics and semantics, and, therefore, logic. Indexicals really do get their values from the surroundings of utterances. In a formal system, we want to represent certain aspects of contexts formally, of course, and that forces us to a certain degree of idealization. But that fact should not lead us into the other extreme, of forcing those representations of contexts so far away from what they are meant to represent that they become unrecognizable. If Predelli's indexes give us some object as the semantic value for "I" merely because, say, the first place in some n-tuple is appointed to that role, in what sense are we analyzing the first person pronoun? The rule we wanted to capture says that uses of "I" refer to the speaker; if we replace that with a simple function, we lose much of what we wanted to achieve. So I advocate an infringement by metaphysics into logic: we need to think about the nature of the contexts of real utterances, and see which

¹⁹See, for instance, MacFarlane (2008), and other papers in that collection. MacFarlane presents a semantics to go with a forking notion of time, so that many present claims about the future are false because the future is unsettled, but when the time that the claim was made about comes, the truth value does get settled, so that the claim, from that point of assessment, is simply true or false, as the case may be. So truth values are relative not just to the context of use, but also to the context of assessment.

of their parts we want to capture in our semantics, and which are extraneous.

This is a programmatic argument, and programs are only convincing if they have good results. I take it that I have done some work in that direction in Ch. 1, where I showed that we should not shy away from context changes as they occur in real conversations, and in Ch. 2, where, based on thinking about what kinds of features of context play a semantic role for different types of context sensitive expressions, I proposed a new way to distinguish between indexicals and demonstratives.

But I can provide an illustration even without going back that far, and so I offer my second argument against Predelli's conclusion. The following is an argument just like (23), except that we have an argument with a premise and a conclusion where before we had a biconditional sentence:

(26) [For any sentence ϕ] $\phi \therefore @\phi$

Now recall the classical Deduction Theorem:

Classical Deduction Theorem: $(\forall n \geq 0)$ an argument $\phi_1, \phi_2, \dots, \phi_n, \psi \therefore \chi$ is valid iff $\phi_1, \phi_2, \dots, \phi_n \therefore (\psi \rightarrow \chi)$ is valid.

Take the case where we only have one premise on the left hand side, and we get the result that $\phi \therefore @\phi$ is a valid argument iff $(\phi \rightarrow @\phi)$ is a logical truth.

The Deduction Theorem is central to classical logic, but it does not apply to LI. As we noted in §1.4.4, there are two reasons for this: first, in LI an argument is made up of context-sentence pairs, not just sentences, and, second, the form that does hold in LI needs the constraint that the contexts of the last steps in the argument are the same:

Indexical Deduction Theorem: $(\forall n \geq 0)(\forall i)$ an argument

$\langle [c_1, \phi_1], [c_2, \phi_2], \dots, [c_n, \phi_n], [c_{n+1}, \psi], [c_{n+1}, \chi] \rangle$ is valid iff $\langle [c_1, \phi_1], [c_2, \phi_2], \dots, [c_n, \phi_n], [c_{n+1}, (\psi \rightarrow \chi)] \rangle$ is valid.

So my argument will have a rather complicated structure: first, I will draw some conclusions that Predelli would presumably accept. Then, I will show that in LI we can draw some related, but different conclusions. The reason for the differences lies in a disagreement about what ought to constitute an argument. Since one consequence of the differences is that we get different lists of logical truths and analytic truths, the conclusion will be that we decide that something is a logical truth not just on the basis of word-level semantics, but also on the basis on word-independent theses about what an argument is. So Predelli is wrong to think that logical and analytic truths are those that are true only due to the meanings of the component expressions.

Start with the classical Deduction Theorem. My argument can be run as being about logic directly, or about analyticity. The former is simpler, so I will begin there. First, note that (26) is valid in LD because of the assumption that in every argument, the premises and the conclusion are in the same context. I ask: why did Kaplan and Predelli focus on just those context sequences? One kind of reason is the one offered by Soames (2010): logic cannot handle context changes. I've tried to show that to be incorrect. The other kind of reason would be to claim that this is what the proper purview of logic is; that thinking about argument structures tells us that logic should only see arguments with the same context throughout. I disagree: we argue with each other, and across time. Logic cannot take that away from us. And the only way the Deduction Theorem remains valid is if we disregard these facts, just as in LD, by requiring that every argument take place within a single context.

But we don't need to revisit that dispute right now. All I need is to note that the disagreement between Predelli and me, and between LI and LD, is about what counts as an argument. And this can easily be seen in the different forms that the Deduction Theorem takes in the two logics. The disagreement, crucially, is not about the meanings of indexicals, nor is it about what truth in a context amounts to. Just as philosophers can disagree about the components of circumstances of evaluation, they can disagree about

what kinds of context sequences are appropriate for our attention. So the decision about whether to count $\phi \therefore @\phi$ as an argument, which is a pre-condition to counting it as a *valid* argument, depends not only on the meanings of indexicals.

Logic is, in part, about what counts as an argument; the fact that some of these decisions are virtually universally accepted is not a reason to forget that they are made, they can be challenged, and they matter.

My preliminary conclusion, then, is that part of what makes us count a truth as analytic is a decision about matters that are not immediately about the meanings of any expressions, but about what counts as a truth, and what counts as an argument.

I must admit that Predelli talks of analyticity, not about logical truth, but he clearly thinks logical truths are all analytic.²⁰ Fortunately, we can make a similar argument about analyticity. First, we generalize from the notion of analytic truth to analytic consequence.²¹ Let us say that ψ *follows analytically* from a (possibly empty) set of premises Γ iff meanings alone guarantee that if all the members of Γ are true, ψ is true. We call such arguments *analytically valid*. I assume that an Analytic Deduction Theorem is also desirable for Predelli:

Analytic Deduction Theorem: $(\forall n \geq 0)$ an argument $\phi_1, \phi_2, \dots, \phi_n, \psi \therefore \chi$ is analytically valid iff $\phi_1, \phi_2, \dots, \phi_n \therefore (\psi \rightarrow \chi)$ is analytically valid.

Just as above, one consequence is that $\phi \therefore @\phi$ is an analytically valid argument iff $(\phi \rightarrow @\phi)$ is analytically true. And the argument given above applies here just as well: accepting the Analytic Deduction Theorem depends just as much as the regular Deduction Theorem on decisions about the nature of arguments.

So how do we decide? I argued in Ch. 1 that some of the arguments we intuitively judge to be valid cannot be made valid unless we change something quite dramatic about the scope of logic; unless, that is, we move to something like LI. But there is an even easier

²⁰See, for instance, Predelli (2011, p. 291), where he equates Kaplan's notion of validity with analyticity.

²¹The idea of analytic validity is also explored in Russell (2008, pp. 101–102).

way to illustrate my point: LI requires arguments to take place in the same possible world, i.e. that all contexts in an argument have an identical possible world. The justification for this requirement went back to the overall reason for being interested in a logic that allows context changes: we do give arguments during which the speaker or the day changes, we do take into account the context changes, and the semantics of indexicals allows this. But we never argue with people in other possible worlds; how could we? And we do not argue with ourselves in another possible world either. We are all stuck here. This fact is not as obvious as the fact that days change, that time passes, or that conversation participants take turns in a conversation. Not obvious, that is, if put in terms of possible worlds; but it is a fact that our arguments depend on nonetheless: why else would we assume that the truth of one premise is relevant to the truth of another?

The most important lesson I want to draw from this discussion is that the move to truth *in a context* brings along with it a new conception of what an argument is. This fact is easy to miss in LD, since it is stipulated that all contexts in an argument must be the same, and, from the point of view of validity, no two contexts are interestingly different: they all settle on a speaker, a time, a place, a world, etc. So we are left with the illusion that arguments are made up solely of sentences, just as before. We should have seen the signs: the richer we make Kaplan's properness condition for indexicals, the more special truths we are likely to get. But the newness of it all comes out even more clearly in our discussion of Predelli's methodological argument: by thinking about the nature of contexts, and what to represent formally, we influence the amount of truths we deem special, be it analytic or logical. Thinking about indexicals gets us this new notion; but thinking about our tools for handling indexicals is just as important.

3.3.4 The Empirical Argument

So far, I've argued that some features of real conversations ought to be allowed to make it into their formal representation. Two questions arise. One is: which features of conversa-

tions really are universal, or, at least, typical enough to deserve our attention? Requiring universality may well be too much; for instance, few semantic theories of proper names try to deal with uses like “You’re no Einstein”, taken as an insult, not as a remark that the addressee is not *the* Einstein, or named “Einstein”. So we probably don’t need to represent conversations with the goal of dealing with any possible use of indexicals. But what counts as typical enough, or central, as opposed to derivative?²²

The second question is this: do we want any and all features shared by central cases of conversations captured in our formal semantics? Clearly not. Here is a trivial example²³: pick a sentence that nobody has ever, will ever, or could ever utter (perhaps because it is too long). Then there is, in fact, no possible context of utterance in which that sentence is false. So any argument that has that sentence as the only premise comes out valid, since there is no possible context such that the premise is true and the conclusion is false. So any contradiction follows from this premise. That is, of course, not desirable.

Other examples are not quite this trivial. This is where Predelli’s second, empirical argument comes along.

Begin with (25). Suppose I am holding office hours, but prefer to close my office door unless a student comes by. In order to prevent students from being discouraged, I place a post-it note on my door, saying “I am here”. Admittedly, I would be more likely to write something more explicit, like “I am in my office; please knock”, and perhaps sign it. But the less explicit note would probably do the job. Three quarters of an hour later, nobody has come. I am a bit tired, so I decide to go get a cup of coffee, but I forget about my note. If, while I am away, a student comes by and reads my note, and discovers that I’m not there, he can accuse me of having conveyed a false message. I was not where I said I was at that moment, my message was false. This shows that there are false uses of “I am

²²With application to proper names, the issue is discussed in Jeshion (2012). Of course, it is not trivial to apply her proposals to indexicals.

²³See [Demons], p. 522 for a quick version of this argument.

here”.²⁴ Assume I died when I left my office, and if I had added “I exist” on my note, that would be false too, if existence entails being alive.²⁵

What do we do with these data? First, the examples seem to support the conclusion of the methodological argument: since there are false uses of “I am here”, it cannot be an analytic or a logical truth. The semantics proposed by Predelli is quite simple. “I” would get its semantic value from the agent feature of the index. The post-it data would only be relevant for the question about how the agent feature gets its value from real world contexts of utterance. This move would make it easy for semantics to deal with post-it notes, if semantics was merely about an indexical getting a semantic value from a tuple of objects, ready to go. But this simplicity is misleading: semantics ought to be concerned with how that tuple got to have those features, and not others. In other words, a semanticist would be just as, if not more, interested in the way in which an object related to the real world utterance gets to play a certain role in the semantic representation of the utterance, the n-tuple. Note that we have continued to speak of “utterances” of these sentences when what is at issue is *readings* of the sentences long after any utterances were committed to paper.

Since the methodological argument is not correct, we don’t even get the Pyrrhic victory of pushing everything down to the level of issues about representing utterances, and we’re back to wondering how semantics ought to deal with my post-it note. And the matter is not settled by the empirical argument alone (as Predelli acknowledges²⁶). What is needed is a theory of recorded messages, and work has been done on this issue.²⁷ When

²⁴Examples like this have generated a fair amount of discussion; see Predelli (2011), and references therein. Replacing post-it notes with answering machines, it all began with [Demons], p. 491, fn. 12.

²⁵Predelli (2005) discusses other recalcitrant uses too, such as “here, the mountains appear on the horizon”, as appearing in a guidebook, where “here” is plausibly taken to pick out the area being talked about, not the writer’s location at the moment he wrote the sentence. A non-recorded message example are demonstrative-like uses of ‘here’: “Here, the Gauls made their last stand against the Romans”, said while pointing on a map at one of the presumed locations of the Battle of Alesia.

²⁶Predelli (2005, p. 64).

²⁷For a good and recent review of the literature, which is also relevant to our discussion, see Predelli (2011).

we settle on a preferred theory, we need to proceed to a larger question question: should recorded uses of indexicals be treated under the same theory as the more common, live uses? Can we accommodate recorded messages in an extension of LI? Or, maybe, can we accommodate LI within a theory of recorded messages, perhaps as a special case where recording and broadcast coincide? What should the scope of our theory be, and why? Considerations of analyticity will play a role in answering this question: if Kaplan is right about the specialness of “I am here”, and yet there are false recorded messages made up of that sentence, does that show that the specialness is a mirage, or does it show the need for separate theories?

And what kind of specialness do we look for? It cannot be that all English speakers believe that “I am here” is special. As Williamson (2008, p. 89) pointed out, an English speaker may have a mistaken view about the meaning of “vixen” which makes him think that “all vixens are vixens” is false, and yet this would not stop him from being a speaker of English, even if that sentence has a good chance of being an analytic truth. All the more for “I am here”, of course. So work on the nature of truth in a context, and on validity in a sequence of contexts (what I earlier called a “conversational situation”), needs to be done to figure out what features of contexts are important, and which can be abstracted away.

3.4 Special Features of LI

Conceptually, the distance between LI and LD is smaller than the distance between LD and FOL. Truth in a context is not new to LI; logical truths generated from theses about contexts, as we just saw, are a semi-hidden feature of LD. But the discussion about Predelli showed that some of the revolutionary aspects of LD were masked by its insistence on keeping contexts fixed throughout an argument. In this section, I explore several ways in which LI differs both from FOL and from LD, with the final goal of seeing exactly why one should not doubt the logicity of LI. Full disclosure: it will turn out that there are

several kinds of logical truths, only some of which will serve all the purposes we expect logic to be helpful with.

The innovation of LI with respect to argument structures is to include the structure of conversational situations as part of the form of an argument. This gives us a new source of validity. FOL gives us some logical truths; LD gives some more, based in part on theses about truth in a context, and in part on thinking about the nature of contexts. By looking at conversational situations, LI opens up the possibility of logical truths which are partly based on forms of argument that are not purely sentential.

3.4.1 Logical Truths and Valid Arguments in LI

The most dramatic example of the difference between LI and LD is the fact that the Deduction Theorem fails in LI, which forces us to keep logical truths and validity apart. So unlike FOL and LD, apart from special logical truths, we will also get forms of valid argument where we didn't expect them to appear.

For instance, a sequence of sentences (what I earlier called a "conversational thread") may be part of a valid argument in one conversational situation, and an invalid one in another.²⁸ Furthermore, this is unavoidable, if we are to allow contexts to vary within an argument. As Soames pointed out, if we don't look at types of arguments, pretty much no argument will come out valid, and we have a logic no more.

Before we discuss the new and troubling ways to be valid, I should mention that there is a special class of conversational threads such that no matter what (proper) conversational situation they are set in, a valid argument will be generated.²⁹ I will call these "super-valid arguments", and, in the case of single-sentence threads, "super-logical truths". These are closest to the logical truths of LD, since their validity is guaranteed by

²⁸Here's the proof. Let ϕ be "I am right-handed", and ψ be "You are right-handed". Let c_1 be (S, A, D, W) , and c_2 be (A, S, D, W) , such that $A \neq S$; that is, c_2 is exactly like c_1 , except that the speaker and the addressee are reversed. Then, $\langle [c_1, \phi], [c_1, \psi] \rangle$ is an invalid argument, whereas $\langle [c_1, \phi], [c_2, \psi] \rangle$ is valid.

²⁹Thanks to David Christensen for pressing me on this issue.

the sentences themselves, as long as they are placed in a conversational situation.

One simple example is logical truths of FOL; for instance, $(\xi \rightarrow \xi)$ is a super-logical truth in LI, since no context can make that sentence false. Another are the LD logical truths expressible in the poorer language of LI, such as the one we discussed above, $(\phi \leftrightarrow @\phi)$.³⁰

There are also some super-valid arguments. Trivially, any argument that has the above super-logical truths as conclusion is valid, no matter what its conversational situation is. And, of, course, if one of the premises is a contradiction, the argument will be super-valid. Apart from such trivial super-valid truths, we will also get entailments between sentences whose truth value does not vary from day to day. For instance, ignoring tense, the following would be valid: “I was born on January 1st, 1980”, in any context, would super-entail “Someone was born on January 1st, 1980”. If the language was a bit richer, containing temporal operators, we would get more interesting results. For instance, presumably \langle “I sing”, “At some time, someone sings” \rangle would be a valid argument, since the day of the premise is irrelevant because of the universal temporal quantifier, and the day of the conclusion, though relevant for the truth value of the conclusion, is irrelevant given the content of the premise.

One might feel that super-validity should be as far as logic should go. Admittedly, this is another way to answer Soames’s challenge. His cases were ones on non-super-valid arguments. If we were only aiming for super-valid ones, we could allow context to vary as much as they wanted, and for validity we could as well quantify over all contexts, of any form. We would still be liberalizing LD, by allowing context to vary; but validity would be similar to LD, because the form of conversational situations would not be of any interest.

³⁰Other examples: $(\phi \leftrightarrow \text{Today } \phi)$ and, more controversially, $(\exists x)x = x$. The latter is a logical truth in most presentations of FOL, because the logic is much easier if empty domains are not allowed in the model theory. But this is no more than convenience; whereas for LD and LI, the prohibition of empty domains comes from deeper within the theory, namely from the requirement that there always be a speaker in any context. Note that the prohibition only applies to the actual world of any model; empty possible worlds are allowed, and they will matter for modal matters, since what is said by an actual utterances can be evaluated relative to them.

But I don't think that this is the right way to respond to Soames's challenge. The balance of payments would not be in our favour. We would regain the freedom to generalize over all contexts, and thus we would regain the nice feature of LD that looking only at the sentences in an argument would be sufficient for assessing validity, at least once we settled on the right notion of context. But we would lose the connection to actual arguments and conversations, and we would ignore the features of indexicals that we depend on when we participate in those arguments. We would leave parts of the semantics of indexicals out of our logic; that is too high a price to pay.

3.4.2 LI and Analyticity

So far, I have been talking about analyticity and truth in virtue of meaning fairly loosely. I did narrow matters some, since I argued that it really should amount to truth in virtue of meaning plus other semantic theses about truth, contexts, circumstances of evaluation, etc. Still, I don't think I have strayed too much from the folk philosophical picture of analyticity.

Even with this simple minded characterization, a problem looms: logical truths are generally thought to be analytic truths, at least by those who find analyticity a useful notion. And if we extend the notion to cover arguments, as I proposed, valid arguments are taken to be analytically valid. *Prima facie*, this doesn't seem to be true of LI. Analyticity is supposed to be a feature of sentences, or perhaps of thoughts; but arguments in LI are partly composed of contexts, which are chunks of the non-linguistic world. They contain people, and places, and moments in time, etc. So should we conclude that logic is not analytic? Or have we hit upon a reason to refuse to call LI a logic?

Before I offer my answer to these good questions, I need to make the notion of analyticity a bit more precise. Fortunately, contemporary literature provides quite a few ready-made notions, and even a superficial tour through the options ought to be enough

to clarify (and, I hope, dispel) the worry.³¹

A good history of analyticity ought to begin with Kant, or perhaps Hume, or Leibniz. But that would be a diversion; instead, I will begin with Frege, since his is the most relevant conception to my purposes:

The problem becomes, in fact, that of finding the proof of the proposition, and of following it up right back to the primitive truths. If, in carrying out this process, we come only on general laws and on definitions, then the truth is an analytic one, bearing in mind that we must take account also of all the propositions upon which the admissibility of any of the definitions depends. If, however, it is impossible to give the proof without making use of truths which are not of a general logical nature, but belong to the sphere of some special science, then the proposition is a synthetic one.³²

So, briefly, an analytic truth is one which can be proven from definitions and logic alone. Under this conception, the question whether logic is analytic is moot, since the analyticity is defined in terms of logic. Nor can we check a system for logicity by checking to see if it is analytic; at best, we can do it the other way round. The part about proofs has not survived as part of contemporary notions of analyticity; but it is instructive to look at Frege because, just like the more common vague notion of “true in virtue of meaning”, his conception can be taken in at least three ways: is it a thesis about a particular relation between any truth (roughly, derivability) and a privileged class? Is it a thesis about the justification that an ideal reasoner could have to believe these truths? Or is it a thesis about how we, non-ideal as we are, can come to know certain truths? Frege certainly was not proposing only the first option, as we can see from the paragraph immediately preceding the one just quoted:

³¹For a good brief historical overview of analyticity, see Juhl & Loomis (2010).

³²Frege ([1884] 1980, p. 4).

[T]hese distinctions between a priori and a posteriori, synthetic and analytic, concern, as I see it, not the content of the judgment but the justification for making the judgment. Where there is no such justification, the possibility of drawing the distinction vanishes.³³

We will come back to apriority in §3.4.3. For now, I want to focus on the claim that analyticity is an epistemic feature, related to the reasons for taking thoughts to be true.³⁴ Depending on how we develop this idea, it may turn out that we can test the thesis that LI is a logic by figuring out if it prescribes correct *laws of thought*, and hence whether what it deems to be valid arguments really deserve that status.

The important question here is not what Frege meant; my main interest is whether analyticity is an inherently epistemic notion. The most straightforward way to think of it as epistemically important is to think that it carries with it a guarantee that any regular Joe, if he is in possession of the right concepts, could be justified in believing an analytic truth, and perhaps even a guarantee that he could come to know it, if only he applied himself.³⁵ If this is incorrect, one could think that analyticity only provides such guarantees for an ideal reasoner, whereas the question of how we, humans, come to such truths would be a further, perhaps empirical, one.³⁶

But the connection between analyticity and epistemology can be severed. This was noted, for instance, by Dummett:

If [...] we agree with Frege in adopting a realistic interpretation of the sentences of our language, there is room for non-epistemic versions of these notions as well. To adopt a realistic interpretation is to hold that the sense of our

³³Frege ([1884] 1980, p. 3).

³⁴See also Frege ([1893] 1964, p. 12), where he says that the laws of logic “are the most general laws, which prescribe universally the way in which one ought to think if one is to think at all”.

³⁵This seems to be the view put forward in Boghossian (1996, p. 387).

³⁶One model here is Tyler Burge’s interpretation of Frege on analyticity and apriority. See Burge (1998, p. 340) and Burge (2000).

sentences is given in such a way as to relate to their determination as true or as false, independently of our capacity, even in principle, for recognizing what truth-value it has. Opposed to a realistic interpretation of this kind is any view which holds that the senses of our sentences are always given in terms of the means available to us for recognizing them as true or false".³⁷

Dummett makes his point more general than I need (and than I'd be comfortable with); he talks about the truth values of any sentence whatsoever.³⁸ I am only interested in this quote as applied to analytic truths: the meanings of (at least some of) our expressions may sometimes be so arranged that the sentence formed by those expressions is guaranteed to be true in virtue of those meanings alone. The question whether we, the language users, or even the ideal reasoner, can figure out this fact, is separate. We then have a number of options: we can be optimistic, and think that humans in fact can, or at least could, know all analytic truths, or we could be a bit less wide-eyed, and admit some constraints on our actual capacities.

But note that this leaves open the possibility that some analytic truths are simply epistemically uninteresting, for instance that some might be knowable only a posteriori. In fact, this claim has been made by Gillian Russell, who, roughly speaking, thinks of analytic truth as truth in virtue of reference determiner. To get an idea of her conception, it would be best to work on an example, which also relates to our topic. It involves Cassius Clay, who changed his name to "Mohammed Ali".³⁹ Suppose now that, when the new name was introduced, it went like this: "Let's use 'Mohammed Ali' to name Cassius Clay" (Russell (2008, p. 58)). Then, says Russell, the sentence "CC is MA" is an analytic truth, but it is of course only knowable a posteriori, and besides, there is no guarantee that users of the two names would know, or even be able to reason out, that the identity

³⁷Dummett (1981, p. 117).

³⁸The passage continues in a manner that I would definitely not follow, and links analyticity to necessity. As LD showed, there are analytic truths which are not necessary, so Dummett is wrong there.

³⁹Russell (2008, p. 65).

sentence is true.⁴⁰

If Russell is right about the semantics of “Cassius Clay”, we have a clear example of an analytic truth, i.e. a truth in virtue of (one kind of) meaning, which lacks any epistemically interesting properties. And even if she is wrong, I agree with the more cautious claim that I attribute to Dummett, that analyticity does not immediately entail any interesting epistemic property, and that a special argument needs to be made about any type of analytic truth if we want it to be interesting epistemically.

In particular, this shows that Boghossian’s notion of epistemic analyticity is not the only one available:

[A statement S is analytic] if mere grasp of S’s meaning by T sufficed for T’s being justified in holding S true.⁴¹

As Russell notes, the notion of meaning has been shown to be non-monolithic by the work of Donnellan, Kripke, and Kaplan, to unfairly name just a few. The consequence is that truth in virtue of one kind of meaning may have different properties from truth in virtue of another kind of meaning.⁴² So there is no reason to think that for any kind of meaning, merely grasping it is sufficient for anything. In fact, for some kinds of meaning, it is somewhat unclear what “grasping” would come to: just think of meaning in the sense of referent, or the content of a proper name according to direct reference theory.

The time has come to go back to LI, and the question whether its truths are analytically true, and if not, if that makes LI non-logical.

⁴⁰Strictly speaking, according to Russell, the sentence would be analytic but for the fact that the names might lack a referent; at one point, she calls our identity sentence “pseudo-analytic” (p. 105). The reason is that she takes the referent determiner of the original name, “Cassius Clay”, to be either a pointing (“Let’s call him (pointing) ‘Cassius Clay’ ” – p. 58) or a description (“the first born baby of Mrs Clay” – p. 114), and this is taken to entail that there are worlds in which the name lacks a referent. This seems mistaken to me; first, I doubt that the name was introduced by description, and if it was introduced by pointing (or, better, by a use of a demonstrative), it is directly referential, and cannot lack a referent in any world. But these matters are not directly relevant to the status of the identity sentence, so I put them aside.

⁴¹Boghossian (1996, p. 363).

⁴²Russell (2008, Ch. 2).

On the face of it, the epistemic notion of analyticity may well be unconnected to logic; or, at any rate, the connection does not immediately follow from the definition. Furthermore, once a connection is made, analyticity might be useful as a test for logicality.⁴³ This is not the notion I am interested in. It belongs within an important project to explain knowledge of logic, mathematics, and “conceptual truth”, of the kind that philosophy is sometimes said to be mostly about.

But, as I said, this is not the notion that best applies to the project I am working on: I am interested in the semantics of natural language, and its connection to logic. I don’t claim that the two projects are unrelated, of course. But if it turns out that LI truths are not epistemically analytic, that is an important and interesting discovery, which I would describe as discovering that not all parts of logic have a certain epistemic role, rather than a discovery that LI is not a logic. So I find more use for Dummett’s non-epistemic notion of analyticity.

This brings me to this section’s conclusion: if Kaplan’s semantics of indexicals is correct, and if his views about truth in a context go well together with my views about the nature of arguments, the logical truths of LI deserve to be called analytic. The biggest hurdle in the antecedent of my conclusion is the claim that the form of an argument is given partly by the form of its sentences, and partly by the form of its contexts. Talk of non-epistemic analyticity needs to be supplemented by an account of the role of contexts in logic. It is a big hurdle because if it is correct, analyticity is not just about sentences; I admit, that is a major departure from tradition. But it is forced on us by considerations that come from semantics, and those cannot be discarded when talking about non-epistemic analyticity.

⁴³In fact, Boghossian uses it as part of an account of the nature of knowledge of logic, rather than an account of the nature of logic. The program is developed in more detail in Boghossian (2000).

3.4.3 LI and Apriority

Logic is supposed to be a priori. Not by everyone, of course; Quine (1953) presented a picture attractive to many, in which logic is just as empirical a part of science as, say, biology. But the picture does not sway everyone, and those who stand their ground against Quine may object to LI on account of its seeming a posteriori nature.

There is one link between the a priori and logic which holds for LI as much as it holds for any logic: I argued for the semantic definition of validity not by looking at the world, or by polling English speakers, but by thinking about the semantics of indexicals, and by looking at how their semantic powers can be used in arguments. I hope that part of the appeal of LI is that it captures something central about the way English speakers use indexicals; like the logical rules for “and”. LI talks about and has rules for the first person pronoun in exactly the same way as FOL talks about and has rules about conjunction. The a priori credentials of the semantic theory underlying LI are just as good as the ones of FOL.

I have in mind one particular way to object to LI on apriority grounds. Recall the argument I presented above: let ϕ be “I am right-handed”, and ψ be “You are right-handed”. Let c_1 be (S, A, D, W) , and c_2 be (A, S, D, W) , and suppose that $A \neq S$; that is, c_2 is exactly like c_1 , except that the speaker and the addressee are reversed. Now take $\langle [c_1, \phi], [c_2, \psi] \rangle$, which is a valid argument in LI; call it “ \wedge ”.

It might be claimed that in order to know that \wedge is a valid argument, one would need to know something about the world: two contexts are parts of \wedge , and surely you need to know the parts that make up an argument in order to assess it for validity. Strictly speaking, this is not right. For validity, all that matters are certain abstract features of the conversational situation, such as whether the speaker of c_1 is identical to the addressee of c_2 , not who the speaker is, what day it is, and so on. In fact, that is how I presented the argument; I didn’t say what day it was, what possible world, or who the speaker was. In that, I admit I was a bit sneaky. What I presented was an argument *form*. My notion of an

argument is such that you have not specified one before you specify all its parts, e.g. who the speaker of the conclusion is. But I presented the argument to the extent that I gave enough information about it to be able to tell that it is valid. All extra knowledge is surely useful in conversation; it is good to know who you are speaking to, but it is not relevant for validity.

So: knowledge that all arguments of the form Λ are valid does not quite require knowing everything about its contexts. Still, even knowledge of the comparatively abstract features relevant for validity is (based on) empirical knowledge. You cannot know a priori that the speaker of the premise is the addressee of the conclusion; nor that the day has not changed mid-argument. So knowledge of validity is *a posteriori*. Doesn't this show that something is amiss with LI?

No, it does not. The apriority of logic is about knowing or justifying its principles. As we noted above, knowing the principles of LI seems no different from knowing the principles of FOL. And it would not help to claim that the validity of Λ is not a priori enough. Once we are given Λ , that is, once we are told what the form of the conversational situation is, the definition of validity takes over, as a priori and formal as you like.

The empirical appears when we go about in the world, and try to see if what people do constitutes a valid argument. The empirical bit is the part where you *formalize* what happens in the world. That is not a new feature of LI. When we symbolize my saying "John is hiding and he's not coming out" in FOL as $H_j \& \neg O_j$, that is based on taking "he" to be anaphoric on "John" from earlier in the sentence. That is an *empirical hypothesis*; how else could we know that the speaker didn't intend to use the demonstrative "he"? Even if it is derived from general conversational principles, it is still a defeasible hypothesis; I find both readings to be available in most uses of sentences like it.

There is a difference between FOL and LI, of course. Formalization in FOL pertains only to the logical form of *sentences*, to the conversational thread; in LI, it also pertains to certain formal features of non-linguistic bits of the world. But the extent of the in-

volvement of the empirical is the same. When we move from what we do to what logic treats, we move between the realm of the empirical to the realm of the logical. That move is bound to be laden with empirical knowledge. Once we are within the realm of logic proper, everything is just as formal in LI as in FOL. Call the last bit a priori or not, depending on your theories of logic and the a priori. But you cannot separate FOL from LI in this respect.

3.5 A Different Version of LI

There are many ways to achieve the goals of LI, and in this section I want to talk about one of them, and I will call the result “LI*”. The differences are superficial, technically speaking, but LI* allows for an easier comparison between a logic that allows context changes and LD, and it provides an easy way to recapitulate the lessons from §3.4 about the unexpected features of LI.

The three main technical innovations of LI have been the following: redefining the notion of an argument as a sequence of context-sentence pairs (as opposed to just a sequence of sentences), defining similarity between conversational situations (i.e. sequences of contexts), which in turn helped define validity not just as guaranteed truth in all contexts (as in LD), but as guaranteed truth in all *similar* conversational situations.⁴⁴

LI* sticks to the traditional notion of an argument: just a sequence of sentences, with the last one thought of as the conclusion (to avoid ambiguity, I will call these “arguments*”). In my terminology, an argument* is simply a conversational thread.⁴⁵ LI* will keep the notion of similarity from LI; but it will change the definition of validity, so that arguments (i.e. conversational threads) are said to be valid* only relative to a particular conversational situation. More formally,

⁴⁴I use “guaranteed truth” to mean: if the premises are true, the conclusion is true.

⁴⁵Recall that I talk of sequences of sentences, instead of sets, only because it makes exposition easier. I make no use of this peculiarity, philosophically speaking.

Definition 7. An argument* $\langle \phi_1, \phi_2, \dots, \phi_n \rangle$, where we think of ϕ_n as the conclusion, is *valid* relative to*, that is, *in*, a conversational situation $CS = \langle c_1, c_2, \dots, c_n \rangle$ iff for any conversational situation $CS' = \langle c'_1, c'_2, \dots, c'_n \rangle$ similar to CS , if $(\forall i \in [1, n - 1])\phi_i$ is true in c'_i , then ϕ_n is true in c'_n .

The main reason I prefer LI over LI* is that in LI, logic is analyzing the things we take part in, that is, arguments made up of utterances, which take place at a time, a place, etc. Arguments as sequences of sentences in their own contexts thus mimic best those things that our intuitions are about. A second, and much less important, reason is that we can see a clear path from LD-validity to LI-validity: instead of quantifying over all context sequences of a certain, very restrictive type (i.e. those formed by having the same context throughout), we liberalize the domain of quantification (all sequences of contexts), and then restrict it to similar conversational situations, so that we don't run afoul of Soames's challenge. On the other hand, it is much clearer why LD is a special case of LI*: since in LD we only look at validity relative to a single fixed context, and since all contexts are formally alike, this relativity need not be mentioned. Once we move to LI*, we see that the notion of validity relative to a sequence of contexts was implicit in LD.

The biggest advantage of LI* over LI is that we preserve the notion of an argument, familiar from FOL, as made up of sentences, and nothing else. This gets rid of questions about what it takes to know what argument* we are making: it's all about knowing the logical form of the sentences used. Recall that in LI, participants always risked not knowing what argument they were making, by not knowing who was speaking at some point, or what day it was, etc. In LI* we still have a version of that problem, in that participants are not guaranteed to know the logical form of the sentence being used; just think of ambiguities, or, even simpler, hearing problems. But we are used to these issues from FOL, so we put them aside, as customary.

The epistemic problems don't go away by the stroke of a pen, of course; LI* just shifts them from arguments to validity. According to LI*, in order to know if an argument* they

are making is valid, the participants need to know if it is valid *in their own conversational situation*. And, just as before, their epistemic access to the conversational situation is not provided merely by their being participants. The fact that it is their own situation that matters is reminiscent of the definition of truth *simpliciter* in LD, inherited by both LI and LI*: a sentence ϕ uttered in context c is true just in case the proposition expressed is true relative to c . In the same way, when an argument* made in conversational situation CS should be checked for validity relative to CS.

LI* shows nicely where worries about analyticity and apriority come from: since validity* is only a relative notion, and since it is relative to sequences of contexts, our knowledge of validity* depends on knowledge about non-linguistic matters. What we do know a priori is that, given an argument*, it would be valid* relative to conversational situations with certain formal properties, and invalid relative to others. But in order to know the form of a conversational situation here and now, we need to know something about the world, the participants in our conversation, and our positions in the world.

We also see clearly the conceptual gap between LD and a logic that allows context changes: validity* is not defined over all contexts. Instead, it takes a look first at the situation in which the argument* is made, and then quantifies over a similarity class of context sequences. So we pay attention, on the one hand, to a particular conversational situation, and only then do we quantify, and even that is restricted by information regarding the form of that situation that is determined by the facts of that situation. This influence of the original situation on its formal representation is what brings many of the seemingly unsettling logical truths and validities which formed the topic of much of this chapter.⁴⁶

⁴⁶I should also note that LI* comes closer than LI to the notion of validity proposed in Georgi (2011), but there are still important differences. First, Georgi defines validity as a property of sentences relative to a single context, not to a sequence. Second, in my definition of validity*, I quantify over all context sequences which bear certain abstract relations; Georgi takes that very context, and never allows it to vary, depending instead on allowing different interpretations of non-logical vocabulary. Conceptually, this second difference trades on distinct conceptions of logic. However, it is not trivial to figure out how much this second difference matters “on the ground”, i.e. in what gets counted as valid and what does not.

3.6 Two Kinds of Logical Truth, Conversationally Speaking

I have argued that LI logical truths may seem unusual, but there are good reasons to accept them as such. Still, there is something to the feeling of unease towards some of them. My diagnosis is that the source is what we expect from logic, especially in its practical relation to the ways we talk and argue. As I pointed out above, this has more to do with symbolization, broadly speaking, what we might call “logicization”, or “formalization”, than with formal logic by itself, but that is not an objection; formal logic and semantics, as conceived here, may well be about some Platonic realm, but only via sublunar activities and language use. In this section, I propose a way to divide LI logical truths into two groups, based on independently needed differences between kinds of knowledge that we assume conversational participants have.

3.6.1 What We Expect, and What We Assume

My intuitive motivations for LI have been examples like the following:

- (4) Josh, the weatherman, says “It’s raining today”. The next day, Josh says “It rained yesterday”.
- (5) Jerry says to Elaine: “If you want to go to the movies, we’ll go to the movies.” Elaine replies: “I do want to go to the movies.”

Both are cases in which, I claim, the sentences have the makings of an argument because of relations between their respective contexts of utterance. The contexts by themselves, of course, cannot constitute an argument, just as two or three trees cannot constitute an argument. But neither can the sentences by themselves, for two reasons: indexicals need contexts to gain a semantic value, without which the sentences do not get a truth

value; and if the contexts did not fit just so, the truth of “It’s raining today” would have no logical connection to the truth of “It rained yesterday”.

As far as truth preservation due to the form of the sentences and the contexts, the two examples are alike. In this section, I want to draw your attention to one way in which they are not.

Let’s start with Josh. Suppose that, instead of the way I presented the situation, he was wrong, and two days passed between his assertions. Josh normally presents the weather every day. On Monday he correctly says “it’s raining today”. The next time he’s up, he speaks as if he got it right the previous day. But he momentarily forgot that on Tuesday they had a guest star present the weather, so his own services had not been needed. In any case, when he says “It rained yesterday” on Wednesday, he didn’t in fact manage to repeat what he had said on Monday (in Frege’s sense of “repeat”).

The new argument would not be valid, of course, since in some conversational situations it didn’t rain on Tuesday. So the premise would be true and the conclusion false. All the better for LI, since that is indeed its verdict. But I want to ask another question: what kind of mistake did Josh make? It seems to me that it is a plain case of a mistaken belief about the world, and his place in time. It is the same kind of mistake as one makes in thinking they are in Paris when they are in Las Vegas: they look a little alike, unless one looks carefully. Conversationally speaking, nothing went amiss with Josh. As an announcer on TV, he ought to be better informed about these things; but as a mere speaker, he is just wrong about the passage of time. “Was that yesterday, or the day before yesterday?” We all have pondered such questions.

My claim is that, conversationally speaking, most uses of “I” and “you” are governed by a stronger norm. If you don’t keep track of conversational participants, you are in a worse way than Josh. One could think of this as a Gricean norm, not in the sense that it follows from one of his conversational maxims, but in the sense that keeping track of participants is a prerequisite for making contributions suitable to the thrust and purpose

of the conversation.⁴⁷ After all, the purpose of the conversation is related to the individual purposes of the individual participants, and if you don't keep track of the participants, you are likely to make inadequate conversational moves. I hope this already strikes the reader as intuitive. In case it does not, let me make my point more slowly.

Conversations are the natural place for arguments like (5). One could discuss cases of talking to oneself, and we will come back to them shortly. But typically, context changes with respect to the participants can be found when the participants are participating in a conversation (hence their name), addressing each other. People may come into a conversation or leave it; people may be addressed and say nothing in reply; sometimes, someone else may reply for them. But in the simplest case, we have situations like Jerry and Elaine trying to decide whether they should go to a movie.

Before I elaborate, let me make one point clear. LI does not require that arguments be placed in a real conversation. Arguments are merely sequences of context-sentence pairs, however they come to be produced, whether in reaction to one another or not. So if I want to contradict something Kant said, I can do it, even though we are not talking. Even more dramatically, I can contradict Kant without meaning to: I can simply utter the negation of something Kant said (provided the contexts are related in appropriate ways; recall that the negation of a sentence in one context may have the same truth value as an utterance of that sentence in another context, as we saw in the case of Napoleon). So why am I talking about conversations? Because that is the typical setting for using indexicals, especially ones related to conversational roles, like "I" and "you".

Keeping track of who is speaking, and to whom, is, in one sense, a bit of world knowledge. But it is a special part of world knowledge, in that our linguistic actions depend on

⁴⁷Here, I have in mind the Cooperative Principle, as stated in Grice ([1975] 1989b, p. 26): "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". Note also that Gricean norms are also to be followed *ceteris paribus*; mistakes are always allowed, there are ways to recover from them, and there may be conversations in which the requirement to keep track of participants is overridden by some other purpose.

it in a much tighter way than on any other kinds of background knowledge. By contrast, keeping track of days is just like any other feature of the world that we expect people to be aware of. I expect you to remember some of what you did yesterday, and that you did it yesterday; but if you don't, or if you are unsure whether that was yesterday or the day before, we can fix the unclarity, and move on. If Jerry didn't realize that Elaine was replying to him, his failure would be much more grave from the point of view of the conversation (though it too can be fixed, of course).

Let me introduce a technical term: I will say that it is *assumed* that participants in a conversation keep track of all other participants. I don't mean that they know who they are, in some robust sense of *knowing who*, just that when X says "you" to Y, and then Y uses "I", it is assumed that everyone party to the conversation knows this *relational* fact about the two contexts: that the addressee of one step in the conversation is the speaker of the next. My claim can be put this way: it is not assumed, in this technical sense, that speakers keep track of days as a constitutive part of the conversation, i.e. as a conversational norm. I will say that knowing about the passage of days is *expected*, but not *assumed*. When it comes to logic, some LI validities depend on things we assume, and some depend on things we expect. Which brings me to my diagnosis: some LI validities that may seem unlogical depend on things we expect, and since the latter kind of knowledge is of a most ordinary, empirical kind, we feel unease at calling them *logical* validities. But at the metaphysical level, both kinds of truth depend on *formal* relations between contexts.

That there is a difference between conversational roles and relations between days ought to be clear. The first one is something that, to a large extent, we *do*: we take turns in speaking, we signal when we want to keep talking and when we want to interrupt, we perform speech acts. This is systematic action, a good source of data for systematic accounts. It is the topic of Conversational Analysis, as pioneered in Sacks et al. (1974); they point out, for instance, that turn taking is a kind of activity not specific to conversations, since it can also be seen in games, standing in line, etc. The management of turn

taking is a rule-guided, complex activity, and our use of person-indexicals rides on this independently important structuring of conversations.⁴⁸ Days just change, without a care for our conversations or our actions.

There could be two broad kinds of reasons for a scientist's attention to conversational roles, as opposed to changes in days. One is that our role is much larger in changes in speakers and addressees, which are created by our actions, than in day changes, which happen quite independently of us. The latter are really nothing very interesting; astronomy tells us about the Earth revolving around the Sun, and the human practice of dividing time into days is not all that complicated. The way this particular division was set up may constitute an interesting bit of intellectual history, but the way days actually change does not. Conversational roles, though, are, to a large extent, up to us. And not just in the sense that human action is what set up the conditions for there being conversational roles. It is also something we all do, every time we talk to others.

The other reason is that systematic changes in conversational roles affect the way we plan our contributions to a conversation much more than day changes. The latter matter for some indexicals, for tenses, and for any plans that have to do with the calendar. But keeping track of participants comes before it; it is not only important for settling on which indexicals to use, and which way to look when we speak, but also for what we say, what we can presuppose, and what we expect the effect of what we say to be.

Let me give just one example from speech act theory, to illustrate how robust is the notion of participant in a conversation. For instance, Clark & Carlson (1982) have argued that we ought to take into account not just speakers and addressees, but also the hearers, towards whom some speech acts may be directed. Their leading example was from *Othello*:

Othello, to Desdemona, in front of Iago and Roderigo: Come, Desdemona. (Clark & Carlson (1982, p. 332))

⁴⁸See Hutchby & Wooffitt (2008) for a good introduction.

Their claim is that Othello had several goals. One was to get the addressee, Desdemona, to come with him. But he also wanted to inform Iago and Roderigo of his request to Desdemona.⁴⁹ His utterance would have been no good had he been addressing them; the fact that they were merely third parties was what enabled him to perform the speech act of informing them of his request that way.⁵⁰ In English, there are no indexicals specialized on hearers qua hearers; so LI has nothing interesting to say about them. But knowledge about participants will be relevant to various parts of theories about language, and that supports the distinction between what we assume, and what we merely expect.

3.6.2 What We Assume about Ourselves as Participants

At this point, it would be tempting to think that, using the artificial lingo I introduced above, we *assume* that people know everything about participants, and we *expect* people to know about the passing days. So figuring out that certain utterances constitute valid arguments based on tracking participants would be epistemically more available than those tracking days. But things may not be so simple.

As we noted above, our topic is not really logical validities; it is ways to capture logical forms out of conversations. This is fraught with practical difficulties, as we already know from trying to symbolize English sentences. Indeed, even those difficulties come at various levels: we may mis-hear what was said; the speaker may pronounce a word in a way that is unfamiliar to us; due to inattention, we may think the speaker said something other than what she did, etc. And even when all this is taken care of, we still face ambiguities of form or vocabulary. Knowing who the participants are helps figure out the truth

⁴⁹A bit of context: at that point in the play, Othello needs to leave to fight off a threat against Cyprus. He wants to take Desdemona with him, and she agrees. They have an hour before his departure (and he wants her to come with him for that hour), and in the following morning Desdemona, along with the official orders, will be brought to his ship by Iago and his wife; so it is important that Iago know what the plan is. Roderigo is only a rejected suitor, and he plays no role in Othello's plans, though one might suspect Othello wanting to make it clear that Desdemona was to come with *him*, not Roderigo.

⁵⁰For a fuller picture of Clark's views, see Clark (1996).

values of what is said: if one speaker says “I am Barack Obama”, but we don’t know who said it, we cannot evaluate what he said. As emphasized earlier, we don’t need that much to know if an argument is valid or not; we just need to *keep track* of participants, which is just as hard as figuring out the logical forms of the sentences uttered.

A more radical complication is that one might think that not all keeping track of participants is on the same level. I discussed earlier the case of Margaret Thatcher unknowingly talking to herself. If she says “I am identical to you”, LI counts that as a logical truth, one which the prime-minister is not at fault for not knowing. Here’s the problem: should we say that we *assume* that the speaker knows if she is talking to herself, just as we expect her to keep track of all participants, or should we say that we merely *expect* her to, on a par with expecting her to know if a day passed from her last speech in front of a mirror?

I am not sure how to answer that question, in large part because I see good reasons for both options. On the one hand, keeping track of yourself among mirror reflexions just doesn’t seem like it should be part of the kind of knowledge essential to conversational participation. On the other hand, when you keep track of others in a conversation, it might sometimes happen that you need to worry about mirror reflexions, shadows, fogs, strange angles, and other film noir paraphernalia. And, besides, we do almost always recognize ourselves in mirrors; they would be pretty useless if we did not.

But I do have two observations about the issue. First, the answer depends in great part on figuring out how to draw the distinction between assuming and expecting. It seems clear to me that participants, and especially their taking turns, are integral to a conversation in a way in which days and places are not. Knowing the former kind of fact is part of being in the conversation; knowing the latter is merely part of... keeping track of days and places. But this leaves many questions open. What is the relation between assuming, in my sense, and Stalnakerian common ground? Is there more to assuming than keeping track of participants? If Clark is correct, and some hearers are also important for speech act theory, does keeping track of them have the same status as keeping track

of speakers and addressees, even if no indexicals pick them out qua hearers? And how should we think of the first person plural pronoun, “we”? Should we think of assuming on the model of keeping track of anaphoric relations?

My second observation is that if we decide that Thatcher’s situation falls on the side of expecting, rather than assuming, this opens up a series of related questions (they may arise anyway, but perhaps less pressingly). In her case, the question was about the relation between speaker and addressee of a single sentence. What happens when we add more sentences? Is it the case that each speaker only needs to keep track of others, but not of him- or herself?

Take a conversational situation made up of four contexts, with the following relations between participants: every speaker addresses someone else, and the speaker of each of the first three contexts is identical to the addressee of the next context (that is, $(\forall i : 1 \leq i \leq 3) S_i = A_{i+1}$; just imagine a line of 4 people, each addressing the one in front, with the last one saying nothing). As agreed, we will say that the participants need to keep track of the last kind of fact, that the speaker becomes the addressee in the next step. If we should not assume that Thatcher knows if she is addressing herself, why should we assume that S_3 knows that he is the addressee of the second step? What makes facts about a single context different from facts about pairs of contexts? And if we decide to treat them the same, why is keeping track of oneself different from keeping track of others? After all, we are requiring of S_4 to keep track of everyone, as they say their respective sentences. Conversations are public and collaborative; shouldn’t all participants bear the same responsibility? Why should I assume that others keep track of *me*, if I don’t?

The point of these rhetorical questions is this: we (the conversation participants) are all in it (the conversation) together. We assume that everyone tracks each other, and therefore we assume the same of ourselves. And we know we can fail, just as we know the others can fail. But getting the structure of the conversation right is a prerequisite for participating in it successfully, and that places special epistemic constraints on us all.

List of Examples

- (1) John is happy. Therefore, John is happy.
- (2) Today is my birthday. Therefore, today is my birthday.
- (3) I'm hungry. If I'm hungry, I should eat. Therefore, I should eat.
- (4) Josh says "It's raining today". The next day, Josh says "It rained yesterday".
- (5) Jerry says to Elaine: "If you want to go to the movies, we'll go to the movies." Elaine replies: "I do want to go to the movies."
- (6) James Madison says to Napoleon: "I'm taller than you are". Napoleon replies (correctly) to Madison: "You're not taller than I am"
- (7) $\langle [c_1, \text{"It's raining today"}], [c_2, \text{"It rained yesterday"}] \rangle$
- (8) $\langle [c_1, \text{"I am hungry"}], [c_2, \text{"I am hungry"}] \rangle$
- (9) $\langle [c_3, \text{"You are hungry"}], [c_4, \text{"You are hungry"}] \rangle$
- (10) $\langle [c_5, \text{"I am hungry"}], [c_6, \text{"You are hungry"}] \rangle$
- (11) A says to B: "I am your sibling". B says to C: "I am your sibling". As a conclusion, C says to A: "I am your sibling".
- (12) $\langle [c_3, \text{"It will rain tomorrow"}], [c_4, \text{"It rained yesterday"}] \rangle$
- (13) $\langle [(Ernst Mach, Ernst Mach, D, W), \text{"I am identical to you"}] \rangle$

- (14) $\langle [(\text{William James, Ernst Mach, } D', W), \text{“You are a genius”}], [(\text{Ernst Mach, Ernst Mach, } D, W), \text{“I am identical to you”}] \rangle$
- (15) You, you, you, and you can leave, but you stay.
- (16) You owe a dollar to him and to her.
- (17) He owes a dollar to you and you.
- (18) That is bigger than that.
- (19) I am bigger than I am.
- (20) Now you see it; now you don't.
- (21) You're either seeing a dominantly green image now, [you flip the switch] or you're seeing a dominantly green image now.
- (22) [Jumping in one corner] It sounds more hollow here than [now I jump over towards the middle of the room] here.
- (23) [For any sentence ϕ] $\phi \leftrightarrow @\phi$
- (24) I exist.
- (25) I am here.
- (26) [For any sentence ϕ] $\phi \therefore @\phi$

Bibliography

- Almog, J. (2012). Referential Uses and the Foundations of Direct Reference. In J. Almog, & P. Leonardi (Eds.) *Having in Mind. The Philosophy of Keith Donnellan*, (pp. 176–184). Oxford University Press.
- Asher, N. (2000). Discourse structure and the logic of conversation. In *Current Research in the Semantics Pragmatics Interface*, (pp. 1–28). Elsevier.
- Austin, J. (1950). Truth. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 24, 111–128.
- Bach, K. (1992). Intentions and Demonstrations. *Analysis*, 52, 140–146.
- Bach, K. (2005). Context ex Machina. In Z. Szabó (Ed.) *Semantics vs Pragmatics*, (pp. 15–44). Oxford: Clarendon Press.
- Barwise, J., & Perry, J. (1983). *Situations and Attitudes*. Cambridge, MA: MIT Press.
- Bezuidenhout, A. (1996). Truth-Conditional Pragmatics. *Philosophical Perspectives*, 16, 105–134.
- Boghossian, P. (2000). Knowledge of Logic. In P. Boghossian, & C. Peacocke (Eds.) *New Essays on the A Priori*, (pp. 229–254). Oxford University Press.
- Boghossian, P. A. (1996). Analyticity Reconsidered. *Nous*, 30(3), 360–391.

- Braun, D. (1994). Structured Characters and Complex Demonstratives. *Philosophical Studies*, 74, 193–219.
- Braun, D. (1996). Demonstratives and Their Linguistic Meanings. *Noûs*, 30, 145–173.
- Bueno, O., & Shalkowski, S. A. (2012). Logical Constants: A Modalist Approach. *Nous*, (pp. 1–27). Published online on 22 June 2012. DOI: 10.1111/j.1468-0068.2012.00865.x.
- Burge, T. (1973). Reference and Proper Names. *Journal of Philosophy*, 70(14), 425–439.
- Burge, T. (1998). Frege on Knowing the Foundation. *Mind*, 107(426), 305–347.
- Burge, T. (2000). Frege on Apriority. In P. Boghossian, & C. Peacocke (Eds.) *New Essays on the A Priori*, (pp. 11–42). Oxford University Press.
- Burks, A. (1949). Icon, Index, and Symbol. *Philosophical and Phenomenological Research*, 9(4), 673–689.
- Caplan, B. (2003). Putting Things in Contexts. *Philosophical Review*, 112(2), 191–214.
- Cappelen, H., & Lepore, E. (2005). *Insensitive Semantics. A Defense of Semantic Minimalism and Speech Act Pluralism*. Malden, MA: Blackwell.
- Clark, H. H. (1996). *Using Language*. Cambridge University Press.
- Clark, H. H., & Carlson, T. B. (1982). Hearers and Speech Acts. *Language*, 58(2), 332–373.
- Corazza, E. (2004). *Reflecting the Mind. Indexicality and Quasi-Indexicality*. Oxford: Clarendon Press.
- Corazza, E. (2011). Indexicals and Demonstratives. In Marina Sbisa, Jan-Ola Ostman, & Jef Verschueren (Eds.) *Philosophical Perspectives for Pragmatics*, (pp. 131–152). Amsterdam / Philadelphia: John Benjamins Publishing Company.
- Corazza, E., Fish, W., & Gorvett, J. (2002). Who Is I? *Philosophical Studies*, 107(1), 1–21.

- Crimmins, M. (1995). Contextuality, Reflexivity, Iteration, Logic. *Philosophical Perspectives*, 9, 381–399.
- Donnellan, K. S. (1966). Reference and Definite Descriptions. *The Philosophical Review*, 75(3), 281–304.
- Donnellan, K. S. (1970). Proper Names and Identifying Descriptions. *Synthese*, 21, 335–358.
- Dummett, M. (1981). *Frege: Philosophy of Language*. Cambridge, MA: Harvard University Press, 2 ed.
- Elbourne, P. (2008). Demonstratives as Individual Concepts. *Linguistics and Philosophy*, 31, 409–466.
- Evans, G. ([1980] 1996). Understanding Demonstratives. In *Collected Papers*, (pp. 291–321). Clarendon Press. Originally published in *Meaning and Understanding*, edited by Herman Parret and Jacques Bouveresse, Berlin: De Gruyter, pp. 280–303.
- Frege, G. ([1884] 1980). *The Foundations of Arithmetic*. Evanston, Illinois: Northwestern University Press, 2nd ed. Translated by J.L. Austin.
- Frege, G. ([1893] 1964). *The Basic Laws of Arithmetic*. Berkeley, CA: University of California Press. Translated by M. Furth.
- Frege, G. ([1918] 1997). Der Gedanke. *Beiträge zur Philosophie des deutschen Idealismus*, 1, 58–77. Translation in (Ed.) M. Beaney *The Frege Reader* (pp. 325–345) Oxford: Blackwell (1997).
- García-Carpintero, M. (1998). Indexicals as Token-Reflexives. *Mind*, 107(427), 529–563.
- Gauker, C. (2008). Zero Tolerance for Pragmatics. *Synthese*, 165, 359–371.

- Georgi, G. (2011). *Demonstratives in Logic and Natural Language*. Ph.D. thesis, University of Southern California.
- Gómez-Torrente, M. (1996). Tarski on Logical Consequence. *Notre Dame Journal of Formal Logic*, 37(1), 125–151.
- Grice, H. P. ([1975] 1989b). Logic and Conversation. In P. Cole, & J. Morgan (Eds.) *Syntax and Semantics: Speech Acts*, vol. 3, (pp. 41–58). New York: Academic Press. Reprinted in Grice (1989) (pp. 22-40).
- Groenendijk, J., & Stokhof, M. (1991). Dynamic Predicate Logic. *Linguistics and Philosophy*, 14(1), 39–100.
- Hacking, I. (1979). What Is Logic? *Journal of Philosophy*, 76(6), 285–319.
- Higginbotham, J. (1988). Contexts, Models, and Meanings: a Note on the Data of Semantics. In R. K. Kempson (Ed.) *Mental Representations. The Interface between Language and Reality*, (pp. 29–48). Cambridge University Press.
- Hutchby, I., & Wooffitt, R. (2008). *Conversation Analysis*. Malden, MA: Polity, 2nd ed.
- Jeshion, R. (2012). Names not Predicates. Manuscript.
- Juhl, C., & Loomis, E. (2010). *Analyticity*. London: Routledge.
- Kaplan, D. (1979). On the Logic of Demonstratives. *Journal of Philosophical Logic*, 8(1), 81–98.
- Kaplan, D. (1986). Opacity. In L. Hahn, & P. Schilpp (Eds.) *The Philosophy of W.V. Quine*, (pp. 229–289). La Salle, Ill.: Open Court.
- Kaplan, D. (1989a). Afterthoughts. In J. Almog, J. Perry, & H. Wettstein (Eds.) *Themes from Kaplan*, (pp. 565–614). Oxford University Press.

- Kaplan, D. (1989b). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.) *Themes from Kaplan*, (pp. 481–563). Oxford University Press.
- Kaplan, D. (2005). The Meaning of Ouch and Oops. Manuscript.
- King, J. (2001). *Complex Demonstratives*. Cambridge, Mass.: MIT Press.
- King, J. C. (2012). Speaker Intentions in Context. *Noûs*. Published online in 2012.
URL <http://dx.doi.org/10.1111/j.1468-0068.2012.00857.x>
- Korta, K., & Perry, J. (2011). *Critical Pragmatics*. Cambridge University Press.
- Kripke, S. (1977). Speaker's Reference and Semantic Reference. *Midwest Studies in Philosophy*, 2, 255–276.
- Kripke, S. (1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press, 2nd ed.
- Kripke, S. (2008). Frege's Theory of Sense and Reference: Some Exegetical Notes. *Theoria*, 74(3), 181–218.
- Kripke, S. (2011). Vacuous Names and Fictional Entities. In *Philosophical Troubles. Collected Papers, Volume I*, (pp. 52–74). Oxford University Press. First delivered in 1973.
- Larson, R., & Ludlow, P. (1993). Interpreted Logical Forms. *Synthese*, 95(3), 305–355.
- Lewis, D. (1981). Index, Context, and Content. In S. Kanger, & S. Ohman (Eds.) *Philosophy and Grammar*, (pp. 79–100). Dordrecht, Holland: Reidel.
- Locke, J. (1975). *An Essay Concerning Human Understanding*. Oxford: Clarendon Press. (Original work published in 1700.).
- MacFarlane, J. (2008). Truth in the Garden of Forking Paths. In M. García-Carpintero, & M. Kölbel (Eds.) *Relative Truth*, (pp. 81–102). Oxford University Press.

- MacFarlane, J. (2009). Logical Constants. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Fall 2009 ed.
- Mount, A. (2008). Intentions, Gestures, and Salience in Ordinary and Deferred Demonstrative Reference. *Mind and Language*, 23(2), 145–164.
- Pelczar, M., & Rainsbury, J. (1998). The Indexical Character of Names. *Synthese*, 114(2), 293–317.
- Perry, J. (2001). *Reference and Reflexivity*. Stanford University: CSLI.
- Predelli, S. (1998). I Am Not Here Now. *Analysis*, 58(2), 107–112.
- Predelli, S. (2005). *Contexts*. Oxford University Press.
- Predelli, S. (2011). I Am Still Not Here Now. *Erkenntnis*, 74(3), 289–303.
- Quine, W. V. O. (1953). Two Dogmas of Empiricism. In *From a Logical Point of View*, (pp. 21–46). Cambridge, MA: Harvard University Press.
- Quine, W. V. O. (1968). Ontological Relativity. *The Journal of Philosophy*, 65(7), 185–212.
- Quine, W. V. O. (1982). *Methods of Logic*. Cambridge, MA: Harvard University Press, 4th ed.
- Recanati, F. (1993). *Direct Reference: From Language to Thought*. Oxford: Blackwell.
- Recanati, F. (2002). Unarticulated Constituents. *Linguistics and Philosophy*, 25(3), 299–345.
- Recanati, F. (2007). Indexicality, Context and Pretence: A Speech-Act Theoretic Account. In N. Burton-Roberts (Ed.) *Pragmatics*, (pp. 213–229). New York: Palgrave Macmillan.
- Reichenbach, H. (1947). *Elements of Symbolic Logic*. New York: Macmillan.
- Reimer, M. (1992). Three Views of Demonstrative Reference. *Synthese*, 93(3), 373–402.

- Rumfitt, I. (2010). Logical Necessity. In B. Hale, & A. Hoffmann (Eds.) *Modality. Metaphysics, Logic, and Epistemology*, (pp. 35–64). Oxford: OUP.
- Russell, G. (2008). *Truth in Virtue of Meaning*. Oxford University Press.
- Russell, G. (2010). A New Problem for the Linguistic Doctrine of Necessary Truth. In N. J. L. L. Pedersen, & C. D. Wright (Eds.) *New Waves in Truth*, (pp. 267–281). Palgrave Macmillan.
- Ryle, G. (1949). *The Concept of Mind*. London: Hutchinson's University Library.
- Sacks, H., Schegloff, E. A., & Jefferson, G. (1974). A Simplest Systematics for the Organization of Turn-Taking for Conversation. *Language*, 50(4), 696–735.
- Salmon, N. (1998). Nonexistence. *Nous*, 32, 277–319.
- Salmon, N. (2002). Demonstrating and Necessity. *Philosophical Review*, 111, 497–538.
- Scott, D. (1970). Advice in Modal Logic. In K. Lambert (Ed.) *Philosophical Problems in Logic*, (pp. 143–173). Dordrecht: Reidel.
- Sher, G. (1991). *The Bounds of Logic: A Generalized Viewpoint*. Cambridge, MA: MIT Press.
- Smith, Q. (1989). The Multiple Uses of Indexicals. *Synthese*, 78(2), 167–191.
- Soames, S. (2010). *Philosophy of Language*. Princeton University Press.
- Stanley, J. (2000). Context and Logical Form. *Linguistics and Philosophy*, 23, 391–434.
- Starr, D. (2010). *Essays on Perceptual Demonstrative Reference*. Ph.D. thesis, University of California, Los Angeles.
- Stojanovic, I. (2007). *What Is Said. An Inquiry into Reference, Meaning, and Content*. Ph.D. thesis, Stanford University.

- Strawson, P. F. (1950). Truth. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 24, 129–156.
- Strawson, P. F. (1960). *Introduction to Logical Theory*. London: Methuen & Co.
- Tarski, A. ([1936] 1983). On the Concept of Logical Consequence. In J. Corcoran (Ed.) *Logic, Semantics and Metamathematics*, (pp. 409–420). Indianapolis, IN: Hackett, 2nd ed. Translater by J.H. Woodger.
- Tarski, A. (1966). What Are Logical Notions? *History and Philosophy of Logic*, 7, 143–154.
- Travis, C. (1996). Meaning's Role in Truth. *Mind*, 105(419), 451–466.
- Warmbröd, K. (1999). Logical Constants. *Mind*, 108(431), 503–538.
- Wettstein, H. K. (1984). How to Bridge the Gap between Meaning and Reference. *Synthese*, 58, 63–84.
- Williamson, T. (2008). *The Philosophy of Philosophy*. Oxford: Blackwell.
- Wolter, L. (2009). Demonstratives in Philosophy and Linguistics. *Philosophy Compass*, 4(3), 451–468.