The Logic of Indexicals

Alex Radulescu

October 28, 2012

1 Introduction

One of the central arguments in Strawson (1960) against the possibility of a formal semantics of natural language and of a logical companion was that indexicals cannot be dealt with formally.[1] Kaplan's "Demonstratives", among its many accomplishments, provided a successful semantics and a logic for indexicals, and thus proved Strawson wrong about indexicals, and about context sensitive expressions in general. The present paper is motivated by the observation that Kaplan's logic is too limited: it only deals with arguments that take place in a single context. In a way, the situation is as if Strawson had been right all along, except for a very small, and arbitrarily chosen, class of arguments. Furthermore, this restriction goes against the power of Kaplan's semantics: indexicals are special precisely because they get different referents in different contexts, and we use this power in argumentation and conversation.[2]

Kaplan himself never said why he restricted logic in this way. But Soames (2010) did: if "I" can be used twice in an argument with different referents, arguments with a Modus Ponens form would come out invalid. Soames's conclusion is that we must force all indexicals to keep constant throughout an argument, and we can (and should) do it Kaplan's way, by excluding arguments with changing contexts.

In this paper, I offer a logic that does what Soames claimed to be impossible: it provides a definition of validity for arguments in which the contexts change. The response to Soames will require

[1] The nutshell of his argument was already present in Strawson (1950).

[2] In this paper, 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).

a re-thinking of the notion of the form of an argument, in order to take into account not just the logical form of the sentences involved, but also certain abstract relations between the contexts of the argument. Explaining what these relations are, and how to capture them, will occupy most of the paper.

2 How to Start Thinking about the 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.[3]

On a closer look, 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 are saying that he is happy by his standards, but in the conclusion we have moved on to judging him by more common standards, and then he will not count as 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 referentially 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. But they are not seen as presenting a challenge to the idea that "John" and "is happy" are the kinds of things that a logician can deal with.

Let us look at indexicals now. Consider this argument:

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

Is this a valid argument? One difficulty is this: if we are just looking at the two sentences, it is hard to tell what it even means to say that they form a *valid* argument. "Today" only gets its value

[3] We are speaking loosely for now. See [Demons], Remark XIX for a more detailed and careful presentation.

when placed in a context, and so does the first person pronoun. If we focus just on the sentences, we do not get truth values; and without truth values, we cannot discuss validity.

This shows that we must look not only at sentences, but at sentences in a context of utterance, so that, e.g. "today" does not just float freely in logical space, 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.[4] Would that be sufficient? Strawson thought not, and it is instructive to see why.

2.1 Strawson's Argument against a Logic of Indexicals

In Strawson (1960, pp. 211-213) we find one of his reasons for thinking that natural language cannot be treated formally. His argument begins with the claim that part of the job of a theory of meaning is to say of each sentence what other sentences it entails, which he takes to mean that it needs to say if the following conditional holds: if a particular sentence is used truly, then another particular sentence, if used in any circumstance, would also be true.[5] This is Strawson's view about the business of logic.[6]

But, and this is Strawson's central point, this is only part of what we need from a theory of meaning. Logic cannot do the whole job because indexicals have the special ability to get different semantic values when placed in different contexts. Indexicals are governed by what he calls "referring rules": "I" picks out the speaker, whoever he or she turns out to be, the present tense in simple sentences picks out the moment of time of the context, etc. And Strawson thinks that these referring rules are outside the realm of logic, since they involve *uses* of sentences, rather than sentences by themselves. Because they cannot talk about indexicals, entailment rules do not capture everything about our intuitive notion of validity, and, therefore, are not sufficient for a theory of meaning. For example, according to Strawson's definition of entailment, it turns out that (2) is not a valid argument.

[4] I want to avoid issues about what exactly worlds are, since this will be not important for the main issues I discuss. Worlds might be thought of as logically possible worlds or even as models. Relatedly, when I talk about validity as truth in all contexts, I really mean "in all contexts, in all structures", thus including a quantification over all interpretations of non-logical expressions. This should cause no confusion, and Kaplan himself talks this way in [Demons].

[5] I 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.

[6] "Formal logic is concerned with the meanings of sentences only in so far as these can be given by entailment-rules" (Strawson (1960, p. 214)).

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 any other use of that sentence. If the uses are, say, one week apart, but made by the same person, 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 does not depend on contextual features. He does not notice some trivial truths, ($p \vee \neg p$ 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 do we do with our intuitions about (2)? Strawson thinks that the best option would be to keep certain facts fixed (same speaker, same time and place), and thus get a valid argument. But, he adds, keeping facts fixed is not something logic can do: 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 a theory of entailment and referring rules.

2.2 Soames's Argument against Context Change

[Demons] showed that Strawson was wrong: we *can* have a logic of indexicals. The fundamental idea was to capture characters (Strawson's referring rules) formally, and to provide a systematic way in which indexicals get their semantic value from the context, thus allowing logic to get a grip on them.[7]

This innovation allows us to go back to (2) and account for its validity. Suppose that the premise and the conclusion are placed in the same context. This guarantees that the semantic value of the two occurrences of "today" and "my" are the same. This allows us not to worry about the actual truth values of the premise and the conclusion. Instead, we basically get what we had for (1): for any

[7] As 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 consisting in 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 does not correspond entirely to the idea of an actual use of a sentence by a person, since he allows talk about the truth value of a 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.

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.[8]

A crucial assumption in the previous paragraph was that the context stays the same for the whole argument. Strawson also relied on the idea that keeping the context fixed must be what our intuitive notion of validity presupposes. But we have not seen an explicit 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 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 ⊃ 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 us see how this argument works when applied to a simple practical syllogism:

(3) I am hungry. If I am 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 should not affect his lunch arrangements many years from now (around

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

2266, to be precise). The first premise is about me, while the second is about him, and Soames is right in saying that (3), interpreted this way, is not a valid argument.[9]

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 will 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.[10]

It looks like the choice is either to give up on the search for a logic of indexicals, or to force arguments to take place within a single context. But this would make Strawson mostly right. True, it would turn out that we can make referring rules part of a formal system. But in conversation we accept as valid more arguments than this logic seems 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 LD is just more formal work, which does not apply to the richness of "today" and "you".

3 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 need for an account of 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?

I will provide an answer to Soames's simple and powerful argument. But first, I would like to appeal to an intuition which is pulling in my direction, well illustrated by this quote from Frege:

> If a time-indication is conveyed by the present tense one must know when the sentence

[9] Soames 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).

[10] For reasons unrelated to indexicals, Soames in fact believes that there can be no logic of *English*, as 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), Braun (1996), Salmon (2002), and Caplan (2003).

> 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 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 is not 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.[11] But I would like to focus on the simple, intuitive part of Frege's point. Consider these situations:

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

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

(4) is just the kind of situation Frege mentioned. Suppose Josh is a weather reporter in Los Angeles. We are in the middle of July, and the big news is that it is raining. The next day the weather goes back to its regular, predictable sunny self, so Josh prefers to talk about the previous day. But he cannot use the same words; the tense and the indexical need to be changed. We will put tense aside for now, along with such issues as the fact that rain only happens in particular places, so the sentences need to be evaluated relative to the same place as well; just focus on the day-related 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, if the first sentence was true in its context, the second sentence *must* be true in its own context. This relation is remarkably like

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

validity, except that the argument only works given a certain relation between the contexts. But this is merely 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 on the previous day. Logic should not 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 this is true, surely he is 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 would like to focus on: there must be a valid argument lurking somewhere around (4).

The conversation in (5) works only 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 will 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 validly draw the conclusion: then, we'll go to the movies. 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".[12] The two utterances just do not connect; Jerry was addressing 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.

This shows that people take full advantage of the logical properties of indexicals: (5) is as valid as any classical Modus Ponens, even if its form is different. So our intuitions that there are logical relations involved in these context-changing conversations are quite robust, and they essentially involve the kind of changes that Frege was talking about. This is my reason for the search for a logic of indexicals: the goal is not to validate a particular view about semantics (though some of us would

[12]The background, which is not necessary for the logical point, but adds meat to the example, is that, as every fan of *Seinfeld* knows, Jerry really does not like Newman; he has no interest in Newman's desires, and he certainly does not want to go to the movies with him.

like that), but rather to account for the fact that we use indexicals as if our uses had logical properties.

Admittedly, Soames is right in saying that if we allow contexts to change in each and every way, we are 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. The second claim is now independent of the first, and may differ in truth value. The challenge, and my goal in this paper, is to construct a logic that validates my claims about (4) and (5), 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 Kaplan's semantics for indexicals 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").[13]

4 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", for instance, will corefer, because the speaker will be guaranteed to be the same in all the steps of the argument. So Modus Ponens will have the simple form we are accustomed to: if one premise is A and the other $\ulcorner A \supset B \urcorner$, we can safely infer B, whether the sentences contain indexicals or not. In fact, for logical purposes the contexts are unimportant; we only care that there is the one context, which does its job of supplying values for the indexicals. For discerning validity, we do not look at any particular context; instead, we need only look at the semantics of indexicals, and the nature of contexts in general.

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

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, the conclusion 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 should not 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 is more to an argument than the sequence of its sentences is a lesson not yet fully exploited.

We have not 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.[14]

Unlike in LD, 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 are part of the same conversation as the one started by Plato, or Thales. But we do not 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. In fact, it is quite amazing how much work in semantics has been on the semantics of instantaneous monologues.[15]

My proposal is to take seriously this banal feature of conversations: think of an argument as

[14] In 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. I take the notion of a possible world as a model to be an idealization, made for theoretical purposes, of the robust notion.

[15] Which 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).

a special type of conversation, in which the last remark is taken as the conclusion and the earlier remarks are taken as premises. 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:[16]

- An *argument* is a sequence of the form $\langle [c_1, \phi_1], [c_2, \phi_2], \ldots [c_n, \phi_n] \rangle$, where $n \geqslant 1$, and each pair $[c_i, \phi_i]$ is a pair of a context and a sentence.[17]

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".

Definition 1. Given an argument $\langle [c_1, \phi_1], [c_2, \phi_2], \ldots [c_n, \phi_n] \rangle$, the sequence of its sentences $\langle \phi_1, \phi_2, \ldots \phi_n \rangle$ 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 resemble contexts for utterances: they are the environment in which arguments take place.[18]

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

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,

[16] I use sequences, as opposed to a set of premises and a conclusion, strictly because the exposition is more convenient this way. LI does not need the full resources of systems like Dynamic Semantics (see, for instance, Kamp & Reyle (1993)), for which it is very important that it deals with sequences of sentences, not sets. On the other hand, something like LI could, and should, be incorporated in any dynamic system, because dynamic logics are deterministic (i.e. the changes are caused by the sentences being added to the conversation, and not by extraneous facts, like the day changing, or the speaker and addressee changing roles mid-argument), and the arguments LI is designed for are not. I should also note that Rumfitt (2010, p. 36–37) also thinks that steps in an argument are context-sentence pairs, but he does not allow contexts to vary.

[17] Note that I am 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. What we have in the text is just a necessary condition.

[18] "Situation" has another use, stemming from Barwise & Perry (1983). There, it means something like "a small part of a possible world". I will not be addressing their theory here, so there should be no danger of ambiguity.

since my goal is to show that we can represent context change with formal means. But behind the scenes, contexts are just what they really are: parts of the world, not abstract sequences.

In this paper, 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 will discuss later, a possible state of the world. I also do not 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.[19]

- 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.[20]

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:

(6) $\langle$[c_1, "It is raining today"], [c_2, "It rained yesterday"]$\rangle$

Let us 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 us fix on certain specific contexts. Let c_1 = (Josh, Kent, Sep 21, 2012, the actual world) and c_2 = (Josh, Kent, Sep 22, 2012, the actual world).[21,22] In order to say that this is a valid argument, we need a certain amount of generalization. After all, we want to know whether the conclusion *must* be true if the premise is true.

Our intuitions about (6) are simple. It is valid because:

[19] [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, so I leave the argument for another time.

[20] Note 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.

[21] Josh 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.

[22] This formulation is a bit sloppy, since days are represented in contexts by integers, and Sep 22, 2012 is not an integer. I do this for readability: strictly speaking, the day should be represented, say, by 1872359, the number of days since the latest creation of the world according to the Mayan calendar.

> 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 the given conversational thread is truth-preserving in all similar conversational situations. 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 could not 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 (6) would also be valid if we made $c_2 =$ (Fred, Marie, Sept 22, 2012, the actual world).

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. 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 an isomorphic relation between conversational situations.[23]

This idea is a weaker version of the constraint in LD, which 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 did not deserve

[23]Note that this notion of similarity is different from the semantic relationism of Fine (2007). The latter is a theory designed for occurrences of a particular singular term which corefer due to non-worldly facts. There are two clear differences. First, LI is interested in more than just mere coreference, as seen from our talk of the relations between days, as in example (4), where we are interested in the sameness of semantic value between different expressions, "today" and "tomorrow". Second, even when the same indexical is used coreferentially throughout an argument, that is in large part due to the relations between the contexts in which it is used, and these are worldly facts.

any special attention. This is why LD is a sublogic of LI. 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, lest it misses intuitively valid arguments, nor too loose, lest it rules in intuitively invalid arguments.

To see how this works, suppose we are interested in a given argument, in my sense, made up of sentence - context pairs. On the one hand, if we make it too easy for two conversational situations to count as similar, we do not get enough validities, and we are saddled with Soames's problem. On the other hand, if we impose strict conditions for two conversational situations to count as similar, we get too many validities, because we do not 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 (6). It did not rain on either day in Los Angeles; but that does not tell us that the argument is valid. [24]

If this sounds too technical, here is 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 is 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 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 a non-logical property.

[24] In Georgi (2011), 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, that might generate unwanted validities. The relation between modality and 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, as I try to show, technically unproblematic.

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.

5 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 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 paper I am only working with a few indexicals, the notion of similarity between conversational situations relevant for validity will have to look at relations between the respective components of contexts: speakers, addressees, and days.

I did not mention worlds just now, even though they are parts of contexts. Here I will follow [Demons], and require that worlds stay constant throughout an argument. The main motivation for this restriction is that I am 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. But we cannot 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 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 is nobody there.[25] 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:

[25] We 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.

Definition 3. A conversational situation

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

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 do not contain locations, which we do not need because we do not have location indexicals in our language. 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.

Doubts have been expressed about the constraints imposed on contexts in [Demons], and I will not address them here.[26] Some concern answering machine messages, which say "I am 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 do not 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. One solution is to search for a special semantics of recordings, but I will not do this here.[27]

5.1 Similarity and Days

Similarity for conversational participants looks at identities and non-identities between participants. This is the point of the Seinfeld example, so I won't make a special argument for it here. Identities also matter for days. This was already obvious from earlier examples:

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

[27] This is one way to think of as yet unpublished proposals made by Eliot Michaelson.

(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:

(6) $\langle$[c_1, "It is raining today"], [c_2, "It rained yesterday"]$\rangle$

(6) 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 is not enough to look at just one-day differences. Here is an argument which requires us to care about two-day differences:

(7) $\langle$[c_3, "It will rain tomorrow"], [c_4, "It rained yesterday"]$\rangle$,

where the day of c_3 is April 2, 2012, and the day of c_4 is April 4, 2012.

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 do not need to look at relations between more than two 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.

5.2 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

will not worry about changes in the worlds of contexts.

For any context c_i, let D_i be the day of c_i, S_i the speaker of c_i, and A_i the addressee of 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 $\langle c_1, c_2, \ldots c_m \rangle$ and $\langle c'_1, c'_2, \ldots c'_n \rangle$ 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 $\langle [c_1, \phi_1], [c_2, \phi_2], \ldots [c_n, \phi_n] \rangle$, where we think of $[c_n, \phi_n]$ as the conclusion, is *valid* iff for any conversational situation $\langle c'_1, c'_2, \ldots c'_n \rangle$ 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 too strong. Intuitively, the validity of (6) 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, 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.

6 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 discomforting 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 contain sentences which 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 logical truths, which did not come up in LD; the logic of "you" and "I" is more interesting than the logic of "I" alone.

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.[28]

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 LI, a sentence *in a particular context* is a logical truth just in case it is true in all contexts which are similar to that context.

The reason why the new definition does not 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]. One consequence is that not all contexts are similar to each other, because the identities and non-identities between the speaker and the addressee must 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 makes vivid one important novel feature of LI.

[28] Note that we talk of the sequence $\langle c_1 \rangle$, because we have defined similarity in terms of sequences of contexts. 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.

Before I present the problem technically, I would like to relate it to an older problem about self-knowledge. Put yourself in Mach's shoes, who did not 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.[29]

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

(8) $\langle$[(Ernst Mach, Ernst Mach, D, W), "I am identical to you"]$\rangle$

This is a soliloquy: the speaker is identical with the addressee. By claim (d) in our definition 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 (8). 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 is how one may become unsettled: begin by granting that it is a necessary truth; after all, true identities with directly referential terms are always necessary. If you are in a conciliatory mood, you may admit that most of the time, the speaker knows when she is addressing herself. But Mach would have denied (8); is that sufficient to show that Mach did not know a *logical* truth, as my theory claims? In other words, 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 disagree. 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". If this argument is fine when lecturing others (using "you"), why should not it be good when lecturing myself? The issue is complicated, since it presupposes a particular conception of the relation between LI and live conversations, and about the proper scope of

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

logic. I leave this for further work. 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 have 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 (8) 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 (8), it was a logical falsehood; now, it is logically contingent.

It might look as if we 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 (8) for a conclusion, but adds some premises, remains valid:

(9) $\langle$[(William James, Ernst Mach, D', W), "You are a genius"], [(Ernst Mach, Ernst Mach, D, W), "I am identical to you"]$\rangle$

Note that what is said in the premise is irrelevant, as are 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, (9) 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, the conclusion is guaranteed to come out true.

Now I can make my point: is (9) any different from (8)? What did Mach not know just before his realization? That he was looking at himself, in Perry's example. According to the original definition

of similarity, that makes him ignorant of two logical facts: that (8) was a logical truth, and that (9) was a valid argument. According to the version with (d*), he would have not known a logically contingent truth, namely that (8) is true, but he would still have not known that (9) 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 does not guarantee self-recognition, or tracking of days.

7 Why Not a Logic on Contents?

One of the distinctive features of LD was that some logical truths turned out to be contingent. Kripke (1980) had already presented some truths which he argued were necessary, but a posteriori, such as "Water is H_2O". Since this example does not really look like a logical truth, this already shows that necessity and logical validity are not coextensive.

> One surprising consequence of [Demons] is that not only are there some necessary truths which are not logical truths, but the entailment does not 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".[30] These sentences are logical truths in LD, 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 contingent logical truth.

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. As we saw earlier, LD 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.

[30] More precisely, any sentence of the form $\phi \leftrightarrow @\phi$ is a logical truth in LD.

Necessity, according to [Demons], applies not to sentences, but to what sentences say, i.e. propositions. For instance, if I say "I am here now" in Los Angeles on October 1, 2012, I express the proposition that I, Alex, am in that city, on that day. But I could have been somewhere else on that day. 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 interesting 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 LI 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?[31]

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. Each deserves a longer discussion, but for now, let me make two simple points, which suffice to show that LI has some advantages.

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, every argument of the following form is valid in LI: $\langle [c_1, @\phi], [c_1, \phi] \rangle$. 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 is 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 ϕ is true

[31] There 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 express the same content.

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.[32]

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.[33] 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.[34]

Here is another example, using an indexical: "I am Renee Kaplan", in the context $\langle$Renee Kaplan, David Kaplan, $D', W\rangle$, expresses a proposition of the same form as above, $x = x$, 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). There are two independent reasons. First, there are contexts in which the speaker is not Renee Kaplan, so the sentence is not true in all contexts, even if we keep the interpretation of the proper name fixed. And second, because we do not keep the interpretation fixed; we define validity as truth preservation in all contexts *in all interpretations*.

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.

[32]This 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.

[33]This 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.

[34]If you are not a Millian, pick any two expressions your favorite theory counts as synonymous, and run the same argument. If your favorite theory is so strict that it allows no synonymous expressions, it is wrong, and should not be your favorite to begin with.

8 Conclusion

In this paper, I have showed that Kaplan's LD is not the right logic for indexicals, because it deals only with arguments which take place in a single context. I then proposed a new logic, LI, which has LD as a special case, and which allows us to look at the full range of uses of indexicals, both by themselves, and as part of arguments.

This paper is a proof of concept, since LI only deals with a few expressions, and consequently it looks at just a small part of the kinds of similarity between contexts. For instance, the vocabulary lacks some prototypical indexicals ("here", "now"), it has no modal or temporal operators, and, perhaps even more glaringly, no demonstratives. Adding some of these expressions is not trivial, since they all bring their own problems. For instance, if we think of "now" as picking out an interval of time, not just a moment, and of "here" as picking out a region of space, should we think of them as demonstratives, or as indexicals? And what are the conditions under which two occurrences of "now" corefer?

Even if we focus only on the vocabulary of LI, it is obviously not true that sentences change truth value only once a day. So, for instance, an argument which just contains the sentence "It is raining" repeated twice is not valid merely because the two sentences were uttered on the same day; what is needed is that they be uttered at the same time, or in the same interval of time. I left these problems aside in this paper, for the sake of brevity. But they remain issues that need attention.

Furthermore, the kind of results that LI gives us can (and should) be embedded in dynamic semantics, by enriching the notions of context, of context change, and of validity. One important feature of LI, in fact, is that it is not dynamic in the sense of dynamic semantics. Changes of context in LI are not caused by the addition of sentences to a conversation, like in DRS, but by brute changes in the state of the world: the time passes, or the speaker moves, or changes roles with the addressee. Both kinds of formal system are needed, and it is not trivial to make them work together.

Many issues remain unsolved. But I hope to have pointed to the right path: there are many context sensitive expressions, and many relations between contexts, and LI teaches us that they need to be treated together, if we are to have any hope of capturing formally all uses of indexicals, and all the ways in which what we say depends on the context of utterance.

References

Asher, N. (2000). Discourse structure and the logic of conversation. In *Current Research in the Semantics Pragmatics Interface*, (pp. 1–28). Oxford: Elsevier.

Barwise, J., & Perry, J. (1983). *Situations and Attitudes*. Cambridge, MA: MIT Press.

Braun, D. (1996). Demonstratives and Their Linguistic Meanings. *Noûs*, *30*, 145–173.

Caplan, B. (2003). Putting Things in Contexts. *Philosophical Review*, *112*(2), 191–214.

Corazza, E., Fish, W., & Gorvett, J. (2002). Who Is I? *Philosophical Studies*, *107*(1), 1–21.

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.

Fine, K. (2007). *Semantic Relationism*. Oxford: Blackwell.

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).

Georgi, G. (2011). *Demonstratives in Logic and Natural Language*. Ph.D. thesis, University of Southern California.

Kamp, H., & Reyle, U. (1993). *From Discourse to Logic*. Dordrecht: Kluwer.

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: Oxford University Press.

Kaplan, D. (1989b). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.) *Themes from Kaplan*, (pp. 481–563). Oxford: Oxford University Press.

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.

Larson, R., & Ludlow, P. (1993). Interpreted Logical Forms. *Synthese*, *95*(3), 305–355.

Perry, J. (2001). *Reference and Reflexivity*. Stanford University: CSLI.

Predelli, S. (1998). I Am Not Here Now. *Analysis*, *58*(2), 107–112.

Quine, W. V. O. (1982). *Methods of Logic*. Cambridge, MA: Harvard University Press, 4th ed.

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: Oxford University Press.

Salmon, N. (2002). Demonstrating and Necessity. *Philosophical Review*, *111*, 497–538.

Smith, Q. (1989). The Multiple Uses of Indexicals. *Synthese*, *78*(2), 167–191.

Soames, S. (2010). *Philosophy of Language*. Princeton University Press.

Strawson, P. F. (1950). On Referring. *Mind*, *59*(235), 320–344.

Strawson, P. F. (1960). *Introduction to Logical Theory*. London: Methuen & Co.