Lessons from the Logic of Demonstratives:
what indexicality teaches us about logic, and vice versa

Gillian Russell

(penultimate draft—final version is forthcoming in *New Waves in Philosophical Logic*, Palgrave MacMillan 2012.)

Logic and linguistically-informed philosophy of language are becoming increasingly specialised disciplines. I think this is a consequence of the fact that each is making rapid progress, but it has *as* a consequence that it has become less routine for someone who is doing cutting-edge work on central topics in the one to also be doing cutting-edge work on central topics in the other. As a result, the topics which receive a lot of attention from logicians may be things that the average philosopher of language regards as fringe topics—e.g. logics of belief or substructural logics—and some central topics in linguistically informed philosophy of language—e.g. the structure of propositions, adverbs or contextualism—may be largely ignored in logic. The aim of the present paper is to provide some contamination in both directions. I will take one seminal idea from the philosophy of language and explore some consequences for logic (or at the very least, for the philosophy of logic) and then take one method from logic, and use it to prove a theorem that can help to explain and clarify some ideas in the philosophy of language.

My central topic is Kaplan's approach to context-sensitivity, as it is presented in "Demonstratives" and developed in his logic for demonstratives, **LD**. (Kaplan, 1989b;a)[1] Though that monograph contains some controversial claims, I think it is true that the central picture—according to which a sentence has a character which interacts with features of the context of utterance to generate the sentence's content—is now the standard model of indexicality; many of the disputes surrounding it concern, not *whether* it is correct, but rather for which expressions it is correct.[2] My overall thesis will be that the insights gained from

[1] It's common to use "context-sensitivity" in a slightly broader way than "indexicality" so that "context-sensitive" can apply to any mechanism whereby an expression conveys some information in one context that it doesn't convey in others, whereas "indexicality" is construed more narrowly to mean that the content of the expression varies with context. MacFarlane (2009) suggests even finer gradations in terminological use and I think those suggestions are good ones but I'll only be interested in one flavour of context-sensitivity in this paper, namely, indexicality.

[2] In my opinion the most formidable challenge comes from (King, 2001) which argues that Kaplan's model is wrong for complex demonstratives, such as "that man drinking a martini." Other authors have suggested that Kaplan's model applies to more expressions than just the indexicals and demonstratives for which it was originally intended, including knowledge ascriptions, truth-ascriptions and names.

considering this idea in the context of logic can aid both camps. In section 1 I will present Kaplan's ideas about indexicality informally, and in section 2 I will introduce his formal system, **LD**. Kaplan wrote "the most important and certainly the most convincing part of my theory is just the logic of demonstratives itself. It is based on just a few quite simple ideas, but the conceptual apparatus turns out to be surprisingly rich and interesting. At least I hope that you will find it so." (Kaplan, 1989b: 487–8) One of the features of **LD** that Kaplan liked was the (well-motivated) failure of necessitation:

> One of things that delighted me about indexicals was the convincingly deviant modal logic. As shown in *Demonstratives* the rule of Necessitation:
>
> If ϕ is valid then $\Box\phi$ is also valid.
>
> fails in the presence of indexicals. (Kaplan, 1989a: 593)

In section 3, I'll argue that this insight can be straightforwardly extended to a deviant account of logical consequence (except that I don't *really* think it is deviant.) Section 4 argues that **LD** also allows for a neat argument against the Linguistic Doctrine of Necessary Truth, and finally section 5 formulates and proves a theorem according to which (very roughly) no set of non-indexical sentences ever entails an indexical one. The hope is that this proof may be of use in the philosophy of language, both in understanding what the so-called "essentiality" of indexicals amounts to, and in narrowing down the set of indexical expressions.

1 Kaplan on Indexicals

Here is the picture of language with which we will be concerned. Sentences—most obviously sentences containing indexicals, such as *I*, *here* and *now*—can be used to express different propositions when uttered in different contexts. Thus it is expedient to think of the meaning of a sentence, not, in the first place, as a proposition, but as something which determines a function from contexts of utterance to propositions, even while we think of propositions themselves as determining functions from possible worlds (or more generally, circumstances of evaluation) to truth-values.

At an intuitive level, you can think of contexts of utterance as the situations in which someone is uttering a sentence (writing it, speaking it or maybe even signing it.) In the formal system we will simplify and represent them as quadruples $\langle a, p, t, w \rangle$ in which a is an agent (e.g. speaker), p a location, t a time and w a possible world.

Back at the intuitive level, a circumstance of evaluation is a situation against which you can assess a proposition for truth. The most obvious component of such a thing is a possible world, but some people also think that the truth-values

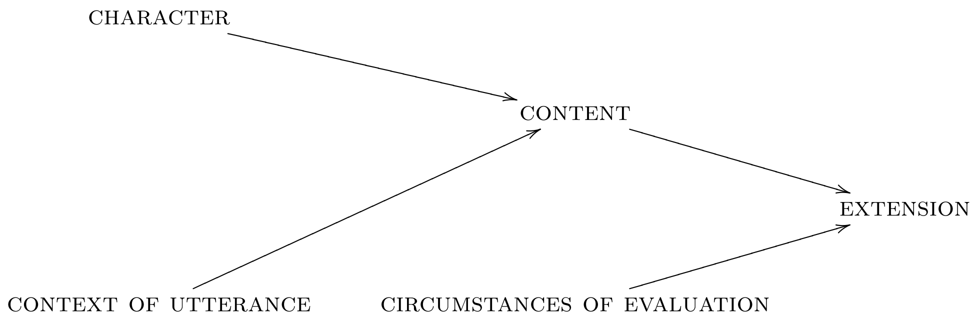

Figure 1: The Big Picture

of propositions vary with time, or even standards of taste, so that you have to assess a proposition for truth relative to an ordered pair of a world and a time, or a world and a time and a standard of taste. For now we'll follow Kaplan and assume that the truth-values of propositions vary with time and possible world only, and hence our circumstances of evaluation will be ordered pairs of a time and a possible world $\langle t, w\rangle$, and our propositional contents will be represented by functions from such pairs to truth-values.

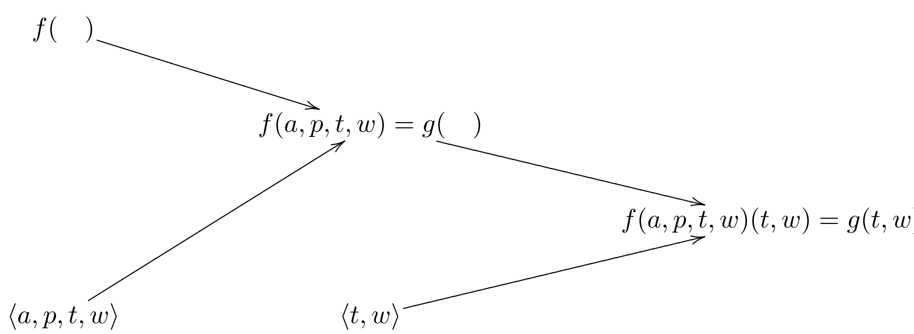

Figure 2: A More Formal Picture. f is the function determined by the expression's character, g the function determined by its content in the context (a, p, t, w))

I will illustrate this general picture with a more concrete example. Consider the sentence "I run." As a simple sentence of English, this is well-formed and meaningful—it has a character—but it won't express a proposition except relative to a context. Let's imagine that it gets uttered in a context in which I am the agent, the time is noon, the place is Kingston Town and the possible world is the real one. On Kaplan's view, the character for 'I' is given by the rule (which speakers learn) that 'I' refers to the agent of the context. Put this together with the context we've specified and we find that the content of 'I' is me, Gillian Russell. Since I am a person, and not say, a sense or an obviously abstract object, this tells us that Kaplan's contents are Russellian—like sets, they are abstract objects which may have non-abstract objects as components. For simplicity, let's assume that 'run' is not an indexical i.e. that its character yields the same content for each context of utterance, namely the property of running. Then the content of 'I run' relative to the context in question is the proposition represented by the ordered pair of me, and the property of running.

$$\langle \text{Gillian Russell}, \text{the property of running} \rangle$$

Relative to some circumstances of evaluation—the ones in which I am running—this has the extension *true*. Relative to others—such as the ones in which I am sitting, or in which I don't exist—it has the extension *false*.

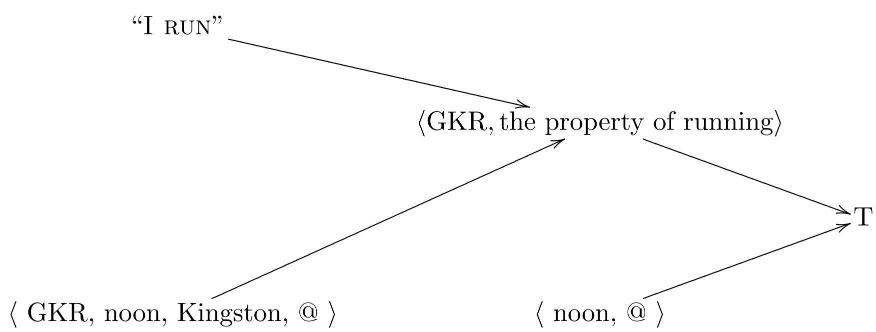

Figure 3: An illustration

The defining feature of context-sensitivity, on this view, is that the content of the expression varies with context of utterance.

2 The Formal System

In an attempt to clarify and systematise his view of context-sensitivity, Kaplan developed a formal logic, that is, a formal language with a model theory that allows a definition of logical truth for the language.

I will begin with a list of the different sorts of expressions in the language and the rules for combining them into well-formed formulas.[3] Our language will include two set of variables, those that range over locations, V_p, and the more familiar sort that range over individuals, V_i.[4] Atomic formulas are formed by combining the appropriate number of the appropriate sort of variables with a predicate. Since their are two sorts of term, the arity of a predicate is given by a pair of numbers $m - -n$, in which the first member is the number of individual-variables, and the second the number of location-variables the predicate takes to form a formula. There are three special 'logical' predicates: the $2 - -0-$place identity predicate, $=$, the $1 - -0-$place predicate, *Exists*, the $1 - -1-$place predicate *Located*. We also have an infinite number of $m - -n-$place i-functors (functors which form terms denoting individuals) and an infinite number of $m - n-$place p-functors (functors which form terms denoting locations.) The language contains the usual sentential connectives $\neg, \rightarrow, \leftrightarrow, \wedge, \vee$, and the quantifiers $\forall$ and $\exists$, the modal operators $\Box$ and $\Diamond$, tense logic operators F, P, G, and finally some expressions peculiar to **LD**: the $0 - -0-$place i-functor, I, the $0 - -0-$place p-functor, $Here$, and the operators N (now) and A (actually). We will also have the unusual $1-$place functor *dthat*.

Our formation rules are as follows:

1. (i) If $\alpha \in V_i$, then α is an i-term
   (ii) If $\alpha \in V_p$, then α is a p-term

2. If π is an m-n-place predicate, $\alpha_1, ..., \alpha_m$ are i-terms, and $\beta_1, ..., \beta_n$ are p-terms, then $\pi\alpha_1, ..., \alpha_m\beta_1, ..., \beta_n$ is a formula.

3. (i) If η is an m-n-place i-functor, $\alpha_1, ..., \alpha_m$ are i-terms, and $\beta_1, ..., \beta_n$ are p-terms, then $\eta\alpha_1, ..., \alpha_m\beta_1, ..., \beta_n$ is an i-term
   (ii) If η is an m-n-place p-functor, $\alpha_1, ..., \alpha_m$ are i-terms, and $\beta_1, ..., \beta_n$ are p-terms, then $\eta\alpha_1, ..., \alpha_m\beta_1, ..., \beta_n$ is a p-term.

4. If ϕ, ψ are formulas, then $(\phi \wedge \psi), (\phi \vee \psi), \neg\phi, (\phi \rightarrow \psi), (\phi \leftrightarrow \psi)$ are formulas.

5. If ϕ is a formula and $\alpha \in V_i \cup V_p$, then $\forall\alpha\phi$ and $\exists\alpha\phi$ are formulas.

[3] In fact we only need a subset of **LD**'s expressions for the present paper, and I will mention only what I need. The full system is on pages 541–542 of (Kaplan, 1989a)

[4] The role of individual constants (names) of the sort familiar to many logic students, is played by 0-place functors.

6. If α, β are either both i-terms or both p-terms, then $\alpha = \beta$ is a formula.

7. If ϕ is a formula, then $\Box\phi$, $\Diamond\phi$, $F\phi$, $P\phi$, $G\phi$, $N\phi$ and $A\phi$ are formulas.

8. If α is an i-term, then dthat$[\alpha]$ is an i-term and if α is a p-term, then dthat$[\alpha]$ is a p-term.

A standard Kripke-style structure for the language of *quantified modal logic* is a quadruple $\langle W, R, D, I\rangle$, in which W is a set of possible worlds, R an accessibility relation, D the domain of the model and I an interpretation function. The interpretation function assigns extensions to all the simple predicates and functors in the language, relative to a possible world from W (that is, it assigns intensions to the non-logical expressions.) That function is then extended to one which assigns intensions to every expression (including complex expressions and variables) via an assignment of objects to the variables in the language, and the rules for computing the values of complex expressions, given the values of their parts. In **LD** we will treat our modal operators in a very simple way and them rules which make no reference to an accessibility relation. Hence we can simplify the structures by dropping R. We do have two types of variables however and because of this we will require two sets to serve as domains of quantification, the set of individuals, U, and the set of locations P, giving us structures that are at least quadruples $\langle W, U, P, I\rangle$. Our language also contains tense operators, requiring us to specify a set of integers (times) T, for their interpretation, and finally the whole point of our present adventure is to consider what adding context-sensitive expressions to our language will do to the logic. Context-sensitive expressions create sentences whose truth-values are sensitive to context. Hence **LD** structures will also contain a set of contexts, C. For formal purposes, we take a context $c \in C$ to be a quadruple $\langle a, p, t, w\rangle$ where $a \in I$, $p \in P$, $t \in T$ and $w \in W$. The upshot is that a structure for **LD** is a sextuple $\langle C, W, U, P, T, I\rangle$. Each of the first five elements of that tuple must be a non-empty set. The final element, I, is the function which assigns appropriate intensions to the non-logical elements of our language. When it comes to context-sensitive expressions, these intensions will be assigned relative to a context:

Definition 1 (Interpretation (I)) *I is function which assigns to each predicate and functor an appropriate intension as follows:*

1. *if π is an m-n-predicate, I_π is a function such that for each $t \in T$ and $w \in W, I_\pi(t, w) \subseteq (U^m \times P^n)$*

2. *if η is an m-n-place i-functor, I_η is a function such that for each $t \in T$ and $w \in W$, $I\eta(t, w) \in (U \cup \{\dagger\})^{(U^m \times P^n)}$ (Note: $\dagger$ is a completely alien entity, in neither U nor P, which represents an 'undefined' value of the function. In a normal set theory we can take $\dagger$ to be $\{U, P\}$)*

3. *If η is an m-n-place p-functor, I_η is a function such that for each $t \in T$ and $w \in W$, $I_\eta(t, w) \in (p \cup \{\dagger\})^{(U^m \times P^n)}$*

4. $i \in U$ *iff* $(\exists t \in T)(\exists w \in W)(\langle i \rangle \in I_{Exist}(t, w))$

5. If $c \in C$, *then* $\langle c_A, c_P \rangle \in I_{Located}(c_T, c_W)$

6. If $\langle i, p \rangle \in I_{Located}(t, w)$, *then* $\langle i \rangle \in I_{Exist}(t, w)$

Truth and Denotation:

We write:	$\vDash^{\mathfrak{M}}_{cftw} \phi$	for	**ϕ, when taken in the context of utterance c (under the assignment f and in the structure $\mathfrak{M}$), is *true with respect to* the time t and the world w.**
We write:	$\lvert\alpha\rvert^{\mathfrak{M}}_{cftw}$	for	**the denotation of α, when taken in the context of utterance c (under the assignment f and in the structure $\mathfrak{M}$), *with respect to* to time t and the world w.**

In general we will omit the superscript 'M', and we will assume that the structure M is $\langle C, W, U, P, T, I \rangle$.

Definition 2 (Assignment) *f is an assignment (with respect to $\langle C, W, U, P, T, I \rangle$) iff:*

$$\exists f_1 f_2 (f_1 \in U^{V_i} \text{ and } f_2 \in P^{V_p} \text{ and } f = f_1 \cup f_2)$$

Definition 3 (Assignment-variants)

$$f^{\alpha}_{x} = (f \sim \{\langle \alpha, f(\alpha) \rangle\}) \cup \{\langle \alpha, x \rangle\}$$

(i.e. the assignment which is just like f except that it assigns x to α.)

This interpretation function is extended to one which assigns an intension to every (simple and complex) expression in the language via the following rules:[5]

1. If α is a variable, $|\alpha|_{cftw} = f(\alpha)$
2. $\vDash_{cftw} \pi\alpha_1...\alpha_m\beta_1...\beta_n$ iff $\langle |\alpha_1|_{cftw} \ldots |\beta_n|_{cftw} \rangle \in I_\pi(t, w)$
3. If η is neither 'I' nor 'here',
   $$|\eta\alpha_1 \ldots \alpha_m\beta_1 \ldots \beta_n|_{cftw} = \begin{cases} I_\eta(t,w)(\langle |\alpha_1|_{cftw} \ldots |\beta_n|_{cftw} \rangle) \text{ if none of } |\alpha_j|_{cftw} \ldots |\beta_n|_{cftw} \\ \text{are } \dagger; \\ \dagger, \text{ otherwise} \end{cases}$$

[5]For the following recursive definition, assume that $c \in C$, f is an assignment, $t \in T$ and $w \in W$:

4. i) $\vDash_{cftw} (\phi \wedge \psi)$ iff $\vDash_{cftw} \phi$ & $\vDash_{cftw} \psi$
ii) $\vDash_{cftw} \neg\phi$ iff $\sim\vDash_{cftw} \phi$
etc.

5. i) If $\alpha \in V_i$, then $\vDash_{cftw} \forall\alpha\phi$ iff $\forall i \in U, \vDash_{f_i^\alpha tw} \phi$
ii)If $\alpha \in V_p$, then $\vDash_{cftw} \forall\alpha\phi$ iff $\forall p \in P, \vDash_{f_p^\alpha tw} \phi$
iii)Similarly for $\exists\alpha\phi$

6. $\vDash_{cftw} \alpha = \beta$ iff $|\alpha|_{cftw} = |\beta|_{cftw}$

7. i) $\vDash_{cftw} \Box\phi$ iff $\forall w' \in W, \vDash_{cftw'} \phi$
ii) $\vDash_{cftw} \Diamond\phi$ iff $\exists w' \in W, \vDash_{cftw'} \phi$

8. i) $\vDash_{cftw} F\phi$ iff $\exists t'$ such that $t' > t$ and $\vDash_{cft'w} \phi$
ii) $\vDash_{cftw} P\phi$ iff $\exists t'$ such that $t' < t$ and $\vDash_{cft'w} \phi$
iii) $\vDash_{cftw} G\phi$ iff $\forall t'$ such that $t < t', \vDash_{cft'w} \phi$

9. i) $\vDash_{cftw} N\phi$ iff $\vDash_{cfc_Tw} \phi$
ii) $\vDash_{cftw} A\phi$ iff $\vDash_{cftc_W} \phi$

10. i) $|\text{dthat}[\alpha]|_{cftw} = |\alpha|_{cfc_Tc_W}$

11. $|\text{I}|_{cftw} = c_A$

12. $|\text{Here}|_{cftw} = c_P$

Now that we can say whether a sentence of **LD** is true on some interpretation, we use that notion to define logical truth in **LD**. Where Γ is either a term or a formula, we write:

$\{\Gamma\}^{\mathfrak{M}}_{cf}$	for	the Content of Γ in the context of utterance c (under the assignment f and in the structure $\mathfrak{M}$)

Definition 4 (Content)

If ϕ is a formula, $\{\phi\}^{\mathfrak{M}}_{cf}$ = that function which assigns to each $t \in T$ and $w \in W$, Truth, if $\vDash^{\mathfrak{M}}_{cftw} \phi$, and Falsehood otherwise.

If α is a term, $\{\alpha\}^{\mathfrak{M}}_{cf}$ = that function which assigns to each $t \in T$ and $w \in W, |\alpha|_{cftw}$.

Definition 5 (Truth with respect to Contexts) *ϕ is true in the context of utterance c, in the structure $\mathfrak{M}$ iff for every assignment f, $\{\phi\}^{\mathfrak{M}}_{cf}(c_T, c_W) =$ Truth.*

Definition 6 (Logical Truth) *ϕ is a logical truth ($\vDash \phi$) iff for every* ***LD*** *structure $\mathfrak{M}$, and every context c of $\mathfrak{M}$, ϕ is true with respect to c (in $\mathfrak{M}$).*

Kaplan immediately lists some logical truths of **LD** and these illustrate the system's "convincing deviance": $\vDash \alpha = dthat[\alpha]$, $\vDash \phi \leftrightarrow AN\phi$, $\vDash N$ (Located I, Here), and yet $\nvDash \Box\alpha = dthat[\alpha]$, $\nvDash \Box(\phi \leftrightarrow AN\phi)$, $\nvDash \Box N$ (Located I, Here).

At an intuitive level, here is what is going on: in order to be a logical truth, a sentence has to be such that no matter the context in which it is uttered, it is true at the world of that context. A sentence will be true at the world of the context, in the context, if and only if the proposition that it expresses relative to that context is true. So one way for a sentence to be a logical truth is for it to always express a proposition which is necessarily true. But if that were the only way, we would expect all logical truths to express necessary truths. Indexicality opens up another way to be a logical truth: a sentence may express different propositions relative to different contexts, but always be such that the proposition expressed relative to that context of utterance is true at the world of the context of utterance (even if it is not true relative to other worlds.)

$NLocated(I, Here)$ is an example of this second type: uttered in a context in which Sam is the speaker, the location is St Louis and the time is noon, it expresses the proposition that at noon Sam is in St Louis. Uttered in context in which Mary is the speaker, the location is Chicago and the time is 1pm, it expresses the proposition that at 1pm Mary is in Chicago. Neither proposition is a necessary truth; either person could have been somewhere else at that time. But the former proposition will be true with respect to the former context, and the latter true with respect to the later.

This revolutionary approach to logical truth is something that Kaplan draws out quite explicitly. But the concept of logical truth is often taken to be a special case of logical consequence, and so one might wonder what we can learn about logical consequence from **LD**.

3 Logical Consequence and Indexicality

A standard, if informal, definition of logical consequence is as follows: a sentence A is a logical consequence of a set of premises Γ if and only if *it is impossible for all the members of Γ to be true and A false.* **LD** shows us that this is a mistake, by which I mean, not just that it is a somewhat imprecise characterisation that requires more scholarly formal explication, but that is a step in the wrong direction.

Sometimes philosophers analyse the claim that it is impossible for the premises to be true and the conclusion false using possible worlds, for example they say that A is a logical consequence of Γ iff and only if every possible world in which every member of Γ is true is one in which A is true as well, or equivalently, that there is no possible world in which every member of Γ is true but A is false. It is a matter of great controversy what a possible world is exactly, but in logic this question is bracketed, much as the question of exactly what a number is can be bracketed when studying arithmetic. Some care is still required when speaking of possible worlds though. There *is* a sense in which the argument $Fa \vDash Fb$, in which a and b are names for the same object, is such that there is no possible world in which Fa is true but Fb false, even though Fb is not a logical consequence of Fa, namely: there is no *metaphysically* possible world in which Fa is true but Fb is false. A standard move here is to distinguish between metaphysically possible worlds and logically possible worlds, and to say that the definition of logical consequence invokes *logically* possible worlds. There might not be a metaphysically possible world in which Fa is true but Fb is not, but there is a logically possible one.

That is a response that threatens to lead quickly to circularity: Q: which are the logical truths? A: The ones true in all logically possible worlds. Q: Which are the logically possible worlds? A: They are the ones which do not violate any logical truths. A better approach is to move from talk of possible worlds to talk of models, where a model is something we can specify in set-theoretic terms. Tarski's model-theoretic definition of logical consequence is now quite standard, and it runs as follows:

Definition 7 *The sentence X follows logically from the sentences of the class K iff and only if every model of the class K is also a model of the sentence X. (Tarski, 1983: 417)*

A model of a sentence is a certain kind of set-theoretic construction, based on the formal language, with respect to which the sentence is true. Of course there may also be similar set-theoretic constructions with respect to which the sentence is false, but on this use of 'model' they are not models of the sentence. We need a neutral name for these more general kinds of set theoretic constructions, and I will call them 'structures'. A model of a sentence K (or a set of sentences X) is then a structure with respect to which K is true (or every member of X is true.)

The lead up above, along with the natural reference to "a model" or "a structure" can seem to suggest that a structure is a kind of object, perhaps a precise, neat, mathematical and independently apprehended correlate to a logically possible world. But *is* that what it is a correlate for?

A structure for a sentence A specifies a function from the non-logical expressions in a formal language to appropriate extensions for each expression in way that makes A true. For example, suppose we have a sentence from first-order logic like this:

$$Fa \rightarrow \neg Fb$$

This sentence contains two logical expressions, '$\rightarrow$' and '$\neg$'. The non-logical expressions are the predicate F and the individual constants a and b. A structure has to assign appropriate objects to the entire language, but let's just look at the part we're interested in. Here's one model of $Fa \to \neg Fb$:

Structure A:

$$D \text{ (the domain of the model) } = \{*, !, \#, ?\}$$

$$|F| = \{*, !, \#\}$$

$$|G| = \{\#, ?\}$$

$$|a| = !$$

$$|b| = ?$$

A makes $Fa \to \neg Gb$ true by making Gb true. And here's another:

Structure B:

$$D = \{!, \#, ?\}$$

$$|F| = \{!\}$$

$$|G| = \{\}$$

$$|a| = \#$$

$$|b| = ?$$

B makes $Fa \to \neg Gb$ true by making Fa false. And here's an assignment of objects to the non-logical expressions which does not make the sentence true, and hence isn't a model of it:

Structure C:

$$D = \{*, !, \#, ?, \&\}$$

$$|F| = \{*\}$$

$$|G| = \{\&\}$$

$$|a| = *$$

$$|b| = \&$$

It is picturesque to think of these objects via little Venn diagrams in which objects and sets are labelled with names and predicates:

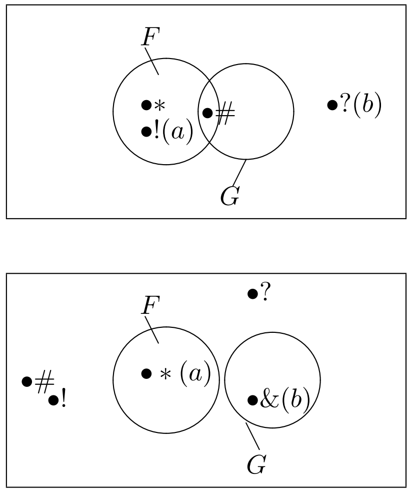

Figure 4: Diagrams for structures A and C: A is a model of $Fa \to \neg Gb$ and C is not.

A reasonable heuristic for using Venn diagrams to find a model for a sentence is this: push objects in and out of the circles in these diagrams, adding a few objects to the rectangle here, taking a few away there, and swapping labels on objects, until you have a diagram which represents a structure with respect to which the sentence is true. A sentence A is a logical consequence of a set Γ if and only if messing with the diagrams in this fashion can never produce a diagram representing a model that makes each member of Γ true, but A false.

All this talk of manipulating objects can encourage the idea that models are precise, mathematical replacements for logically possible worlds. As (Etchemendy, 1999: 23) points out, people sometimes think—even *Kaplan* sometimes thought (Kaplan, 1999: 159)—of Tarski as having reduced logically possible worlds to, or explicated them in terms of models, or more generally, structures. But is this how Tarski thought of them? Consider for a moment that the structures of model theory might not represent different possible worlds, but rather different possible languages. Not different ways the world might be, but different ways the formal language might be. The sentence Fa is true in some languages, but false in others. The sentence $Fa \lor \neg Fa$ is true in all languages, provided we keep the meanings of the logical constants fixed, and assign appropriate extensions to the non-logical expressions (e.g. sets to predicates, elements of the domain to names etc.) On this view, different structures represent different interpretations of the non-logical parts of the formal language. Hence another common gloss on "logical truth": *true on all interpretations* and logical consequence:

Definition 8 (Logical Consequence (interpretations)) *A sentence A is a logical consequence of a set of sentences Γ just in case every interpretation that makes every member of Γ true also makes A true. (See e.g. (Bostock, 1997: 7), (Quine, 1935: 81))*

This is the conception of model theory that Etchemendy refers to as "interpretive semantics" (Etchemendy, 1999) as opposed to the approach using possible worlds, which he dubs "representational semantics." On the former, checking for logical consequence is a matter of reasoning about different ways the language might be, on the latter it is a matter of reasoning about different ways the world might be.

One might wonder whether representational and interpretive approaches are two equally acceptable ways of thinking about the same formal machinery. Models are, after all, just set-theoretic constructions that are useful in characterising logical consequence. So long as that characterisation is accurate[6]—that is, it tells that that an argument is valid whenever it is valid and not otherwise—couldn't the correct way to think about models just come down to a matter of personal preference?

Perhaps not. There is a reason to prefer the interpretive approach to the representational one, even before we begin to consider context-sensitivity. Considering different structures involves considering what happens when names are assigned different extensions. Thus suppose that things really are as in structure 1 above, and the denotation of a is be ! and the denotation of b is be ?. Now consider an alternative structure 1^* (not above), in which a and b are both names for !. 1^* is a model of the sentence $a = b$ but 1 is not a model of that sentence. The models-represent-logically-possible-worlds view says we should think of 1^* as representing a logical possibility. But it doesn't: remember that when we consider alternative possible worlds, and describe them, we describe them in the language we actually speak. If a and b are not names for the same object, there is not even a *logically* possible world which satisfies $a = b$ interpreted as we use it. Considering what happens when a and b name the same object isn't imagining the world being different; it's imagining the language being different.[7]

Kaplan doesn't define logical consequence in "Demonstratives" but instead contents himself with the related property, logical truth. We can generalise from this to logical consequence. A sentence A is a logical truth in LD if it is true in all contexts of all structures. We will say that A is a logical consequence of

[6]Etchemendy famously doesn't believe that Tarski's definition is accurate, but see (Gomez-Torrente, 1998) for a response on Tarski's behalf.

[7]This already showed up the oddity of my way of describing changes that one would make to the venn diagrams above to get them to represent different possible worlds. I said that one was permitted to i) add and delete objects from the domain, ii) alter the extensions of predicates and iii) 'swap labels on objects'. The first two could be thought to correspond to considering how things would be if there were more or less objects, or if those objects had different properties. But in what sense does swapping labels on names correspond to a way the world might have been different? It is much better understood as a way the *language* might have been different.

Γ just in case there are no contexts in any structure in which every member of Γ is true but A is false. Now it is a celebrated feature of **LD** that there are logical truths ϕ such that $\Box\phi$ is not a logical truth, and not for some special technical reason, but because there are logical truths ϕ such that $\Box\phi$ is false—there are contingent logical truths. Examples include $Located(I, Here)$ (the translation of 'I am here' into the language of **LD**), $Exists(I)$ (I exist) and $A\phi \leftrightarrow \phi$ (Actually ϕ if and only if ϕ.)

If we take the $\Box$ to represent metaphysical necessity, as is standard (Burgess, 1997), it follows immediately that there are logical truths which express propositions that are not true in all metaphysically possible worlds and a fortiori that there are logical truths that express propositions that are not true in every logically possible world, since every metaphysically possible world is logically possible.

When we turn to logical consequence, we find that the following are valid arguments even though there are possible worlds where the premises are true and the conclusion false. Consider:

$$\frac{\text{AFa}}{\text{Fa}}$$

NFa is a logical consequence of Fa, both intuitively and in the sense that it fits the definition of logical consequence in **LD**; there is no context of utterance in which AFa is true, but Fa is not. However AFa receives the same truth-value with respect to all possible worlds, in particular, if it is true, then it is necessary, whereas Fa is contingent, hence if Fa is true relative to a context of utterance, then AFa is necessary but Fa is not—hence there is a possible world where the premise is true, but the conclusion is not.

This state of affairs will seem paradoxical if you think we simply stipulate that logical consequence is preservation of truth across possible worlds. But Tarski did not think of himself as stipulating the meaning of "logical consequence", but as proposing a precise explication of an intuitive notion that we possessed antecedently (Tarski, 1983: 409) and if we approach model theory as concerning interpretations, then **LD** simply represents a richer conception of an interpretation—if you like, a more general way of thinking about how extensions are assigned to expressions. Expressions may get their extensions relative to worlds and times, or they may get them relative to *contexts*, worlds and times. The result is both a generalisation of our definition of logical consequence thus far, and one that better tracks our intuitive, informal understanding of logical consequence when it comes to arguments containing context-sensitive expressions.[8]

[8] This is also something that should make one hesitate to equate content with "inferential role". In **LD**, 'A' and 'N' are inferentially equivalent in the following sense: ϕ always has $A\phi$ as a logical consequence and $A\phi$ always has ϕ as a logical consequence. Moreover ϕ always has $N\phi$ as a logical consequence and $N\phi$ always has ϕ as a logical consequence. But ϕ, $A\phi$ and $N\phi$ all mean different things, and to see this it is sufficient to note that the propositions they express have different modal and temporal profiles. $N\phi$ is true at all times, if true now. ϕ need not be like that. And $A\phi$ is necessary if true. ϕ need not be.

4 Context and the Linguistic Doctrine of Necessary Truth

There are difficult epistemological and methodological problems associated with logical truths. Many have remarked on their a priority and (supposed) necessity, and wondered what makes it the case that a sentence—whether of logic, or mathematics—is not merely true, but necessarily true, and how one could come to know such truths independently of experience. One traditional answer is the linguistic doctrine of necessary truth. According to this view, the truths of logic and mathematics are analytic—true in virtue of their meanings—and our knowledge of them is derived from the knowledge in virtue of which we are competent speakers of a language. Indexicality provides a new challenge to this kind of view. Consider the following sentence:

(1) I exist

(1) is usually uttered in contexts in which it expresses a contingent truth, but nothing in Kaplan's account prevents it from expressing a necessary one. Suppose for a moment that it were to be uttered by some necessarily existing object—perhaps God—then the proposition it expressed would be a necessary one. Hence there are sentences that express necessary propositions in some contexts, but contingent ones in others.

I like to use God in this example because I think it makes the point in a straightforward way, but perhaps you don't believe in God and object to his use in philosophical examples. There are other examples that will make the same point. Consider this one:

(2) That can be halved.

In a context where the object demonstrated is something which is only contingently halvable (say, a cake—there are possible worlds in which that same cake is utterly indivisible, but in our world cakes are usually halvable.) But in a context where the object demonstrated is the number 4, (perhaps our speaker points to the written numeral and it is clear from the context what he means to refer to) the claim made is necessary; 4 can always be divided into 2 and 2. Hence there are contexts with respect to which (2) expresses a contingent proposition, and contexts with respect to which it expresses a necessary one.

Or consider one more example:

(3) Fthat exists.

Fthat is a new indexical, one which directly refers to whatever object is the speaker's favourite. Some people's favourite objects are necessarily existing things, like numbers. Sally's favourite object is π, for example. Other people's favourite objects are, sadly, only contingently existing things; Dave's favourite object is his teddy bear. In contexts in which Sally is the speaker, (3) expresses a necessary truth, but in contexts in which Dave is the speaker, it expresses a contingent one.

I hope I have now said enough to make it plausible that Kaplan's framework allows there to be sentences which express necessary truths in some contexts, but contingent ones in others.

The existence of sentences like the ones above provides a fresh argument against the linguistic doctrine. The intuitive idea is straightforward: if a sentence expresses a necessary proposition in some contexts, but not in others, then its meaning alone cannot account for its necessity in the former contexts since that is constant between the two contexts.[9] Slightly more formally:

1. If the linguistic doctrine of necessary truth is correct, then if a sentence expresses a necessary truth, its meaning is sufficient to make it the case that it expresses a necessary truth.

2. There are sentences which express necessary propositions in some contexts, but contingent ones in others.

3. If a sentence expresses a contingent proposition in some contexts, then its meaning is not sufficient to make it the case that it expresses a necessary truth.

4. So there are sentences which express necessary truths whose meaning is not sufficient to make it the case that they express necessary truths.

5. So the lingustic doctrine of necessary truth is not correct.

5 An Indexical Barrier to Implication

The last consequence of **LD** that I wish to draw out concerns whether or not there are any valid arguments in which all the premises are non-context sensitive, but the conclusion is constant—that is, whether or not there is an indexical barrier to implication, on the model of (Restall and Russell, 2010). A first pass at such a barrier might look like this:

(4) No set containing only non-indexical sentences entails an indexical one.[10]

(4) is not true. Here are four sorts of counterexample.[11] First, we can get sentences containing 'actually' and 'now' from non-indexical sentences:

[9]This argument appears at greater length, and with responses to anticipated objections in (Russell, 2010b).

[10]Normally we speak of indexical expressions—like *I*, and *actually*—not indexical sentences, but an indexical expression is one whose content can vary with context and the content of a sentence is determined by the content of its parts, and hence a sentence containing an indexical expression will be an indexical expression itself.

[11]It might seem like belabouring the point to produce *four* different sorts of counterexample when one will do the job, but seeing the problems with this naive barrier will make it easier to see why the final restricted barrier theorem has the shape it does.

$$\frac{\phi}{A\phi} \qquad \frac{\phi}{N\phi}$$

$$\frac{\text{Sam is a frog.}}{\text{Actually Sam is a frog.}} \qquad \frac{\text{Sam is a frog.}}{\text{Sam is a frog now.}}$$

Second, there are counterexamples in which the conclusions are logical truths containing indexicals, or in which the premises are logically inconsistent non-indexical sentences:

$$\frac{Fa}{Located(I, Here)} \qquad \frac{Fa}{\alpha = dthat[\alpha]}$$

$$\frac{\text{Sam is a frog.}}{\text{I am here now.}} \qquad \frac{\text{Sam is a frog.}}{\text{The actual shortest spy is the shortest spy.}}$$

Third, there are counterexamples that make use of the fact that if a property holds quite generally and universally, then it holds of the referent of an indexical, no matter which context we are in:

$$\frac{\forall x Fx}{FI} \qquad \frac{\forall p Rp}{RHere}$$

$$\frac{\text{Everything is a frog.}}{\text{I am a frog}} \qquad \frac{\text{It's raining everywhere.}}{\text{It's raining here.}}$$

Finally, there are counterexamples that we might call Prior-style counterexamples, after A.N.Prior's well-known objection to Hume's Law (Prior, 1960):

$$\frac{Fa}{Fa \vee GI} \qquad \frac{\begin{array}{l} Fa \vee GI \\ \neg Fa \end{array}}{GI}$$

$$\frac{\text{Sam is a frog.}}{\text{Sam is a frog or I am a newt.}} \qquad \frac{\begin{array}{l} \text{Sam is a frog or I am a newt.} \\ \text{Sam is not a frog.} \end{array}}{\text{I am a newt.}}$$

The idea with the Prior-style arguments is that the disjunctive sentence is of unclear status. Should we treat it as genuinely indexical, or genuinely non-indexical? Perhaps we could do either, but Prior's point is that if we do the former, then the arguments on the left hand side of the page are counterexamples to the barrier thesis, but if we do the latter, then the arguments on the right hand side are.

So the naive thesis is false, and one might wonder why anyone would pursue the project further. But there have been a number of thought experiments in the philosophy of language that have suggested their might be some kind of barrier to implication between indexical and non-indexical sentences. (Castaneda, 1968; Lewis, 1979; Perry, 1988) For example, John Perry tells the story of a shopper who discovers a trail of sugar on the floor whilst pushing his cart around the grocery store. The shopper suspects that someone has put a bag of sugar with a hole in it into his or her cart, whcih they are now pushing around the store, unaware that they are leaving a trail of sugar behind them. Our shopper decides to follow the trail, catch up with the Messy Shopper, and enlighten them, but after following the trail in circles, unable to catch up as the mess gradually gets worse, our shopper suddenly realises that he has been making the mess all along. He realises: "*I'm* the Messy Shopper!" At this point, his behaviour changes. He looks in his *own* cart to find the bag. Perry holds that until the Messy Shopper accepts the indexical description of the situation, using *I*, we wouldn't expect this behavioural change, and a natural hypothesis about why is that it doesn't follow from the non-indexical sentences he accepts before that that "I am the messy shopper!" is true for him.

That is a bit rough and ready. Surely the Messy Shopper knows the truth of some indexical sentences, such as 'I am in a grocery store' and 'I am following the Messy Shopper.' Lewis' Two Gods thought experiment (Lewis, 1979; Russell, 2010a) fits better with the demand for a barrier thesis, but the success or otherwise of thought experiments such as these in eliciting intuitions about barrier theses will be less relevant once we have proved one, and it is to this task that I now turn.

In **LD** sentences are true relative to a context. Some aspects of the context also serve as aspects of the circumstance of evaluation, namely the world and the time of the context. We are going to be particularly interested in indexicals which are sensitive to changes in aspects of the context which are *only* aspects of the context, and not also aspects of the circumstances of evaluation. We define the inter-contextual relation of *partial context shift* as follows:

Definition 9 (Partial Context Shift)

A context $c^* = \langle a^*, p^*, t^*, w^* \rangle$ *in a structure* M^* *stands in the partial context shift relation to a context* $c = \langle a, p, t, w \rangle$ *in a structure* M *iff* $M^* = M$ *and* $t^* = t$ *and* $w^* = w$ *(i.e. the structures remain identical and contexts are allowed to shift only in their agent and place elements.)*

Here is an illustration of partial context-shift. Suppose we have a context C, within some structure M, in which John is the agent, the place is St Louis, the time is 3pm, and the world is the actual world. Now consider three variations on this context. C1 is different in that someone else—Sally—is the agent. This represents a change in part of the context which is not also part of the circumstances of evaluation, and so C1 stands in the partial context-shift relation to the original context. C2 is different in that the time is 4pm (though John is still the agent.) C2 does not stand in the partial context-shift relation to the original context because it represents a change in a part of the context which *is* also a part of the circumstances of evaluation. Finally, C3 is different in that the agent is now Sally, but also in that the structure that C3 is a part of is M3, not M, and in M3 there are additional possible worlds. C3 does not stand in the partial context-shift relation to C, this time because their containing structures are different. The intuitive idea here is that partial context-shift can be used to isolate a certain sub-class of indexical sentences: those whose truth-values change with aspects of the context of utterance which are not also part of the circumstances of evaluation. Call such indexicals *type 1* indexicals. [12]

Now we'll use the intercontextual relation defined above to define two classes of sentences:

Definition 10 (Type 1 Constant Sentences) *A sentence A is constant iff whenever* $(M, c) \vDash A$*, and* (M^*, c^*) *is a partial context shift of* $(M, c), (M^*, c^*) \vDash A$.

Definition 11 (Type 1 Indexical Sentences) *A sentence A is indexical iff there is some structure-context pair* (M, c) *and some partial context shift of* $(M, c), (M^*, c^*)$ *such that* $(M, c) \vDash A$ *but* $(M^*, c^*) \nvDash A$.

The idea here is that a type 1 constant sentence is one such that changing the part of the context which is not also a part of the circumstances of evaluation will never affect the truth-value. Examples of type 1 constant sentences include Fa, AFa, NFa, $Located(I, Here)$, $Fa \wedge \neg Fa$ and $FI \vee \neg FI$. Examples of type 1 indexical sentences include FI, $RHere$, $Fa \rightarrow FI$, $\forall x GxI$ and $Ga \vee FI$. Some of the type 1 constant sentences might naturally be referred to as 'indexical sentences' on less technical readings of the expressions 'constant sentence' and 'indexical sentence', but what I am trying to do here is to isolate the distinctive feature of indexical sentences that underlies the intuition that they are not entailed by constant ones. I'm suggesting that that distinctive feature is the ability to change their truth-value when we change the context (without changing the circumstances of evaluation.) It may sometimes turn out that sentences with indexical expressions in them nevertheless do not have that distinctive feature, e.g. 'I am here now' which is true in all contexts. Nonetheless the more strictly formulated barrier thesis can help to explain our somewhat inchoate intuitions that one ought not to be able to deduce indexical sentences from non-indexical ones *and* the exceptions to that rough and ready rule.

We require one more definition before formulating our restricted indexical barrier thesis. Let $A(v/\alpha)$ be the result of replacing all occurrences of the indexical α in A with the variable v.

Definition 12 (Complete Indexical Generalisation) *An indexical generalisation of an sentence A with respect to an indexical term* α *is a sentence* $\forall \xi (A(\xi/\alpha))$ *where* ξ *does not already occur in A. For example* $\forall v(Fv \wedge GHere)$ *is an indexical generalisation of* $FI \wedge GHere$ *with respect to 'I'. A complete indexical generalisation of A is the result of repeating this process until there are*

[12]If you think about it, nearly every sentence can change its truth-value with aspects of the context which are aspects of the circumstances of evaluation, since nearly every sentence changes its truth-value when you change the possible world aspect of the context.

no more indexicals in the sentence, e.g. $\forall p \forall v(Fv \wedge Gp)$ *is a complete indexical generalisation of* $FI \wedge GHere$.[13]

We can now formulate and prove our indexical barrier theorem:[14]

Theorem 13 (Restricted Indexical Barrier Theorem) *No consistent set of (type 1) constant sentences* X *entails a (type 1) indexical sentence* A *unless* X *also entails all of* A*'s complete indexical generalisations.*

Proof 14 *Suppose* $X \vDash A$ *and let* A' *be a complete indexical generalisation of* A*. We show that* $X \vDash A'$*. Let us suppose we number the indexicals in the sentence in turn from left to right:* $\alpha_1 ... \alpha_n$ *Note that* A' *will be (or will be equivalent to) the last in a finite sequence of formulas* $A, A_1 ... A_n$ *such that* A_j *is* $\forall \xi_j A_{(j-1)}(\xi_j/\alpha_j)$

Induction Hypothesis*: for all* A_m *where* $m < j$*,* $X \vDash A_m$*.*

Induction Step*: We show that* $X \vDash A_j$*. Let* $\langle M, c\rangle$ *be an arbitrary structure-context pair. Suppose* $\langle M, c\rangle$ *makes every member of* X *true. Each member of* X *is constant, and hence for all* $c' \in C$ *(where* C *is the set of contexts in the structure* M*),* $\langle M, c'\rangle$ *will make every member of* X *true as well. Since by the induction hypothesis* $X \vDash A_i$*, it follows that for all* $c' \in C$*,* A_i *will be true at* $\langle M, c'\rangle$*.*

Now suppose there were some assignment, f*, of objects to variables with respect to which* $A_m(\xi_j/\alpha_j)$ *is false for some* $\langle M, c^*\rangle$*. Then with respect to a context which has* $f(\xi_j)$ *as its first member (2nd member, if* α_j *is a p-term instead of an i-term)* A_i *would be false at* $\langle M, c^*\rangle$*. But this contradicts what we have already found. Hence there is no assignment which makes the open formula false. It follows that* $\forall \xi_j A_i(\xi_j/\alpha_j)$*—that is* A_j*—is true with respect to* $\langle M, c\rangle$*. Hence* $X \vDash A_j$*.*

It follows by complete induction that if $X \vDash A$*, then* $X \vDash A'$*.* □

Philosophers of language use a number of familiar tools for adjudicating disputes about the meanings of particular expressions. One of these tools exploits the fact that the validity of an argument depends on what the sentences contained in the argument mean. Hence one way to test a theory of meaning involves examining the valid and invalid arguments in which the expression appears. If the argument is valid, the theory of meaning for the expressions containing in it had better not predict that it is not, and if the argument is invalid, the theory of meaning for the expressions containing it had better not predict that it is.

[13] Given the formation rules for **LD**, specifying that $\forall \xi(A(\xi/\alpha))$ has to be a *sentence* ensures that ξ is of the correct term-type (i.e. position or individual) for that argument place in the predicate.

[14] This proof first appeared in (Russell, 2010a).

These days, there are many disputes about whether particular expressions in natural languages are indexical. For example, some have contended that truth-ascriptions, such as *It is true that snow is white* or vague expressions like *blue* and *red*, or knowledge ascriptions, such as *John knows that 2+2=4* or names, like *John* and *Mary* are indexical. Others disagree. I am hopeful that the restricted indexical barrier theorems could provide an additional tool which could help us to make some headway within these debates.

However, seeing how to connect up the formal barriers with an argument for or against a particular semantic view is not as easy as one might like. In the abstract, an ideal use of the barrier theorem could work as follows: theory A holds that expression ϕ is an indexical. We find a valid argument in the natural language which has a set of purely non-indexical premises, and a sentence containing ϕ as its conclusion. We point out that, given the barrier theorem, such an argument could not be valid if ϕ were indexical. We conclude that ϕ is not indexical after all.

This all sounds very reasonable, but there are some pitfalls to be avoided in the application and I will finish up this section by pointing to some of these.

Case Study 1: vague predicates It is commonly accepted that certain vague predicates are indexical. For example, the received view of gradable adjectives, like 'tall' and 'rich' is that they pick out different properties, given different contexts of utterance (for example, when we're discussing 12 year olds, or when we're discussing professional basketball players.) (Kennedy, 2011) So we might expect the barrier theorem to confirm that sentences like *John is tall* are never entailed by non-context sensitive sentences, such as *John is 5' 5"*, *John is a 12 year old* etc. Yet **LD** does not actually contain any indexical predicates at all. Its indexicals are limited to the two terms *I* and *Here*, the 1-place functor *dthat* and a few operators such as *actually* and *tomorrow*. Of these only the terms are of the kind alluded to by the barrier theorem, i.e. they are sensitive to aspects of the context which are not also part of the circumstance of evaluation. **LD** doesn't actually model the relation of logical consequence on sentences containing indexical predicates. So in order to use the theorem to confirm the received view of gradable adjectives, we would need to first, add some indexical predicates to the formal language and second, add something to the context-sequence that could plausibly be the thing that the extension of these predicates varied with. (Presumably it is neither the agent, nor the location which determines the denotation of *tall*.) Although see no reason why this couldn't be done, it does need to be done and it adds to the work required to apply the barrier theorem in practice.

Case study 2: time and tense Here is a second cautionary tale. In **LD**, the circumstances of evaluation are represented by a world-time pair $\langle t, w\rangle$. Many philosophers believe that propositions do not change their truth-values with respect to time, but only with respect to different possible circumstances, e.g. (King, 2003). Other, more liberal, philosophers believe that propositions change their truth-values with respect to epistemic standards, or even standards of taste. (MacFarlane, 2005; 2009) In principle, someone might think that propositions change their truth-values with respect to location. It would be straightforward to adapt **LD** to account for the first and third of these views: for the first, we let a circumstance of evaluation be a lonely possible world w and a context of utterance would remain a quadruple $\langle a, p, t, w\rangle$. Since any tense operators would now be vacuous we would take the tense operators F, G, Now etc out of the language, and instead introduce first order variables which range over times and, optionally, an indexical singular term Now, on the model of $Here$, which denotes the time of the context. Call the resulting system (the details of which are left as an exercise for the reader) $\mathbf{LD}^{-}$.

Alternatively, if we thought that propositions could change their truth-values with locations, as well as time and possible world, we could instead add to the circumstances of evaluation, making it a triple of location, time and possible world. Again contexts of utterance could remain quadruples $\langle a, p, t, w\rangle$. Since propositions can change their truth-values with location, it would be natural to add operators like 'at all places ϕ' and 'at some places ϕ' and replace our singular term $Here$ with an indexical operator $Here\phi$. Call this system $\mathbf{LD}^{+}$

In $\mathbf{LD}^{+}$, location, time and possible world are aspects of the context of utterance which are also aspects of circumstances of evaluation. In $\mathbf{LD}^{-}$, neither time nor location are aspects of the circumstances of evaluation. So the barrier theorem, (suitably adapted to the new systems) will tell us that if $\mathbf{LD}^{+}$ is the correct logic for English, certain arguments containing Now and $Here$ may be valid, whereas if $\mathbf{LD}^{-}$ is the correct logic for English, certain arguments of that kind cannot be valid. Given that we have independent intuitions about which English arguments are valid, this should allow us to decide between $\mathbf{LD}^{+}$, **LD** and $\mathbf{LD}^{-}$ and thereby tell us whether or not the propositions expressed by English sentences express propositions whose truth-values vary over times and locations. So consider the following English arguments:

It is snowing.
____
It is snowing here.

It is snowing.
____
Actually, it is snowing.

It is snowing.
____
Now, it is snowing.

I hold that the first argument is not valid but that the second two are. If an argument is needed, consider what we would say about the following conversations: Mike's mum is in her study and can't see out the window from where she is. She calls out to the teenage Mike, who is in the living room and can see outside: "Is is snowing?" Mike says "duh, of course it's snowing." Mike's mum goes through and looks outside, but sees no snow. She says, "I thought you said it was snowing?" Consider what we would say about the following three responses from Mike:

(5) Well, it *is* snowing. It's snowing in Chicago. It just isn't snowing *here.*

(6) Well, it is snowing. It snowed last winter. It just isn't snowing *now*.

(7) Well it is snowing. It could have been snowing. It just isn't snowing *actually*.

The first makes sense, even if it is a little pedantic. That's because you can truthfully say 'it is snowing' so long as it is snowing somewhere. Pragmatic considerations—plausibly violated by the uncooperative Mike—suggest that Mike's response is unhelpful, but not that he has actually said anything false. Hence the premise in argument 1 doesn't entail the conclusion. But responses 2 and 3 make no sense at all. I think that's because 'actually it is snowing' and 'it is snowing now' follow from 'it is snowing' (that is, arguments 2 and 3 are valid) making responses 2 and 3 incoherent.

Putting this data together with the indexical barrier theorem suggests the following: if $\mathbf{LD}^-$ were the correct logic for English, then time, like location, would be an aspect of context of utterance which was not also an aspect of circumstances of evaluation, and hence we would expect arguments with non-indexical premises but conclusions that contained the indexicals *Now* and *Here* to be invalid (modulo certain special circumstances where the properties attributed were universal.) The argument above using Now is not invalid. Hence $\mathbf{LD}^-$ is not the correct logic for English, and moreover the propositions expressed by English sentences can vary their truth-values over time.

So far so neat, but here comes the cautionary part of the tale: if one thinks, *pace* the above argument, that propositions do not change their truth-values with respect to times, but only with respect to possible worlds, then one already needs a view about how a sentence like 'it is snowing' can be true on some occasions, but false on others. The natural view is that that happens because the sentence expresses different propositions, ones with different truth-values, on different occasions. For example perhaps it expresses the proposition that it is snowing at time A, when uttered at A, and the proposition that it is snowing at time B, when uttered at time B. If you hold this view, then you think that 'it is snowing' is indexical—it expresses different content on different occasions, and hence the validity of the argument 3 above is no violation of the indexical barrier theorem after all. $\mathbf{LD}^-$ is saved.

Well, perhaps. One thing that the indexical barrier theorem can make clearer are the constraints on adopting a position such as this consistently. The barrier thesis says that if time is not an aspect of circumstances of evaluation, and $\phi \vdash N\phi$ is valid for all values of ϕ, then all values of ϕ are themselves indexical, that is, the $\mathbf{LD}^-$ fan will be obliged to hold that all sentences are indexical. There might be some independent reason to hold this view anyway: perhaps all sentences contain verbs, all verbs are verbs are tensed, and tense is a form of indexicality. But again, sometimes linking views together more clearly with their commitments is itself a useful tool for making progress.

There is a lot more to say about both of the above examples, but my hope for now is just that they illustrate the following point: that while I hope the indexical barrier theorem can be a useful tool in the philosophy of language, it requires a certain amount of skill and work to wield it. But I hope that this won't put us off too quickly.

6 Conclusion

In this paper I have argued that considering Kaplan's model of indexicality and his logic *LD* in the context of logic gives us two rather surprising insights into the nature of logic, and can help us develop logical tools for use in the philosophy of language. The consequences for logic were that logical consequence is not necessary truth preservation and a new, very intuitive, challenge to the linguistic doctrine of necessary truth. The tool for for philosophers of language was the restricted indexical barrier theorem—a tool that goes some way towards explaining intuitions about the essentiality of the indexical, and could perhaps also have application in semantics.

References

Bostock, D. (1997). *Intermediate Logic*. Oxford University Press, Oxford.

Burgess, J. P. (1997). Quinus ab omni nævo vindicatus. *Canadian Journal of Philosophy, Supplementary Volume 23: Meaning and Reference*.

Castaneda, H.-N. (1968). On the logic of attributions of self-knowledge to others. *Journal of Philosophy*, 65(15).

Etchemendy, J. (1999). *On the Concept of Logical Consequence*. CSLI, Stanford.

Gomez-Torrente, M. (1998). On a fallacy attributed to tarski. *History and Philosophy of Logic*, 9.

Kaplan, D. (1989a). Afterthoughts. In Almog, J., Perry, J., and Wettstein, H., editors, *Themes from Kaplan*. Oxford University Press, New York.

Kaplan, D. (1989b). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives. In Almog, J., Perry, J., and Wettstein, H., editors, *Themes from Kaplan*. Oxford University Press, New York.

Kaplan, D. (1999). What is russell's theory of descriptions? In Irvine, A. D., editor, *Bertrand Russell: Language, Knowledge and the World*. Routledge, London.

Kennedy, C. (2011). Adjectives. In Fara, D. G. and Russell, G., editors, *The Routledge Companion to the Philosophy of Language*. Routledge, New York.

King, J. C. (2001). *Complex Demonstratives*. MIT Press, Cambridge, Mass.

King, J. C. (2003). Tense, modality, and semantic values. *Philosophical Perspectives*, 17(1):195–246.

Lewis, D. (1979). Attitudes de dicto and de se. *The Philosophical Review*, 88:513–543.

MacFarlane, J. (2005). The assessment sensitivity of knowledge attributions. In *Oxford Studies in Epistemology 1*, pages 197–233. Oxford University Press, Oxford.

MacFarlane, J. (2009). Non-indexical contextualism. *Synthese*, 166:231–250.

Perry, J. (1988). The problem of the essential indexical. In Salmon, N. and Soames, S., editors, *Propositions and Attitudes*, Oxford Readings in Philosophy. Oxford University Press.

Prior, A. N. (1960). The autonomy of ethics. *The Australasian Journal of Philosophy*, 38:199–206.

Quine, W. V. O. (1935). Truth by convention. In *The Ways of Paradox and other essays*, chapter 9, pages 70–99. Random House, New York.

Restall, G. and Russell, G. (2010). Barriers to Implication. In Pigden, C., editor, *Hume, Is and Ought: New Essays*. Palgrave MacMillan.

Russell, G. (2010a). Indexicals, context-sensitivity and the failure of implication. *Synthese*, pages 1–18. 10.1007/s11229-010-9756-9.

Russell, G. (2010b). A new problem for the linguistic doctrine of necessary truth. In *New Waves in Truth*, pages 267–281. Palgrave MacMillan.

Tarski, A. (1983). On the concept of logical consequence. In Corcoran, J., editor, *Logic, Semantics and metamathematics*, pages 409–420. Hackett, Indianapolis, 2nd edition.