Where Monsters Dwell

David Israel and John Perry

August 26, 1994

1 Introduction

> Are there such operators as 'In some contexts, it is true that', which when prefixed to a sentence yields a truth if and only if in some context the contained *sentence* (not the content expressed by it) expresses a content that is true in the circumstances of that context?
>
> Operators like 'In some contexts it is true that', which attempt to meddle with character, I call *monsters*. I claim that none can be expressed in English...
>
> I am not saying we could not construct a language with such operators, just that English is not one. And such operators *could not be added to it.* (Kaplan 1989, pp. 510f.)

Kaplan says that monsters violate **Principle 2** of his theory. **Principle 2** is that indexicals, pure and demonstrative alike, are directly referential. In providing this explanation of there being no monsters, Kaplan feels his theory has an advantage over double-indexing theories like Kamp's or Segerberg's (or Stalnaker's), which either embrace monsters or avoid them only by ad hoc stipulation, in the sharp conceptual distinction it draws between circumstances of evaluation and contexts of utterance. We shall argue that Kaplan's prohibition is also essentially stipulative, and that it is too general. The main difference between ourselves and Kaplan is that the basic carriers of a truth-value is a sentence-in-a-context; our account is utterance-based.

Our utterance-based theory, which we call the *reflexive-referential theory* differs from Kaplan's in a couple of important respects, which, we claim, are crucial for a correct understanding of issues related to monsters. Here is a summary of the similarities and differences:

- On both approaches, monsters are formally possible; that is in Kaplan's formal theory, coherent definitions can be written for monsters, and this is also true on our theory.
- Conceptually, in Kaplan's "direct reference" semantics, the prohibition against monsters has what one might call a deep semantical motivation, for the basic semantical unit, the *content*, simply does not include the parameters on which monsters would operate.

- In contrast, our basic semantic unit, which we call *reflexive content*—one example of which is *indexical content*—is one among many levels of content that we recognize. Reflexive content does include the parameters on which monsters operate, and there is no deep semantic motive for excluding them on our theory.
- On our account, there are many places where monsters might, but don't dwell, and the reasons for their absence are basically pragmatic.
  (i) We are usually interested in what we call *incremental content*, and at this level of content the parameters monsters need to thrive are unavailable
  (ii) Reflecting these interests, some important operators, like "says that", operate only on incremental content.
- Still, we think monsters, in particular *epistemic* or *cognitive*, as opposed to *alethic*, are possible and (who knows?) even actual.

The structure of the rest of the paper is as follows. In §2, we sketch the background of Index and Double-Index theories of modality. In §3 we give a brief account of Kaplan's framework and of his theory of indexicals. In §4, we return, with fresh motivation, to Double-Index accounts. In §5 we end with a look ahead to our own account.

2 Modality And Monsters

2.1 Index Theory

A central notion of semantics in the style of Tarski is that of the truth (satisfaction) of a sentence (formula) *in a model*. This relativity of truth can seem quite artifactual, since it may seem we are interested in truth *simpliciter*. Of course, logicians aren't interested in truth simpliciter, but in logical truth. In the model-theoretic tradition, this latter notion is captured by that of truth in all models. When we move to modal logic, the extension of the classical relativity to models, in terms of truth in a world in a model, seems actually less artifactual than the base case, for we are after logical aspects of a notion that involves relativizing truth to a space of alternative possibilities. Something similar holds for temporal logic: here it is not alternative possibilities, but simply different moments in time that we conceive truth as relative to. And so on for other modal modelings. Moreover, we may have occasion to model relativities along more than one dimension.

Consider the following progression:

- $\mathcal{M}, w \models \Box\Phi$
  - iff $\forall w' : wRw' \rightarrow w' \models \Phi$
- $\mathcal{M}, t \models H\Phi$[1]
  - iff $\forall t' : t' < t \rightarrow t' \models \Phi$

[1] 'H' is the Priorean 'throughout history' or 'always in the past' operator.

- $\mathcal{M}, \langle w, t\rangle \models \Box(H\Phi)$
  - iff $\forall w', t' : wRw' \,\&\, t' < t \rightarrow \langle w', t'\rangle \models \Phi$

Why stop with two? Indeed, the advice of a great logician tells us there is no good reason to stop at all.

> For more general situations one must not think of the $i \in I$ as anything as simple as instants of time or possible worlds. In general we have
>
> $$i = \langle w, t, p, a \ldots\rangle$$
>
> where the index i has many *coordinates*...All these coordinates can be varied, **perhaps independently**, and thus affect the truth-value of statements which have indirect reference to these coordinates. (Scott 1970)

The spirit of this advice is to give a unified treatment of modality and indexicality. In the beginning sentences are evaluated with respect to or in a model; the meanings (intensions) of sentences can be thought of as properties or sets of models: the meaning of a sentence being the set of models that make it true. When we move to the alethic modalities, sentences are evaluated relative to a model and a world of that model. Fixing a model, then, we identify the meaning of a sentence with the set of worlds of that model at which the sentence is true.

Imagine one starts, where our little progression of relativities ends, with a language with both alethic and temporal modalities, but without indexicals. The meaning of a sentence Φ, $[\![\Phi]\!]$, is a set of ordered pairs of worlds and times (or a function from such pairs to $\mathcal{B} = \mathbf{2}$). If we consider adding indexicals such as the personal pronoun 'I' and the locative adverb 'here' to the language, Scott advises us simply to extend the structure of the indices or *points of reference*. The meanings of sentences of the new language are now sets of quadruples of worlds, times, individuals (people) and places. These quadruples are the *circumstances* within (against) which a sentence must be evaluated for truth and falsity.

Kaplan presents a dilemma for this *pure index theory* approach. Consider the two sentences:

(1) a. I am here now.
    b. John Perry is in Moraga on June 15, 1994.

If we consider the quadruple i that consists of the actual world, June 15, 1994, JRP and Moraga and consider $[\![\text{I am here now}]\!](i)$, we can see that the proposition expressed by (1a) at that index is the same as the proposition expressed by (1b) at that index and at many others. Indeed, we can imagine or stipulate that the only relativity in (1b) is pure alethic relativity; its truth depends only on what world is being considered; it is true in some, false in others. What of (1a)? Consider the index i' that consists of the actual world, June 15, 1994, Napoleon Bonaparte and Moraga. Surely

$[\![\text{I am here now.}]\!](i') = 0$. Thus (1a), too, is contingent—true at some indices, false at others. Nothing (so far) can be said against this index, for Scott explicitly allows the possibility of of independent variation of the coordinates. So unless something more is said, we seem to have lost any chance of a logic of indexicals; for, of course, (1a) differs from (1b) precisely in being a very plausible candidate for a valid sentence in such a logic.

This problem of missed validities motivates the move to restriction to *proper* indices. In our example, proper indices are those $\langle w, t, a, p\rangle$ such that in w, a is located at p at t. Now if we restrict our structures to those *proper structures* in which all indices are proper, (1a) comes out logically true: true at every index in every proper structure. Consider (2)

(2) Necessarily, I am here now.

$\Box\Phi$ is true at an index in a structure if Φ is true at every index in that structure. Now if we assume the principle of modal generalization that if $\models \Phi$, then $\models \Box\Phi$, we seem to be stuck with the logical truth of (2), and this seems wrong. Here we haven't missed a validity; we've created spurious ones.

There is a way around this dilemma. It involves dropping the standard principle of modal generalization, and Montague (arguably) avails himself of it in Montague 1974. There he allows structures with improper indices, but defines logical truth as truth at every proper index in every structure—thereby guaranteeing the logical truth of (1a). As for (2), it is not logically true, because there is a structure with improper indices, that is, a structure such that (1a), though logically true, is not (just plain) true at every index in that structure.

This solution seems to be merely a technical trick; as we shall see, though, it is an attempt to get at something real and important.

2.2 Double-Index Theory

What is needed, but not provided by index theory, is an explanation of the special role of proper indices in the characterization of the logical truths of a language with indexicals. This special role is determined by the different roles played by aspects of context and aspects of circumstance in determining the truth of sentences. The aspects of context—the identity of the speaker, the place and time of the utterance—determine aspects of the proposition expressed by the sentence in the context; the proposition so determined is then to be evaluated for truth and falsity in varying circumstances. We thus get a two-step account:

- from sentences and contexts to propositions,
- from propositions in circumstances to truth-values.

Index theory yields a one-step account: from sentences at indices to truth-values. Kaplan's diagnosis of the problems with index theory is thus exactly right:

> The difficulty is the attempt to assimilate the role of *context* to that of *circumstance*. The indices $\langle w, t, a, p\rangle$ that represent contexts must be proper in order that (1a) be a truth of the logic of indexicals, but the indices that represent circumstances must include improper ones in that (2) *not* be a logical truth. If one wishes to stay with this sort of index theory, the minimal requirement is a system of *double* indexing, one index for context and another for circumstances. (Kaplan 1989, pp. 509f.)

Double-Index theory was developed by Vlach and Kamp (Kamp 1971), both students of Montague and systematized, though not with complete generality, by Segerberg (Segerberg 1973). We present here only a hint by way of examples. The basic idea is to model the distinction between contexts and circumstances via a two-dimensional logic of (one family) of index sets. In the simplest case, the index set is a set of points, as in abstract versions of modal logic. We introduce the following notation:

- $w \models_v \Phi$

which, for the applications at hand, is to be read "the sentence Φ, uttered in world w is true at world v. Consider now the following contrasting pairs:

(3) a. Necessarily, Φ
   - $w \models_v \Box\Phi$ iff $\forall v' : vRv' \rightarrow w \models_{v'} \Phi$

   b. Actually, Φ
   - $w \models_v A\Phi$ iff $w \models_w \Phi$

(4) a. On the next day, Φ
   - $t \models_u O\Phi$ iff $t \models_{u+1} \Phi$

   b. Tomorrow, Φ
   - $t \models_u T\Phi$ iff $t \models_{t+1} \Phi$

How does double-indexing deal with (1) and (2)? Rather than look at that rather complicated case, it suffices to look to see how double-indexing allows one to guarantee the validity of $\Phi \leftrightarrow A\Phi$ without yielding that of $\Box(\Phi \leftrightarrow A\Phi)$. The idea is precisely that deployed in the single index account by Montague. Let indices now be pairs of worlds and call the diagonal pairs $\langle w, w\rangle$ proper. Now define indexical validity or indexical logical truth as truth at all proper indices in all structures. Validity or logical truth simpliciter is truth at all indices. Thus the truth-clause for A guarantees the indexical validity of $\Phi \leftrightarrow A\Phi$, but necessity requires truth at all indices, proper or not.[2] Thus for the necessitation to be indexically valid, the original biconditional must be true at all indices, proper or not; but we can have $w \models_{w'} A\Phi$ without having $w \models_{w'} \Phi$.

Kaplan does not take note of the device Montague exploited to solve the problem for index theory posed by (1) and (2); nor does he mention

[2] We assume, for simplicity, that necessity is a universal, S-5 operator.

the similar device available to double index theory. His complaint against the latter is that it, too, blurs the distinction between contexts and circumstances. This seems false, but it does lead Kaplan into his discussion of monsters.

> However, mere double indexing, without a clear conceptual understanding of what each index stands for [of the conceptual difference between context and circumstance?] is still not enough to avoid pitfalls. (Kaplan 1989, p. 510)

The pitfall is the begetting of monsters. What are monsters? Kaplan's example is:

- In some context it is true that Φ

This is a monster if it understood as yielding a truth upon being prefixed to a sentence just in case in some context, *not* in some circumstance, the embedded sentence expresses a true proposition in the circumstances associated with that context. Kaplan notes that there is a construction in English that allows us to say what we seem to want to say with this monster:

- In some context, "Φ" is true.

or more fully, and in the style of double-indexing:

- In some context, c, the proposition expressed by "Φ" in c is true in c.

This semantic ascent brings out a parallel between monsters and a famous brand of puzzle:

- If you call a horse's tail a leg, how many legs does a horse have?

On a monstrous reading, the answer is: 1. This is the reading—assumed here to exist—on which the puzzle is paraphrasable as:

- In a context in which the word 'leg' means what 'tail' actually means, how many legs does a horse have?

Now we can see the aptness of Kaplan's comment that monsters attempt to operate on the meanings of sentences, as opposed to the content (proposition) expressed by a sentence in a context. Of course there is a such a difference only in an account in which the meaning of a sentence differs from the proposition expressed by it. There is no such difference on pure index theory. The meaning of sentence (in a structure), $[\![\Phi]\!]$ is that function from indices of that structure to $\{0,1\} = \mathcal{B}$ whose value for $i \in I$ is 1 just in case Φ is true at i, and this is what is usually taken to be the proposition expressed by Φ.

In Double Index Theory, the meanings of sentences are functions from indices (in their role as contexts) to propositions, which are themselves

functions from indices (in their role of circumstances of evaluation) to truth-values:

- Meanings of sentences as functions:
- $\mathcal{I} \rightarrow (\mathcal{I} \rightarrow \mathcal{B})$

A two-dimensional operator corresponds to a function $\mathcal{F}$: $(I \rightarrow (I \rightarrow B)) \rightarrow (I \rightarrow (I \rightarrow B))$. **All such operators are monsters!** Of course, some of these monsters are benign, in that they don't really change meanings: they operate only on proposition expressed. Some, however, are not so innocent.

To make this contrast more intelligible, we first introduce some notational conventions.

- $[\![\Phi]\!](i)(i') = 1$ iff $i \models_{i'} \Phi$.

We want to be able to get at the proposition expressed by a sentence at an index. To do this, we superscript the context-index:

- $[\![\Phi]\!]^i(i') = [\![\Phi]\!](i)(i')$

Now let $[\![\varnothing\Phi]\!]^i$ be the proposition such that $[\![\varnothing\Phi]\!]^i(i') = [\![\Phi]\!](i)(i)$. Thus, for any index (= context) i, $[\![\varnothing\Phi]\!]^i$ represents the diagonal proposition for i, so $[\![\varnothing\Phi]\!]$ is the 'diagonalizing' of the meaning $[\![\Phi]\!]$ of Φ. Now consider the following two operators, the first due to Kamp, the second to Vlach:

- $[\![\ddagger\Phi]\!](i)(i') = [\![\varnothing\Phi]\!]^i = [\![\Phi]\!](i)(i)$
  - Doesn't change the context parameter
- $[\![\dagger\Phi]\!](i)(i') = [\![\varnothing\Phi]\!]^{i'} = [\![\Phi]\!](i')(i')$

This last comes to: the proposition expressed by Φ is true; where "the proposition expressed by Φ" is *non-rigid*, varying with circumstance of evaluation. This is a nonbenign monster.

Kaplan's point is just that Double Index Theory makes no principled distinction between these two operators:

> [Double Index Theory] allows a simple and elegant introduction of many operators which are monsters. In abstracting from the distinct conceptual roles of played by contexts of use and circumstances of evaluation the special logic of indexicals has been obscured. Of course restrictions can be put on the two-dimensional logic to exorcise the monsters, but to do so would be to give up the mathematical advantages of that formulation. (Kaplan 1989, p. 512)

This is as good a place as any to enter a caveat about identifying Montague as a pure index theorist. Montague distinguishes within what we might call a 'generalized index' two parts, an index proper—not to be confused with a proper index—and a context of use. The *meanings* of closed sentences are functions from generalized indices to $\mathcal{B}$; but the *senses* of

sentences are functions only from indices proper to $\mathcal{B}$. This means that the analogue of two-dimensional functions are ruled out: functions from contexts (or from generalized indices) into **2** cannot be the senses (contents) of sentences and hence can not be arguments to modal (intensional) operators. No monsters!! Thus Montague's theory, like double index theory, does distinguish between meanings and contents (propositions/senses) and it prohibits monsters–at least monsters of the type Kaplan discusses. Moreover as we shall see this prohibition of monsters does not differ in form all that much from Kaplan's.

3 Kaplan's Theory

We have noted that a central idea in model-theoretic semantics is that of the relativity of truth. In general what a semantic account provides is a type of truth-valuable entities and a circumstance of evaluation. In the case of modal logic and its generalization in pure index theory, the truth-valuable entities are sentences and the circumstances of evaluation are indices (within structures). In the case of double index theory, the truth-valuable entities are propositions (functions from indices to $\mathcal{B}$) and the circumstances of evaluation are again indices. Given what Kaplan says in the descriptive, philosophical sections of Kaplan 1989, one might expect something similar in his logic of demonstratives, except that a clear conceptual distinction would be made between contexts and circumstances. That is, one would expect a two-step theory, modeling the meanings—Kaplan's term is 'character'—of sentences as functions from contexts to propositions, which are in turn functions from circumstances to $\mathcal{B}$. In fact, though, he present a version of a single index theory that, like Montague's, divides the one vector into two parts: context and circumstance, context itself being modeled as a quadruple consisting of an agent, a time, a position, and a world. Again, as in Montague, circumstances (indices proper) are pairs of worlds and times.[3] So the truth-evaluable entity in Kaplan's account are sentences-in-a-context:

> The Content of a sentence in a context is, roughly, the proposition the sentence would express if uttered in that context. This description is not quite accurate. First, it is important to distinguish an *utterance* from a *sentence-in-a-context*. The former notion is from the theory of speech acts, the latter from semantics. Utterances take time, and utterances of distinct sentences cannot be simultaneous (i.e., in the same context).[4] But to develop a logic of demonstratives it seems most natural to be able to evaluate several premises and a conclusion all in the same

[3] We are, of course, ignoring relativity to assignments.

[4] This is gratuitous; one can chose the granularity of the temporal dimension to suit one's purposes.

> context. The the notion of Φ *being true in* c *and* $\mathcal{A}$ does not require an utterance of Φ. In particular, c_A (the agent of the context) need not be uttering Φ in c_W at c_T. (Kaplan 1989, p. 546)

We remind the reader of Kaplan's two basic principles about demonstratives and indexicals:

Principle 1 The referent of a pure indexical depends on the context, and the referent of a demonstrative depends on the associated demonstration.

Principle 2 Indexicals, pure and demonstrative alike, are directly referential.

What Kaplan means by Principle 2 is that the referential relation between, e.g., an indexical, as occurring in a sentence Φ in a context and its referent is not mediated by the content of the sentence in that context. We return to this below (maybe).

We present a brief sketch of Kaplan's *Logic of Demonstratives* in outline form:

The Formal System $\mathcal{LD}$:

- $\mathcal{A}$ is a $\mathcal{LD}$ structure iff there are $\mathcal{C}, \mathcal{W}, \mathcal{U}, \mathcal{P}, \mathcal{T}, \mathcal{I}$ such that:
  - $\mathcal{A} = \langle \mathcal{C}, \mathcal{W}, \mathcal{U}, \mathcal{P}, \mathcal{T}, \mathcal{I} \rangle$
  - $\mathcal{C}$ is a nonempty set (of contexts)
  - If $c \in \mathcal{C}$, then
    - $c_A \in \mathcal{U}$ (the *agent* of c)
    - $c_T \in \mathcal{T}$ (the *time* of c)
    - $c_P \in \mathcal{P}$ (the *place* of c)
    - $c_W \in \mathcal{W}$ (the *world* of c)
  - $\mathcal{W}$ is a nonempty set (of worlds)
  - $\mathcal{U}$ is a nonempty set (of all—actual and possible—individuals)
  - $\mathcal{P}$ is a nonempty set (of positions—common to all worlds)
  - $\mathcal{T}$ is the set of integers (thought of as times, common to all worlds)
  - $\mathcal{I}$ is the interpretation function, assigning pairs of times and worlds to wffs., and meeting the following conditions:
  - $i \in \mathcal{U}$ iff $(\exists t \in \mathcal{T})(\exists w \in \mathcal{W})(\langle i \rangle \in \mathcal{I}_{\text{Exist}}(t, w))$
  - If $c \in \mathcal{C}$, then $\langle c_A, c_P \rangle \in \mathcal{I}_{\text{Located}}(c_T, c_W)$
  - If $\langle i, p \rangle \in \mathcal{I}_{\text{Located}}(t, w)$ then
$$\langle i \rangle \in \mathcal{I}_{\text{Exist}}(t, w)$$
- **So all contexts are proper.**
- **Truth and Content**

  Truth $\models^{\mathcal{A}}_{ctw} \Phi$ for: Φ,in context c is true with respect to time t and world w.

Denotation $|\alpha|_{cftw}$ for *the denotation of* α, in context c (under f) with respect to time t and world w.

Content Where Φ is a wff.,$\{\Phi\}_c^{\mathcal{A}}$ for the **content** of Φ in c.

- $\{\Phi\}_c^{\mathcal{A}}(t,w) =$ TRUTH iff $\models^{\mathcal{A}}_{ctw} \Phi$.

Truth in a context Φ is *true in* c in $\mathcal{A}$ iff $\{\Phi\}_c^{\mathcal{A}}(c_T, c_W) =$ TRUTH.

Validity Φ is *valid in* $\mathcal{LD}$ iff for every $\mathcal{LD}$ structure $\mathcal{A}$ and every $c \in \mathcal{A}$, Φ is true in $\mathcal{A}$.

Character $\{\Phi\}^{\mathcal{A}}(c) = \{\Phi\}_c^{\mathcal{A}}$

- **The crucial clauses of the definition of satisfaction**
  1. $\models_{cftw} R\alpha_1 \ldots \alpha_n$ iff $\langle |\alpha_1|_{cftw}, \ldots, |\alpha_n|_{cftw}\rangle \in \mathcal{I}_R(t,w)$
  2. $\models_{ctw} \Box\Phi$ iff $\forall w' \in \mathcal{W} : \models_{ctw'} \Phi$
  3. $\models_{ctw} A\Phi$ iff $\models_{ctc_W} \Phi$
  4. $\models_{ctw} N\Phi$ iff $\models_{cc_T w} \Phi$
  5. $|I|_{ctw} = c_A$
  6. $|\text{Here}|_{ctw} = c_P$
- **Crucial cases**
  1. $\models (\Phi \leftrightarrow AN\Phi)$
  2. $\models N$(Located, I, Here)
  3. $\models$ Exist I
  4. $\not\models \Box(\Phi \leftrightarrow AN\Phi)$
  5. $\not\models \Box N$(Located, I, Here)
  6. $\not\models \Box$Exist I

Now what of monsters? In Kaplan's theory, the meaning or character of a sentence Φ is a function: $\mathcal{C} \rightarrow (\mathcal{I} \rightarrow \mathcal{B})$ where $\mathcal{I} = (\mathcal{T} \times \mathcal{W})$. To rule out monsters is to rule that there are no functions of the type: **Character** $\rightarrow$ **Character**. This is to stipulate that that no operator can effect (any component of) the c component of the index $\langle cftw\rangle$. This rules out the analogue of the † operator above and is directly analogous to Montague's prohibition of operators on contexts or generalized indices.

4 The Veil of Ignorance

As Kaplan makes clear, the notion of Content, of what is said, is central to his account. At the level of sentential content, of proposition expressed, the content/reference determining features of the context have been applied and as, in function application generally, no trace of them remains. The features of the sentence-in-a-context that determine what is said are not part of what is said. Operators such as the alethic and temporal modalities only apply to propositions expressed: this is precisely what the stipulation against monsters comes to. But what of other attitudes—in particular what of the propositional-attitude operators, more narrowly construed. Though Kaplan has much to say about issues concerning the interaction between indexicals/demonstratives and such operators, he does not introduce them into $\mathcal{LD}$.

Consider knowledge. Imagine we decide to add a family of unary modal operators K_a, indexed by agents $a \in \mathcal{U}$ to $\mathcal{LD}$ and to introduce (structurally identical) associated binary accessibility relations for them, where the intuitive reading of the relationship is that wK_aw' iff w' is an epistemic alternative for a relative to w. But now we should remind ourselves of the following facts about actual utterances and the contexts in which they are produced:

- One might not know who the agent of c is.
- One might not know when the time of c is.
- One might not know what the place of c is.
- One might not know what the world of c is.

Indeed, a speaker himself might be ignorant of the fact that he was the speaker of a given utterance. Consider the case of echoes, especially as produced at a famous and much-visited location. So given a type for an utterance, that is, given a sentence Φ, other contexts for Φ are epistemic alternatives. To see what may be involved here, let us return to a simple double index account, in which the basic indices are just worlds. Here the 'context' index would represent the epistemic perspective of the agent and the circumstance index would, as usual represent, the world about which the knowledge claims are made. The clause for K (now suppressing indexing by agent) would be as follows:

- $w \models_v K\Phi$ iff $\forall w', v' : \langle w, v\rangle R_K \langle w', v'\rangle \rightarrow w' \models_{v'} \Phi$.[5]

Notice that this operator involves quantification over contexts or generalized indices. It is a non-benign monster.

What would this look like in $\mathcal{LD}$? In conformity with Kaplan's restriction, and supposing for simplicity that indices proper—circumstances—are just worlds, all he would allow us is this:

- $cw \models K\Phi$ iff $\forall w' : wR_Kw' \rightarrow cw' \models \Phi$.

But to capture the facts about ignorance, what we need is rather more like this:

- $cw \models K\Phi$ iff $\forall c', w' : \langle c, w\rangle R_K \langle c', w'\rangle \rightarrow c'w' \models \Phi$; where $c' = \langle c'_a, c'_t, c'_p, c'_w\rangle$.

This, of course, is monstrous.

In sum:

- Perhaps there is something right about Kaplan's prohibition, but it is not quite right. Perhaps there could not be pure modal monsters, but there can be epistemic (and deontic, etc.) monsters.

[5]See Rabinowicz and Segerberg 1994 for a similar treatment of knowledge in a two-dimensional context, motivated by very different concerns.

- Double indexing has no explanation of the lack of modal monsters; Kaplan's theory does not allow the epistemic ones.

We claim that an utterance-based theory explains why there can be epistemic monsters, but no modal monsters, and also clarifies Kaplan's fundamental distinction between contexts and circumstances of evaluation.

5 Utterances

5.1 The Reflexive-Referential Theory: A Look Ahead

Utterances are the fundamental truth-evaluable entities. Utterances are acts, concrete nonrepeatable particular events.[6] Utterances are not to be confused with tokens. Tokens are also concrete particulars, but they are objects, not events. Tokens are reusable, in way that utterances are not. The distinction between utterance and token is fairly easy to see with respect to written tokens. Written tokens typically have longer duration that the act—the utterance—that produced them. They are composed of chalk or graphite; they can be erased or underlined. Not so the utterance. In speaking we produce more-or-less evanescent tokens whose perceptible existence doesn't much outlive the duration of the utterance that produced them. Still these tokens can be recorded and then manipulated in various ways. Utterances, speakings, can be recorded, too; indeed they can be filmed without sound. For many years, that's what movies largely consisted in: recordings of utterances, without any recording of the tokens produced. In the case of computer files, the distance between utterance and token is quite large; indeed, it is a little mysterious as to what the token produced is. It is utterances, not tokens, that are the primary bearers of truth and falsity.

As Kaplan notes, utterances take time. An utterance also has a particular speaker and a place. All actual utterances have the same world, namely the actual world. We claim, without argument, that the agent, time and place of an act are metaphysically essential features of that act. It makes no sense to say that its world is also metaphysically essential to it.

Consider now what a competent speaker/hearer of English knows about the truth-conditions of an utterance of "I am tired" solely on the basis of his linguistic knowledge, that is, in the absence of knowing who said it and when.

- A utterance u of "I am tired" is true iff *the speaker of u is tired at the time of u*.

The italicized condition yields a proposition when predicated of a particular utterance. So consider a particular utterance **u**; it is true iff *the speaker of* **u** *is tired at the time of* **u**. We call this proposition the reflexive

[6] For more on acts and actions, see Israel et al. 1991, Israel et al. 1993.

content of **u**—more particularly, its *indexical content*.[7] What Kaplan calls the content, what is said by the utterance, we call the *incremental content*; it is generated from the reflexive content, given all the features of the utterance and context that determine reference. In the case of pure indexicals, these are features of the utterance itself. In the case of demonstratives, these will include features of the wider context of utterance.

As we have seen, the incremental content does not (usually) involve the utterance or its context; these have been used in determining what is said, they are not a part of it. But that need not be true for the other contents.[8] Assume a competent speaker/hearer of English encounters **u**, but is ignorant of who said it and when. He knows its truth-condition, its pure reflexive content. In this case, he is ignorant of its content in two different ways: he is ignorant of what the content is, of what proposition the utterance expresses, because he is ignorant of the context of utterance, and he is ignorant of the truth-value of that proposition. Formalizing such facts motivates the monstrous treatment of knowledge sketched above.

5.2 Monsters, Revisited

We have claimed that agent, time and place of an utterance are essential features of it. This explains why there cannot be metaphysical monsters. You can't 'take' an utterance to a metaphysical alternative and leave its reference and truth determining features behind. That is, there are no real modal alternatives with respect to the context of utterance.

One can, on the other hand, take a sentence uttered in one index to another index, and one can take a sentence and an agent to another world and another time. Thus if, like Kaplan, one takes as the prime truth-evaluable entity a sentence-in-a-context, where a context is an n-tuple but without a representative of the particular utterance, it is hard to justify the stipulation that no expression of the language can involve a shift in context. Indeed we have seen that because none of the metaphysically essential features of utterances are epistemically transparent features, in modeling knowledge it seems that we want to be able to shift context. We should note that the claim of asymmetry between modal and epistemic monsters does not commit us to the view that there are real monsters in English.

In the full paper, we will develop the referential-reflexive account in sufficient detail to justify our claim that it generates adequate answers to the problems posed by monsters.

[7] This is precisely to leave open the possibility of there being other kinds of reflexive content, e.g., that associated with uses of proper names.

[8] Note that in general where there are n independent dimensions of context-relativity exploited in an utterance, there will be $2^n - 1$ reflexive contents.

References

Israel, David, John Perry, and Syun Tutiya. 1991. Actions and Movements. In *Proceedings of IJCAI'91*. Sydney. Morgan Kaufmann.

Israel, David, John Perry, and Syun Tutiya. 1993. Executions, Motivations, and Accomplishments. *Philosophical Review* 102:515–540.

Kamp, Hans. 1971. Formal Properties of 'Now'. *Theoria* 37:227–274.

Kaplan, David. 1989. Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals. In *Themes from Kaplan*, ed. Joseph Almog, John Perry, and Howard Wettstein. 481–614. New York: Oxford University Press.

Montague, Richard. 1974. Pragmatics. In *Formal Philosophy*, ed. Richmond Thomason. 95–118. New Haven: Yale University Press.

Rabinowicz, Wlodek, and Krister Segerberg. 1994. Actual Truth, Possible Knowledge. In *Proceedings of TARK V*, ed. Ron Fagin, 122–137. Asilomar. Morgan Kaufmann.

Scott, Dana. 1970. Advice on Modal Logic. In *Philosophical Problems in Logic*, ed. Karel Lambert. 143–173. Dordrecht: Reidel.

Segerberg, Krister. 1973. Two-Dimensional Modal Logic. *Journal of Philosophical Logic* 2:77–96.