

Appeared in *Journal of Philosophical Logic* 24 (1995), pp. 227-240.

What is Character?

David Braun

University of Rochester

In "Demonstratives", David Kaplan argues that indexicals and other expressions have a sort of meaning, character, that is distinct from their more commonly recognized sort of meaning, content. In my opinion, Kaplan's arguments for this conclusion are successful and extremely significant.

But I am troubled by Kaplan's attempts to describe character. For one thing, Kaplan is not consistent in his mathematical representation of character. In some places, he represents character with one sort of function; in other places, he uses another sort of function. I show this in section one below, and argue that one of these representations is clearly better than the other. But I believe that even the better of these two representations is inadequate for capturing the kind of meaning that Kaplan calls 'character'. In section two, I argue that no extensional function can do the job adequately.

The problems I find with Kaplan's theory of character are near duplicates of well known problems with traditional possible worlds semantics. By this I mean theories of the sort inspired by Carnap (1956), together with the Kripkean idea of treating worlds as "points" of evaluation. The basic idea of such theories is that an expression's meaning (in the sense of intension) is determined by its value (extension) at a set of points of evaluation (worlds). Kaplan's theory of

character is identical in form: an expression's meaning (character) is determined by its value (content) at a set of points of evaluation (contexts). Given these similarities, it is perhaps not surprising to find problems with Kaplan's theory that resemble those that afflict traditional possible worlds semantics. I make the comparisons explicit in section three. I conclude with some reflections on the recurring inadequacies of "point semantics".

1. Which Function?

Kaplan says "it is convenient to represent character by functions from possible contexts to contents" (p. 505---all page references to Kaplan 1989). But Kaplan does not stick consistently with this characterization of character. There is a mismatch between this description of character, and the definition of character he gives in his formal system.

To show this, I need a few preliminary observations. If character can be represented by a function from contexts to contents, then any two expressions that have the same content in every context have the same character. For instance, 'attorney' and 'lawyer' have the same content in every context; thus they have the same character, on this representation of character. The context-sensitive sentences 'I am an attorney' and 'I am a lawyer' also have the same content in every context, and therefore have the same character. Co-referring proper names, like 'Cicero' and 'Tully', have the same character because (as Kaplan points out, p. 502) they refer to the same person in every context and so have the same content in every context. Thus the sentences 'Cicero is Cicero' and 'Cicero is Tully' have the same character.¹ All of this seems fine, if character is a kind of meaning. After all, there is no reason why two expressions (even two simple expressions) should not agree at all levels of meaning, including character, and so be

synonymous in the strictest sense.

Since Kaplan says that character can be represented as a function from characters to contents, we might expect that in his formal logic, every character will be represented within a structure by a function, from contexts in that structure, to contents in that structure.

Consequently, we might expect the character of an expression to vary from structure to structure, because the set of contexts (along with the sets of individuals, times, and worlds) varies from structure to structure. So one would expect sentences like 'I am an attorney' and 'I am a lawyer' (or their formal surrogates) to have the same character in one structure (for instance, in the "intended structure"), but different characters in other structures. Similarly for 'Cicero is Cicero' and 'Cicero is Tully'.

But surprisingly, Kaplan's formal system does not satisfy these expectations. I quote his definition of 'Character' (p. 548).

Definition: Where Gamma is either a term or a formula, the Character of Gamma is that function which assigns to each structure A, assignment f, and context c of A, $\{\text{Gamma}\}_A, cf$.

(Here $\{\text{Gamma}\}_A, cf$ is the content of Gamma in A and context c, under assignment f.) Notice that this is not a definition of the term 'the Character of Gamma in A'. The term defined is simply 'the Character of Gamma'. So the Character of an expression is not relativized to structures. Moreover, the definition says that the Character of an expression is a function from structures, and contexts in those structures, to contents.² So, contrary to Kaplan's informal remarks, these functions are not functions from contexts to contents within structures; they are defined over structures. Furthermore, on this definition of Character, an expression's Character

does not vary from structure to structure. Indeed, the idea that the Character of an expression could vary in this way makes no sense on this definition.³

Could I be misinterpreting Kaplan? The definition might be open to alternative interpretations; perhaps Kaplan's intended definition is really the one we expect. But Remark 12 (pp. 551-2) confirms the suspicion that Kaplan's formal definition of Character really is "cross-structural". I quote again.

Remark 12: The strongest form of logical equivalence between two formulas ϕ and ϕ' is sameness of Character, $\{\phi\} = \{\phi'\}$. This form of synonymy is expressible in terms of validity.

$\{\phi\} = \{\phi'\}$ iff $\models \square$ etc.

Notice that Kaplan is not trying to state conditions under which two formulas have the same character in a single structure. Rather, he is trying to tell us when two formulas have the same Character, period. Moreover, this Remark says that ϕ and ϕ' have the same Character iff the modal formula on the right hand side is valid: that is, true in every context of every structure. Thus it's clear that Character, like validity, must be a notion that is defined over structures, rather than within structures.⁴

So Kaplan's formal definition of Character is not what we would expect from his informal descriptions of character. Moreover, his formal definition has consequences that clash in important ways with his informal characterizations of character. For instance, according to the formal definition of Character, the sentences 'I am an attorney' and 'I am a lawyer' cannot have the same Character. The reason is that in some structures the predicates 'attorney' and 'lawyer' have different contents (in every context); thus by the formal definition of Character, these

expressions have different Characters.⁵ It follows that the modal formula mentioned in Remark 12 is not valid when (the formal surrogates of) these sentences are substituted. Similarly, for 'Cicero is Cicero' and 'Cicero is Tully'. In some structures the names do not co-refer; in those structures, the names differ in content in every context. Thus they do not have the same Character, and the modal formula is not valid.⁶

So we have a mismatch between the notion of character that Kaplan presents in his informal remarks, and the notion of Character he presents in his formal system. Which is the more appropriate functional representation of the intuitive notion of character? Should we prefer the informal notion, and represent character as a function from contexts to contents within a structure (that is, within the "intended structure")? Or should character be represented as a function from structures and contexts to contents, as in the formal definition? To me, it is obvious that the first, informal representation, better characterizes the kind of meaning that Kaplan wishes to capture. The only question for me is, why does Kaplan choose the second sort of function in his formal system?

The first sentence of Remark 12 gives us a hint. Kaplan wants sameness of character to be a kind of logical equivalence.⁷ But if he defines character in the way his informal remarks suggest, then this will not be so. For there are some sentences that have the same character within some single structure (for example, within the intended structure), but which do not have the same character in all structures: for instance, the sentences 'Cicero is Cicero' and 'Cicero is Tully'. So long as Kaplan defines validity as he does, these pairs of sentences are not logically equivalent.⁸

Perhaps Kaplan should change his notion of validity, or (better) give up the idea that

sameness of character is a kind of logical equivalence. I discuss this further in section 3. But in any case, if we want to represent the intuitive notion of character with an extensional function, we had best represent it by a function from contexts to contents within a single structure (the intended structure).

2. A Function?

But now I want to argue that no extensional function is fully adequate for capturing the kind of meaning that Kaplan calls 'character'. I can best explain this by starting with Kaplan's notion of content.

In "Demonstratives", Kaplan introduces the notion of a structured content: a content that is structured much like a linguistic expression and that has individuals and relations as constituents. Kaplan uses these structured contents to make certain "fine-grained" distinctions in content. For instance, he uses them to distinguish between propositions that are necessarily equivalent. And, most importantly, he uses structured contents to distinguish between rigid designators that are, and rigid designators that are not, directly referential.

An important feature of these propositions (and of contents in general) is that the relations that appear in them are primitive, unreduced, intensional entities. They are not sets; they are not Carnapian functions from circumstances of evaluation to sets. They are not extensional entities. (From here on, I use 'property' and 'relation' to mean such primitive intensional entities---not their extensional surrogates.)

Kaplan never makes an analogous proposal about character. He never suggests that characters are non-extensional entities, and he never suggests that characters are structured. But

once we have contents that are structured entities containing primitive non-extensional entities, it's natural to think that characters are similar sorts of entities. In particular, characters of simple expressions could be relations between contexts and contents, rather than extensional functions; characters of complex expressions could be structured entities containing relational characters as constituents. This view is consistent with much of what Kaplan says. When he says, "it is convenient to represent character by functions from possible contexts to contents" (p. 505, my emphasis), Kaplan at least leaves it open that characters are not simply extensional functions on contexts.⁹

I believe that there are certain fine-grained distinctions in meaning, and in character in particular, that we cannot make using only functions-in-extension. Thus I believe we should accept the view that every character is either a primitive relation, or a structured entity containing these primitive relations as constituents. I will present two cases where such fine-grained distinctions seem to be called for.

Here is the first case. Kaplan claims, plausibly enough, "To supply a synonym for a word or phrase is to find another with the same character . . ." (p. 521). But if the character of a complex expression is an extensional function from contexts to contents, then this claim is almost certainly mistaken. Suppose we have a language that contains Kaplan's dthat-terms. (Kaplan seems to have this sort of language in mind when he makes the preceding claim.) Now consider the following two dthat-terms.

dthat[the x such that: x=Nixon]

dthat[the x such that: (Snow is slight & x=Nixon) or (-Snow is slight & x=Nixon)].

The content of a that-term, in a context, is simply the referent of the enclosed singular term, in that context. Thus the above that-terms are logically guaranteed to have the same content in every context.¹⁰ So if character is merely an extensional function from contexts to contents, then these terms have the same character. But they are not synonymous. So if sameness of character is supposed to entail synonymy (that is, sameness of meaning), then we should not take character to be a mere extensional function on contexts.¹¹

This first problem could be solved using extensional, set-theoretic tools. We could hold on to the idea that the characters of simple expressions are merely extensional functions on contexts, yet also maintain that the characters of complex expressions are structured entities that contain the basic extensional characters as constituents.¹²

But adding structure is not enough to allow us to adequately represent character, I believe. For functions-in-extension do not adequately represent the characters of even simple indexicals, like 'I'. Let me explain why.

Kaplan maintains that indexicals have "a limited kind of specific descriptive meaning relevant to the features of a context of use" (p. 497). Kaplan repeatedly makes this sort of claim (pp. 486, 498, 505 n. 31, 520 n. 44). He tries to express this "descriptive meaning" of 'I' using phrases like 'the speaker', 'the person speaking', 'the agent' and 'the agent of the context'. It's clear that this descriptive meaning is not supposed to be the content of 'I' in a context, because Kaplan says, "it is relevant only to determining a referent in a context of use, and not to determining a relevant individual in a circumstance of evaluation" (p. 498). So a reasonable hypothesis is that Kaplan is talking about what he calls the 'character' of 'I' when he mentions its "descriptive meaning", and that he wishes to use character to capture this "descriptive meaning" within his

semantics.¹³ Let's suppose that this is Kaplan's intent. Can an extensional function on contexts do the job of capturing this "descriptive meaning"? I think not.

I believe that when we use phrases like 'the speaker' or 'the agent of the context', we are trying to express a particular relation between individuals and contexts, a relation that might be better expressed by the phrase 'is the agent of'. One reason for thinking this is that when we try to describe the meaning of 'I', we are unwilling to use expressions that are co-extensive with 'is the agent of', but which do not express the same relation. Now a function-in-extension can have the same extension as the relation that we wish to express. But it cannot completely "capture" or "represent" such a relation.

Let me explain in a bit more detail. The relation that we wish to express, when we use phrases like 'the agent of the context' to describe the meaning of 'I', is something like the relation described in (1).

- (1) being the agent of: the relation that holds between (individual) x and (context) c iff x is the agent of c .

Now consider the following relation.

- (2) being the agent of, and such that either Pakistan is beautiful or Pakistan is not beautiful: the relation that holds between x and c iff x is the agent of c , and either Pakistan is beautiful or Pakistan is not beautiful.

(1) and (2) are co-extensive relations. But they are distinct. Of course, there are many other such co-extensive relations.¹⁴ But only (1) is a plausible candidate for being a meaning of 'I', or the character of 'I'.

To see this, let's suppose that we tried to describe the meaning of 'I' using a phrase that

unambiguously expressed relation (2). For instance, suppose we said the following: 'I' means 'is the agent of, and such that Pakistan is beautiful or Pakistan is not beautiful'. Or suppose we tried to express this a bit more precisely by means of a reference rule.

In every context c, 'I' refers in c to the x such that (x is the agent of c, and either Pakistan is beautiful or Pakistan is not beautiful).

This gets the reference of 'I' right, but clearly gets something important wrong. A person learning English would learn the wrong meaning for 'I' (in the sense of character) if she thought that this rule explained the meaning of 'I'. Thus we do seem to intend to express relation (1), and not (2), when we try to state the "descriptive meaning" of 'I'.

But we cannot completely capture or represent a relation like (1) with a function-in-extension on contexts. A function-in-extension can represent relation (1) only by being co-extensive with it. But then it equally well represents relation (2), and many others as well. Thus we need to take a step beyond extensional functions to capture the meaning of 'I'. In fact, we should take a step like the one that Kaplan himself took in his theory of content. We should say that the characters of simple expressions are genuine relations.¹⁵

3. Problems with "Point Semantics"

As I said earlier, the problems I am raising for Kaplan's theory of character are near duplicates of problems that arise for "possible worlds theories of content". I think that some general lessons might be drawn from this.

The basic idea of traditional possible worlds semantics is that an expression's meaning is determined by its value (extension) at each of a set of points of evaluation (worlds). The

expression's meaning, its intension, is a function that delivers the right value at each point. So two expressions have the same meaning, or intension, iff they have the same value at each point. Kaplan's theory of character has the same form---but contexts play the role of points, and contents play the role of values. Expressions have values (contents) at points (contexts). The meaning (character) of an expression is determined by these values; two expressions have the same meaning (character) iff they have the same value (content) at each point.

The problems with Kaplan's theory are (at least partly) due to its evaluation-at-points format. This is somewhat surprising, because Kaplan's "three-tiered" system is meant to avoid some of the problems of the earlier "two-tiered" systems, especially when contents are allowed to be structured. But let me explain how the problems are analogous.

The first problem with traditional possible worlds semantics has to do with meaning and logical equivalence. Early proponents of possible worlds semantics, like Carnap, said that two expressions have the same intension, and thus the same meaning (in the sense of content), iff they are logically equivalent. (See Carnap 1956, secs 1-6.) These early proponents of possible worlds semantics endorsed a three way equation: x and y are logically equivalent iff they have the same intension iff they have the same meaning (content).

This equation is likely to be unacceptable to anyone who holds (contrary to Carnap) that the proper names 'Cicero' and 'Tully' are rigid designators. For if these names are rigid designators, then the sentences 'Cicero is Cicero' and 'Cicero is Tully' have the same intension. But one might reasonably doubt whether these sentences are logically equivalent.

The traditional way to avoid this problem involves three steps. First, relativize the intension of an expression to a structure. Second, say that two sentences are logically equivalent

iff (roughly) they have the same intension in every structure. Third, say that two sentences have the same meaning (in the sense of content) iff they have the same intension in the intended structure. Notice, however, that this solution gives up the identification of logical equivalence with sameness of meaning.

Kaplan makes claims about character that are similar to Carnap's claims about meaning (content). Kaplan wants character to be a kind of meaning; he also wants sameness of character to be a kind of logical equivalence. As a result, Kaplan has problems like Carnap's. For instance, according to Kaplan's informal description of character, 'Cicero is Cicero' and 'Cicero is Tully' have the same character, even though it is doubtful whether these sentences are logically equivalent.

Kaplan could avail himself of a solution to his problem that is similar to the old solution to the old problem. He does take one step in this direction: he says that logical equivalence is truth of a certain formula in (every context of) every structure.¹⁶ But he fails to take the other steps: he does not relativize character to structure, and he does not say that sameness of meaning is sameness of character in the intended structure. Since Kaplan takes only the first of these steps, his formal system ends up entailing that 'Cicero is Cicero' and 'Cicero is Tully' have different Characters, contrary to his own intentions.

The second problem for possible worlds theories of content is more familiar. Consider the following definite descriptions.

the x such that: $x=Nixon$

the x such that: $(Snow\ is\ slight\ \&\ x=Nixon)\ or\ (-Snow\ is\ slight\ \&\ x=Nixon)$.

These two expressions have the same extension at each world. So they have the same intension.

But clearly they differ in meaning (in the sense of content). The that-terms that I displayed earlier raise an analogous problem for Kaplan's theory of character: they have the same content at each context, and yet seem to differ in meaning (in the sense of character).

A well known way of escaping this problem for possible worlds semantics adds structure to intensions. (See Carnap 1956, secs. 14-16; and Lewis 1970, sec. 5.) A structuring of characters would also solve the analogous problem for character. But if my argument concerning the "descriptive meaning" of 'I' is correct, then this is not enough to fix Kaplan's theory of character. We still need primitive relations to serve as the characters of simple expressions.

What lessons can be drawn from these comparisons of Kaplan and Carnap? One lesson is that it is risky to equate sameness of meaning, at any level, with logical equivalence. But I believe there is a second, more serious, lesson that might reasonably be drawn from this story. Carnap and Kaplan attempt to represent two quite different sorts of meaning in the same traditional way, by using extensional functions on points of evaluation. But both of their attempts lead to the same problems. This result suggests that there is something fundamentally misguided in the traditional approach to meanings. It strongly suggests that no meaning is an extensional function on points of evaluation.¹⁷

References

Barwise, Jon and Perry, John. 1983. Situations and Attitudes. Cambridge, MA: MIT Press.

Braun, David. Forthcoming. "Structured Characters and Complex Demonstratives."

Philosophical Studies.

Carnap, Rudolf. 1956. Meaning and Necessity (2nd edition). Chicago: University of Chicago Press.

Kaplan, David. 1989. "Demonstratives" and "Afterthoughts". In Joseph Almog, John Perry, and Howard Wettstein (eds.), Themes from Kaplan, pp. 481-614. Oxford: Oxford University Press.

Lewis, David. 1970. "General Semantics." Synthese, 22, pp. 18-67. Reprinted in David Lewis, 1983, Philosophical Papers, Vol. I, pp. 189-229. Oxford: Oxford University Press.

Richard, Mark. 1983. "Direct Reference and Ascriptions of Belief." Journal of Philosophical Logic, 12, pp. 425-52.

Notes

1. I am assuming a simple compositionality principle: very roughly, that the meaning of a complex expression depends only on the meanings of its parts, and not, for instance, on how many occurrences of the same expression occur in the complex expression. In lectures, Kaplan has proposed rejecting this principle, and proposed that 'Cicero is Cicero' and 'Cicero is Tully' express different propositions, which we can respectively represent (roughly) as follows:

< Cicero, self-identity >

< <Cicero, Cicero>, identity >

Notice that even on this view, the sentences 'Cicero is an orator' and 'Tully is an orator' have the same content (in every context), and thus the same character. Since these latter sentences would serve many of my purposes equally well, I will continue to use the identity sentences, and set aside Kaplan's views about them.

2. I ignore relativization to assignments from here on.

3. As I said above, we would naturally expect the notion of character to be represented by different "structure-internal" functions in different structures. Such a structure-internal function can be obtained from Kaplan's "cross-structural" function by restricting the latter function's "structure argument" to a single structure, for instance, to the intended structure. (For this reason, Terry Parsons has suggested to me that the cross-structural function for an expression be called its 'Character Potential'.) But Kaplan does not mention this way of obtaining a structure-internal function, and he does not give the resulting function a name (for instance, 'the Character of Gamma in A'). The apparent reason for Kaplan's choice emerges in the text below.

4. The proof of the equivalence in the left to right direction depends crucially on the fact that two formulas have the same Character (in the sense of the formal Definition) iff they have the same truth value with respect to every context, time, world, and assignment in every structure. So the proof falls through if Character is relativized to structure.

5. Kaplan does not formally define the notion of the content of a predicate in a context and structure. He defines this notion only for formulas and terms. But his structures do assign intensions to predicates. These intensions assign extensions to world-time pairs, in the same way that contents of predicates (in contexts) should. Furthermore, the intension of a predicate does not vary from context to context, and neither should the content of a predicate. So the content of a predicate in a context and structure in Kaplan's formal system is most naturally identified with the intension assigned to it by the structure.

6. As Kaplan points out (Remark 11, p. 551), his formal system does not contain proper names. But Kaplan says that proper names could (in effect) be added to the system. What we would need is a set of distinguished individual constants (0-0-place i-functors), and semantic rules for those constants which would assure that each has a Stable Character and a Stable Content in

every context. (See again Remark 11, p. 551.) This could be done by requiring that structures assign constant intensions to these expressions (the values of which at world-time pairs are individuals), and by stipulating that the denotation of one of these expressions with respect to a structure, context, assignment, world, and time is the value of the intension assigned to the name in that structure, when applied to that world and time (under that assignment). My argument in the main text assumes that different structures may assign different intensions to the same name. So these names are not treated as logical constants, or treated in the way that Kaplan treats the indexicals.

7. For a closely related motivation, see Corollary 3, p. 539, where Kaplan asserts that character is the bearer of logical truth. Kaplan is more hesitant in "Afterthoughts". There he says (p. 596) that validity "...[is] a property of sentences (or perhaps characters)..."

I should note that Kaplan seems to make use of two different notions of logical equivalence in his paper. One is the notion I have been discussing thus far, the notion of sameness of Character defined in Remark 12. The second notion is more standard: ϕ and ψ are logically equivalent iff $[(\phi \text{ iff } \psi)]$ is valid (true in every context of every structure). Kaplan seems to have this second notion in mind when he says, on p. 539, "I call ϕ and ψ logically equivalent when $[\phi \leftrightarrow \psi]$ is logically true." This also seems to be the notion that he has in mind when, in note 65 on p. 539, he says that ϕ and $[\text{AN}\phi]$ are logically equivalent. (See also Remark 3, p. 547.) Notice that according to the first notion of logical equivalence (from Remark 12), ϕ and $[\text{AN}\phi]$ are not logically equivalent, since $[(\phi \leftrightarrow \text{AN}\phi)]$ is not valid. (See Remark 3.) So the two notions of logical equivalence are different. But the second notion, like the Remark-12-notion, is also "cross-structural". So even on this second notion, the sentences 'I am a lawyer' and 'I am an attorney' are not logically equivalent. Thus any attempt to express sameness of Character in terms of this second notion of logical equivalence would have problems similar to those I describe in the main text, as long as the same notion of validity is retained.

8. Kaplan could assure that these sentences are logically equivalent in a restricted class of structures---say, all those structures that satisfy certain meaning postulates regarding 'lawyer', 'attorney', 'Cicero' and 'Tully'. Or he could put constraints on the structures similar to those that he places on the interpretations of the indexicals. But I think these options are not appropriate, given Kaplan's goals. Restricting classes of structures by meaning postulates (or other such devices) is appropriate only if one is after a notion of semantic truth (or semantic entailment or semantic equivalence). But Kaplan is trying to give a theory of logical truth. So it seems he should consider more structures, rather than fewer.

9. Kaplan similarly represents contents as extensional functions on circumstances, without abandoning primitive relations as contents (p. 502). Kaplan points out that the representation of characters with extensional functions on contexts has "...the drawback that equivalent characters are identified." (p. 505) It's not clear which "equivalent characters" Kaplan has in mind. At this point in his paper, Kaplan might be entertaining the idea that co-referring proper names can have distinct characters, even though the characters of such names are represented by the same extensional function on contexts. See p. 505, note 31. But Kaplan later clearly rejects this idea about proper names. See p. 562.

10. I am assuming that 'dthat' is a "demonstrative surrogate", as Kaplan puts it, rather than a rigidifier. See pp. 579-81. Given this assumption, the content of both dthat-terms in contexts where Nixon exists (that is, in contexts where Nixon exists in the world of that context) is simply Nixon himself.

11. There are even more obvious examples of non-synonymous dthat-terms with the same functional character, if we are willing to accept some of the conclusions of section 1. Consider the following two dthat-terms.

dthat[the x such that x=Nixon]

dthat[the x such that: if arithmetic is incomplete and $2+3=5$, then $x=Nixon$ and x is self-identical and x is human]

These terms have the same content in every metaphysically possible context. (By a 'metaphysically possible context', I mean a context whose world is metaphysically possible.) This is so because it is metaphysically necessary that arithmetic is incomplete and $2+3=5$, and metaphysically necessary that Nixon is self-identical and human if he exists. But every context in the intended structure of worlds and contexts is metaphysically possible. So these terms have the same content in every context of the intended structure. Now the discussion in section 1 concluded that the extensional function that best represents character (in the intuitive sense) is a function on contexts within the intended structure. So if character is (represented by) any extensional function on contexts, then the above two terms have the same functional character. But it is clear that these terms are not synonymous.

I present a somewhat expanded version of the above argument in my (forthcoming). In that paper, I also argue that structured characters are useful in a semantic theory for complex demonstratives of ordinary English, such as 'that man who is watching Nixon'.

12. The structure itself could also be set-theoretic in nature. My suggestion that we use structured characters is not new: Richard (1983) uses structured characters in his semantics for de se belief reports. (Richard does not mention dthat-terms as a motivation for introducing structured characters into semantics.)

13. As far as I can tell, Kaplan never straightforwardly asserts that the descriptive meaning of 'I' is its character; nor does he say outright that he wants to represent that meaning with character. In fact, a reader might well get the impression that Kaplan thinks that the descriptive meaning of 'I' in a context is a property, such as being the agent of that context. This property is distinct from both the character of 'I' and the content of 'I' in that context. It also varies from context to context: if Cathy and Carl are contexts, then in Carl the descriptive meaning of 'I' (on this view) is the property being the agent of Carl, while in Cathy it is being the agent of Cathy. This alternative view is more in accord with the common use of "descriptive meaning" to designate a property associated with an expression. But I will assume here that the descriptive meaning of 'I' that Kaplan discusses is just its character.

14. For instance, (3) is also co-extensive with (1).

(3) being the agent of and existing: the relation that holds between x and c iff x is the agent of c and x exists.

Relation (4) is also co-extensive with (1), assuming that all contexts are metaphysically possible. (See note 11.)

- (4) being the agent of, and being such that $3+3=6$ and arithmetic is incomplete: the relation that holds between x and c iff x is the agent of c and $3+3=6$ and arithmetic is incomplete.

15. Barwise and Perry (1983) also hold that the character (or linguistic meaning) of a simple expression is a relation (between "discourse situations" and, roughly, contents). But contrary to the view I express in the text, Barwise and Perry say that the characters of complex expressions are also unstructured relations.

16. As I pointed out in note 7, Kaplan has two different notions of logical equivalence. The two notions use different formulas to define logical equivalence. But both notions are defined "cross-structurally".

17. Thanks to Noa Goldring and Terry Parsons for helpful comments and suggestions. Very special thanks to Joseph Almog, whose detailed comments and suggestions resulted in many significant improvements in this paper.