Building Character

Eliot Michaelson and Jack Woods*

Abstract

Part of knowing a language is knowing of a set of rules for using individual expressions and having a capacity for deriving more complex rules, rules for using whole sentences, from those simpler rules. This capacity helps to explain, among other things, how we can recover information about what a speaker was trying to convey even when we know too little about the surrounding context to determine exactly what they were saying. For all its intuitive interest, this capacity to derive more complex linguistic rules from simpler ones has gone relatively unexplored. We seek to fill this lacuna by offering a compositional model for linguistic rules based on Kaplan's notion of character. To succeed in this, we'll need to tinker with some of Kaplan's own starting assumptions. This tinkering allows us to say some helpful things about logical truth, presupposition, and the meanings of demonstratives—in addition to correcting some persistent misconceptions regarding the compositionality of character.

Nothing spoils fun like finding out it builds character.
—*Calvin & Hobbes*

1 Introduction

Knowing a language is partly a matter of knowing a set of rules. This much is uncontroversial.[1] Some of those rules plausibly pertain to individual ex-

*King's College London and The University of Leeds. This work is entirely collaborative and names are listed in alphabetical order

[1] It remains an open question whether to construe such knowledge as propositional, with the contents being the relevant rules, or rather as some sort of colloquial way of talking about a complex capacity—one which may itself require bits of propositional knowledge, though not necessarily with these rules as contents. See Matthews (2006) for helpful discussion. Nothing in what follows will hinge on this choice.

pressions, others to ways of combining those expressions into more complex expressions, culminating in whole sentences. This much is also commonly taken for granted in both the philosophy of language and linguistics. Yet the nature of these rules is still relatively ill-understood.

To see this, specify everything one might want to know about some context. Formal semantics will then tell us how, for any sentence uttered in that context, to get from word meanings to something like the truth-conditions of this sentence, relative to that context. This, however, doesn't capture everything we know on the basis of knowing the rules of a language. We can infer a great deal about what someone is trying to communicate when they assertively utter a given sentence, merely in virtue of our knowledge of the rules of language—and even when we know nothing about the context.

This fact about contexts and communication has been noted before. In one prominent discussion of it, Stalnaker (1978, p. 80) asks us to consider the following utterance directed at O'Leary:

(1) You are a fool.

According to Stalnaker, an utterance of (1) is to be associated with a PROPOSITIONAL CONCEPT representing a collection of the various propositions it could be expressing, depending on who the token of 'you' refers to (something which Stalnaker represents in matrix form). This collection is intended to model everything that the speaker might be aptly asserting in the context: I'm a fool, O'Leary is a fool, Stalnaker is a fool, etc.

On Stalnaker's way of thinking, we are capable of recovering the propositional concept in virtue of merely hearing the sentence. Then we use it to engage in further reasoning about what was being asserted, even when we lack any real information about the context. That is, we know a reasonable amount about what was asserted by (1) even if we don't know that O'Leary was the target, or who said it, or even where and when it was said.

We take this to be a fairly compelling picture of linguistic reasoning and interpretation: what one can hope to recover on the basis of the utterance alone is something like the range of ways that a sentence could be aptly or felicitously tokened in the context in which it was uttered. On our preferred way of putting things, one should be able to recover a *complex rule of use*, a rule for using the sentence itself, on the basis of an utterance. Thereby, one will be in a better position to reason about how the sentence was being used on this particular occasion, i.e. about what the speaker was trying to convey by means of it. One will also be able see what, if any, information is conveyed over the entire range of felicitous tokenings, i.e. as assertion of (1) conveys to even the most ignorant competent speaker that someone thinks someone is a fool.

To the best of our knowledge, neither Stalnaker nor any of those following in his footsteps has offered a detailed explanation of *just how it is* that listeners are supposed to be able to associate propositional concepts with particular token utterances.[2] We presume that Stalnaker and others who talk about reasoning with propositional concepts will reject the thought that we know the meanings of entire sentences brutely. So, in order to make good on the broader story, we need to account for how we are able to *derive* propositional concepts from the meanings of arbitrary, well-formed strings of words, in conjunction with their mode of combination or syntactic structure.[3]

This is the challenge that we set for ourselves here: explaining how it is that we can *derive* something like a propositional concept, in our view a complex rule of use, *compositionally*—or, roughly, on the basis of the meanings of the constituent parts of the relevant sentence (words) and their mode of combination (or sytactic form) alone.[4] In a sense, our theory is rather simple: linguistic rules are such that, when one has a well-formed sentence, there is a way of fitting the rules of use for each constituent into a rule of use for the whole sentence. In other words, we hold that one important constraint on the shape of linguistic rules is that they had better compose.

Of course, we plan to say more than just this. Our main task below will be to construct a *model* of what linguistic rules might be like, one which allows for complex rules to be generated (fairly straightforwardly) from simple ones and their mode of combination.[5] What's the point of constructing such a model? For one thing, it allows us to make sharpened predictions—something which we take to be a standard desideratum in formal approaches to understanding natural language. More importantly though, the hope is that constructing such a model will both deepen and refine our understand-

[2]We are not the first to note that this is an outstanding challenge for Stalnaker. Lewis (1980, pp. 95–6) also notes this lacuna in the Salnakerian system.

[3]We'll ignore force and mood markers in what follows, as things will be complicated enough without these.

[4]The notion of compositionality will be introduced in greater detail in the next section.

[5]We'll be using some standard tools from model theory to construct our model, but one could construct an equally good model in any number of ways; there is nothing essentially model theoretic about our approach.

ing of the nature of linguistic rules and how they interact with more standard approaches to linguistic content.

The proof of the pudding, as always, is in the tasting. To preview, we take it that the sort of model we'll construct below will help us to say some more sensible things about the meanings of the demonstratives than have typically been said in the past, as well as illuminating some features of LOGICAL TRUTHS and offering a natural account of how the PRESUPPOSITIONS of subsentential expressions pass up to sentential expressions. The sensibleness of these things is, we hope, some indication that our model is a helpful one.

We'll construct our model by drawing on the resources of a rather different tradition to the Stalnakarian one: Kaplan (1989b)'s notion of CHARACTER. We'll be a bit more explicit than usual about typing in Kaplan's system, since the types he assigns to characters will raise some issues for compositionality that we'll have to work to overcome. As will become clear in the course of the discussion, Kaplan puts the so-called true indexicals, like 'I' and 'today', at the heart of his discussion. This, we will argue, makes his model unsuitable for the demonstratives like 'this' and 'that'. But this limitation can be overcome, in part via reflection on what it is that we are trying to model—namely, linguistic rules. The necessary modifications will also allow us to iron out some kinks for Kaplan's system pertaining to his notion of logical truth.

Here is the plan for what remains. In §2, we'll offer a bit of background on the relevant notions of character and compositionality. We'll also say a bit about why one might doubt that characters compose, properly speaking. Next, in §3, we'll introduce a typed theory of character for a small fragment of English. In §4, we'll show how to make these characters compose and, in §5, we'll compare the resulting theory to an earlier one due to David Braun. §6 will explore what to say about the rules for using sentences containing multiple demonstratives. §7 will show how presuppositions can be naturally incorporated into the theory. And §8 will discuss the nature and epistemology of logical truths. Finally, in §9, we will conclude by reflecting on the notion of partial understanding that our system effectively models.

2 Background

We expect that what we have said so far may elicit a degree of head scratching. Don't we already have any number of relatively successful compositional semantic theories—for example, the one outlined in Heim & Kratzer (1998) and then modified and expanded on in a great deal of subsequent work—which model how complex linguistic rules can be derived from simpler ones? In a word, no. The problem is that, while there *are* any number of compositional semantic theories available, and while many of these are highly successful at explaining a number of interesting features of linguistic meaning, they all effectively assume a level of knowledge of the context that is incompatible with the sort of Stalnakerian picture we just outlined above. More specifically, all of these theories assume that composition takes place *after* reference and other context-sensitive features have been resolved. While that's perfectly fine for certain purposes, this assumption means that theories along these lines offer no obvious answer to the question of how knowing the rules of language allows us to reason about a context on the basis our partial understanding of an utterance, rather than vice-versa.[6]

To illustrate, consider once more an utterance of (1). On any standard compositional semantic theory, in order to determine the meaning of this sentence token, we first need to determine what 'you' refers to. That's not a problem for such theories *per se*, but it is a problem if we want to try to use one of these theories to explain how we are able to engage in the sort of reasoning that Stalnaker pointed to, reasoning with something like a propositional concept rather than a proposition. The problem is that, if we take the model at face-value, then we shouldn't be able to derive a compositional semantic value without first knowing the value of 'you'. In the absence of further explanation, it remains unclear how we could adapt such a system to help us understand the nature of composed linguistic rules—things which, following Stalnaker, we take it that we can derive even in the absence of knowing who a particular token of 'you' was referring to.

[6]Note that if, following Lewis (1980) and many others, we were to distinguish between SEMANTIC VALUE and ASSERTORIC CONTENT, then it might seem that a formal theory of semantic value will be well-suited to model the sorts of semantic rules we have in mind. After all, semantic values are typically understood to be complex functions from some set of parameters to truth-values, something which looks much like a model of a rule of use. That said, such entities are typically presented as theoretical posits which explain the behavior of sentences embedded under modals of various kinds, in addition to tense operators, etc. So, while we are certainly open to the possibility that much of what we do in terms of character here might also be done in terms of compositional semantic values, we will have to leave this question unresolved. There are simply too many options to canvass, and none of these other notions was designed to capture the notion of a linguistic rule in the way that Kaplan's character was. Better, it seems to us, to start with the latter.

As noted above, we'll take a different tack, one to be worked out in terms of Kaplan's notion of character as introduced in his 'Demonstratives' (1989b). Kaplan glossed characters as rules for using expressions and sentences; in other words, 'character' was his term for what we have been calling linguistic rules. On Kaplan's conception of things, these are rules which, for any given expression, take a context and return the content of that expression. In the case of indexicals, terms like 'I' or 'here', we can think of these contents as the referents of token utterances of these expressions. Technically, for Kaplan, contents are functions from possible worlds to extensions, rather than extensions themselves. But, with the indexicals, this function is supposed to return the same object as the extension at each possible world. So it will not do any real harm to think of the rules associated with the indexicals as returning a referent, and Kaplan himself often speaks this way (Ibid., p. 505)

In addition to the notion of character itself, Kaplan also offers a way of modeling these: characters, he suggests, can be modeled as functions from contexts to contents (Ibid., pp. 505–6). Shortly after this, Kaplan specifies that his notion of character is meant to be compositional; the character of any complex expression is a function of the characters of the parts of that complex, in conjunction with their mode of combination (p. 507). Unfortunately, Kaplan offers relatively little by way of detailed discussion of how this composition is supposed to work, how it is that we are supposed to be able to derive complex rules of use from simpler ones—or, in the model, how we are supposed to derive more complex functions from simpler ones.

And indeed, the formal machinery that Kaplan employs can make it look as though characters actually *shouldn't* compose. Partly, this is because philosophers of language nowadays are used to a particular kind of compositional theory: namely, a theory that assigns each atomic term in the language to a type, and then derives the types of more complex expressions via function-application.[7] We'll return to why Kaplan's theory is bound to strike many as non-compositional shortly, but first it will help for us to say something about what compostionality amounts to in general.

Following Pagin & Westerståhl (2011), we take a theory of meaning to be compositional iff, to each grammatical rule which can combine one or more expression into a well-formed clause, there corresponds a function which maps the meanings of each of those expressions to the meaning of the whole clause

[7]Again, see Heim & Kratzer (1998) for a paradigm instance of this approach.

(p. 126).[8] Note that nothing here has been said about what the meaning of a complex, such as a sentence, must be. Standardly, these are taken to be truth-values. But we are interested in a different sort of meaning: rules for using sentences. So we are free to think about the relevant sort of output here as a rule rather than a truth-value; doing so in no way undermines the possibility of our semantic theory being compositional. That will just require that the rules of the whole be derivable from the rules of the parts, via some function corresponding to the relevant mode or modes of combination.

To forestall any possible confusion, it is necessary to stress that compositionality requires only that there be *a* function taking us from the meanings of the parts to the meanings of the whole. In many of the formal semantic theories that philosophers are now often familiar with, the relevant function is simply application. So, we are now used to assigning expressions to types: $\langle e \rangle$, or entities, for referring terms; $\langle t \rangle$, or truth-values, for sentences; $\langle e, t \rangle$, or functions from entities to truth-values, for basic property terms; etc. With a simple subject-predicate sentence, we therefore derive the truth-value by applying a function of type $\langle e, t \rangle$ to an entity—just as compositionality would seem to demand. So, the sentence 'Julia flies' will turn out to be true if the function denoted by 'flies' maps the individual picked out by 'Julia' to *true* and false if it maps her to *false*.

Pretty a picture as this is, as we'll see below, even standard formal semantic systems cannot make do with *only* such simple rules for deriving complex meanings. This will be important to bear in mind, since our proposal for modeling the derivation of complex rules from simpler rules will need to be more complicated than simple function-application as well. It will still be fully compositional, however, and with a close analogue in what we take to be the gold standard of semantic compositionality, Heim & Kratzer (1998).

[8] This is just a plain English gloss on Pagin & Westerståhl (2011)'s condition 'Funct(μ)', which they define formally as follows (p. 126):

> For every rule $\alpha \in \Sigma$ there is a meaning operator r_α such that if $\alpha(u_1, ..., u_n)$ is meaningful, then $\mu(\alpha(u_1, ..., u_n)) = r_\alpha(\mu(u_1), ..., \mu(u_n))$.

Here, Σ is the grammar of the language and μ is the semantic theory for it. Note that, on the assumption that every subexpression of a meaningful expression in the language is also meaningful, then Funct(μ) is equivalent to 'substitutional compositionality', according to which two sentences composed of synonymous subexpressions and combined in the same way will also be synonymous (Ibid., p. 127).

Now let us return to the question of why one might doubt that Kaplanian character is going to turn out to be compositional.[9] If we're used to a system which guarantees compositionality by assigning each category of atomic expressions to a type and then composes meanings via function-application, then it will seem natural to think that indexicals like 'I' or 'here' need to be assigned to contents—or, really, extensions—before they can compose. The extension of an indexical is plausibly of type $\langle e \rangle$. That means that it will combine with simple predicates of type $\langle e, t \rangle$ to output truth-values. In other words, assuming that we assign these expressions to their extensions before trying to derive a complex meaning, everything goes swimmingly.

What though if we try to compose characters? Recall that characters, according to Kaplan, are to be modeled as functions from contexts to contents. Even if we simplify and treat these as function from contexts to extensions, we end up with a problem: the types don't match. Consider that the character of 'I' will be modeled by a function of type $\langle c, e \rangle$. A simple predicate like 'am hungry', on the other hand, will be of type $\langle c, \langle e, t \rangle \rangle$—or a function from contexts to functions from entities to truth-values. When we try to compose via function-application, the system crashes. The predicate needs to be fed a context, but the subject fails to provide one. Instead, it provides a function from contexts to entities. What are we to do?

Kaplan (1989b) himself had a suggestion: we derive complex characters not from simple characters and their mode of combination directly, but from contents which are themselves determined by characters. So first we determine the content of each individual expression. Then we compose these. Finally, to derive the character, we abstract over the complex content that

[9] To be clear, we take it to be a settled matter that, supposing that Kaplanian content is compositional, then Kaplanian character will have to be as well. For a proof, see Westerståhl (2012). This runs contrary to King & Stanley (2005), who accept compositionality of content but allege that composed characters will be only 'trivially' compositional—a notion which they fail to explicitly define (p. 123). We take their claim to be primarily aimed at one particular model of character composition, due to Braun (1994), and we will return to what their worries might have been when considering that view below.

The more interesting questions, to our minds, are whether an informative model of character composition can be offered, and whether that model will need to run via content composition. This first question is not settled by Westerståhl (2012), since his proof doesn't construct an informative model of character composition. The second is particularly salient given recent arguments to the effect that Kaplanian content is not, in fact, compositional (e.g. Rabern 2012, 2013). Since our composed characters do not run via contents, this is a worry we can safely set to the side.

we have derived. In other words, we take the content and transform it—somehow—into a function from contexts, etc. to contents (p. 548).[10]

The problem with this suggestion is that 'somehow'. Consider, for example, two sentences as uttered by David Kaplan:

(2) I love Hawaiian shirts.

(3) David Kaplan loves Hawaiian shirts.

In this context, (2) and (3) express the same contents. That is, both these sentences express functions from worlds to truth-values that map those worlds to *true* if David Kaplan loves Hawaiian shirts in that world and to *false* otherwise. For our purposes, however, that's bad: the problem being that, if we are to derive characters by abstracting over contents, then we would seem forced to predict that (2) and (3) have the same character. But clearly they don't. (2) can be uttered felicitously by any speaker who loves Hawaiian shirts, whereas an utterance of (3) can only be used to express something about David Kaplan's love for Hawaiian shirts.

In other words, to move from a content to a character, we need to know *which* character we are looking for. But content alone cannot tell us this; any number of characters can generate one and the same content relative to some context. There is, in other words, no 'way back' from content to character.[11] One can move from content to an arbitrary function from contexts to that content, but not necessarily to the particular function that one was, intuitively, looking for. Kaplan's suggestion, therefore, would seem to be substantively incomplete. We do think that there is better way of developing this idea, one which is still roughly in line with Kaplan's original suggestion. To see that, we're going to need to start building up a slightly different sort of system.

[10] Here is Kaplan's discussion in full: 'Our notation '$\{\phi\}^{\mathfrak{A}}_{cf}$' for the Content of an expression gives a natural notation for the Character of a notation, namely '$\{\phi\}$' (p. 548). Here, c is the context, f is a variable assignment, and $\mathfrak{A}$ is a structure (or a set of contexts, worlds, individuals, positions, times, and intensions). Essentially, what Kaplan seems to be saying is that character can be derived from content by transforming that content into a function from contexts, variable assignments, and structures to contents.

[11] Just as there is no way back from reference to sense, due to the latter's being finer-grained than the former.

3 A Basic Model of Kaplanian Character

To build our compositional theory of meaning rules, it helps to begin by offering a model of character for a small fragment of English. Since Kaplan was himself primarily interested in indexicals and demonstratives, we'll start with those well-known examples.

Let's model a context C as a 5-tuple consisting of an *agent*, *location*, *time*, *demonstratum*, and *world*, where the agent and demonstratum must be located at the location and time in the relevant world (i.e. our contexts are 'proper' in Kaplan's sense). So $C = \langle a, l, t, d, w \rangle$. Using a_c to represent the agent of a given context we can model the character of the indexical 'I' as follows:

(4) $\{\{\mathbf{I}\}\} = \lambda c.a_c$[12]

In other words, we represent the character of 'I' as a function from a context to the agent of the context (this is what the lambda indicates). This reflects the naive view that 'I' refers to whoever is speaking or otherwise designated as the agent.[13]

Similarly, we can represent the characters of 'here', 'now', and 'that' as:

(5) $\{\{\mathbf{here}\}\} = \lambda c.a_c$

(6) $\{\{\mathbf{now}\}\} = \lambda c.t_c$

(7) $\{\{\mathbf{that}\}\} = \lambda c.d_c$

Note that, so far, we are treating these all as functions from contexts to extensions. That's not quite what we wanted, however, since on Kaplan's theory these should really be functions to contents.

Here is a more accurate way of modeling Kaplan's proposed characters for the indexicals:

(8) $\{\{\mathbf{I}\}\} = \lambda c.\lambda i.a_c$

(9) $\{\{\mathbf{here}\}\} = \lambda c.\lambda i.l_c$

(10) $\{\{\mathbf{now}\}\} = \lambda c.\lambda i.t_c$

(11) $\{\{\mathbf{that}\}\} = \lambda c.\lambda i.d_c$

---

[12] We adopt the double-set bracket notation for characters from Glick (2017). Kaplan himself used a single-set bracket to indicate chacacter.

[13] For a summary of reasons to give up this naive view, see Cohen & Michaelson (2013). Since these issues aren't germane here, we'll ignore them.

We use i to represent the circumstance of evaluation, which we will also refer to as the 'index'. These indexes won't really do anything with respect to the indexicals, but they can with respect to other terms when embedded under operators like modals. So it's good to have them around, and to have the characters of indexicals regimented in the official, proper way.

Note that there are plausibly characters which are insensitive to both context *and* index, not just to the index. For example, on the assumption that proper names function like constants, the character of a name will look roughly like this:

(12) $\{\{\textbf{Julia}\}\} = \lambda c.\lambda i.Julia$

Here, '*Julia*' is being used as a meta-language name for Julia.[14]

In their predicative use, property terms will be sensitive to the index in the sense that they will tell us that, relative to some features of the world, a certain set of properties should be mapped to either *true* or *false*. So, for example:

(13) $\{\{\textbf{is tall}\}\} = \lambda c.\lambda i.\lambda x.\ height_i(x) \in TALL_c$

(14) $\{\{\textbf{is red}\}\} = \lambda c.\lambda i.\lambda x.\ color_i(x) \in RED_c$

As is fairly standard, we take the predicate 'is tall' to be true of an object just in case the height of that object is greater than some contextually set standard.[15] Instead of putting the degree explicitly in the semantics, however, we instead appeal to the set of heights greater than that degree in context, i.e. $TALL_c$.[16] An object is tall in a context if it is a member of this set.

Finally, we offer a semantics for adjectives like 'very', which should shift the contextually-relevant standard whereby to a stricter version of it, and for the sentential connective 'and', something which Kaplan failed to define:

(15) $\{\{\textbf{very}\}\} = \lambda c.\lambda i.\lambda P.stricter(P)$

(16) $\{\{\textbf{and}\}\} = \lambda c.\lambda i.\lambda S_1.\lambda S_2.[\![S_1]\!]^{c,i}\ \&\ [\![S_2]\!]^{c,i}$

[14] Of course, it might be that names can be bound, in which case this won't be the right analysis. See, inter alia, Cumming (2008). We do not mean to be taking on any real commitments here; the claim is merely for purposes of illustration.

[15] See Kennedy (2007) and Kennedy & McNally (2010). Note that we've also indexed height and color. Intuitively, this is so that the modals can shift us to worlds where an object is taller or redder than it actually is.

[16] While this choice will obscure certain interesting features of degrees, it will make things simpler for our purposes by allowing us to model all predication via set membership.

'*stricter*' is a meta-language name for a function that maps properties onto restricted versions of those properties. Applied to $TALL_c$, this function will return a subset—intuitively the taller objects. Applied to RED_c, it will return a range of colors closer to prototypical red.[17] 'And', on the other hand, takes a context, index, and two sentences and returns the classical conjunction of those sentences. In other words, 'and' corresponds to a function which maps to the sentence to *true* if both conjuncts are mapped to *true* relative to the relevant context and index, and to *false* otherwise.

We've now given a formal model of character—or context-invariant meaning rules—for a small fragment of English. We won't deal explicitly with modals, though it's fairly easy to see that these should be functions from the contents of sentences, i.e. functions from indexes to truth-values, to contents of sentences. The important thing is that modals won't care about the default index, the one with the actual world, time, location, etc., in determining truth. And, although Kaplan himself barred them from his semantics, it's fairly easy to see how to represent the character of MONSTERS, or operators that shift the context, as well: supposing that these are sentence-level operators, then they will be functions from functions from contexts to contents to functions from contexts to contents. Analogously to the modals, monsters won't care about the actual context in order to determine truth or falsity.

Since allowing monsters into our language will complicate with the relevant notion of logicality that we'll seek to explicate in §7, we'll follow Kaplan in barring them from our language. We'll return to this issue below.

4 Composing Characters

Let's look in a bit more detail at an example:

(17) I am tall

We want to spell out a function on characters such that $\{\{\textbf{is tall}\}\}$ composes with $\{\{\textbf{I}\}\}$ to generate $\{\{\textbf{I am tall}\}\}$.[18]

[17] We could let 'stricter' itself pick out a function which is sensitive to context, but we'll simplify by not doing so.

[18] Returning to the technical discussion from Westerståhl (2012) referenced above (fn. 8), we want to demonstrate that there's a function satisfying the condition Funct(μ). This says, roughly, that there's a function assigning characters to complex expressions given the character of the parts and their syntactic mode of composition. We then want to show that there's a tractable procedure for generating the values of the function μ.

As we said above, $\{\{\mathbf{I}\}\}$ is $\lambda c.\lambda i.a_c$, a function from contexts to functions from indices to the agent of the context. $\{\{\textbf{is tall}\}\}$ is $\lambda c.\lambda i.\lambda x.\ \mathit{height}_i(x) \in \mathit{TALL}_c$. Intuitively, we want $\{\{\textbf{I am tall}\}\}$ to be $\lambda c.\lambda i.\ \mathit{height}_i(a_c) \in \mathit{TALL}_c$—in other words, we want $\{\{\textbf{I am tall}\}\}$ to be a function from contexts and indices to truth-values. In essence, we want to compose underneath the context and index parameters.

This observation is already enough, really, to see that there *is* such a function. After all, it's clear what the value of composing the subsentential expressions of 'I am tall' should be: $\lambda c.\lambda i.\ \mathit{height}_1(a_c) \in \mathit{TALL}_c$. And it's not at all difficult to see that we could figure out what the right result should be from a couple of leading examples. That's enough to justify the claim that there will be such a function. But this stops short of being a systemtic compositional procedure which we can carry out to get the right result in any arbitrary case. That is what we really need if we are to model how our grasp of the rules of a language is able to bear its informational fruit.

Direct composition along the lines of a simple function-application account is ruled out due to the type-mismatch problem outlined in §2. Again, what we really want to do is compose *under* context and index—something which is not allowed by simple function-application. But there is another way available, one that will get us the goods we want. On that procedure, rather than composing via function-application, we just directly define the output we want. In the language of lambda abstraction that we've been using to this point, here is how this procedure will look:

> CHARACTER FUNCTIONAL APPLICATION (CFA):
> if α is a branching node and $\{\gamma, \beta\}$ the set of its daughters, then, for any context c and index i, if $\lambda c.\lambda i.\beta(c)(i)$ is a function whose domain contains $\lambda c.\lambda i.\gamma(c)(i)$, then $\{\{\alpha\}\} = \lambda c.\lambda i.\{\{\beta\}\}^{c,i}(\{\{\gamma\}\}^{c,i})$.

The important thing to notice here is that this procedure doesn't raise any issues for compositionality, since the meaning of the composed character depends on the meaning of the component parts and the corresponding syntactic structure—and those alone. While context and index are mentioned in the course of the procedure, they occur always as variables. So one need know nothing about the actual context or index to apply CFA.

Let's work through an example. To generate the character of 'I am very tall', we start by calculating $\{\{\textbf{is very tall}\}\}$ from $\{\{\textbf{very}\}\}$ and $\{\{\textbf{is tall}\}\}$. Since $\{\{\textbf{is tall}\}\}$, or $\lambda c.\lambda i.\lambda x.\mathit{height}_i(x) \in \mathit{TALL}_c\ (c)(i)$, is of type $\langle e, t\rangle$ and {{**very**}}, or $\lambda c.\lambda i.\lambda.P.stricter(P)(c)(i)$, is of type $\langle\langle e,t\rangle,\langle e,t\rangle\rangle$, we have by CFA that {{**is very tall**}} is:

(18) $\lambda c.\lambda i.\lambda x.stricter(height_i(x) \in TALL_c)$

Since $\lambda c.\lambda i\lambda x.stricter(height_i(x) \in TALL_c)(c)(i)$ is a function of type $\langle e,t\rangle$ and $\lambda c.\lambda i.a_c(c)(i)$ is of type $\langle e\rangle$, we have by CFA that {{**I am very tall**}} is:

(19) $\lambda c.\lambda i\lambda x.stricter(height_i(a_c) \in TALL_c)$

Again, we've just repeatedly composed under context and index, harmonizing the context and index variables so that they're constant between the composed expressions. Nothing strange here.

One way of thinking about what's going on here is to conceive of this procedure as amounting to: first 'saturate' each character with a 'dummy' index and context. Then compose. Then abstract over these dummies, yielding a character once more.

Now, this step-wise procedure might seem to fall prey to the worries we raised for Kaplan's own suggestion that we derive composed character from composed contents—so it is crucial to see why this is not the case. The difference is that Kaplan was proposing to abstract over contents themselves. Our suggestion is to abstract over *names of contents*, or actually names of extensions (since those are the parts which can compose directly, via function-application). While seemingly minor, this difference matters here, since we can have more than one name for the same entity. If the naming scheme we're using systematically preserves certain information (as the lambda calculus does with derivational history), then we *can* hope to recover character from extension—via that additional information (i.e. derivational history). It may well be that this is what Kaplan had in mind in 'Demonstratives'; his discussion is simply too terse for us to be able to know for sure (see fn. 10).

This step-wise interpretation of CFA isn't mandated, however. Another way of interpreting this compositional procedure focuses instead on the role of the linguistic rules being composed. So, for instance, the rule for 'very' tells us to move from a predicate to a stricter version of that predicate. Combined with 'tall', we move to a stricter version of tallness—one which admits fewer objects relative to a context and world. The rule for 'I' tells us to find the speaker of the context. So, combining these, we derive a rule that tells us to look for the speaker of the context and see if she is amongst the very tall things relative to that context and world. If yes, then the utterance is true. This is just the rule that the repeated application of CFA yields.

In a sense, what we're doing is analogous to how semanticists like Heim & Kratzer (1998) treat intensional operators (just in a sense though; don't overread this). They say that intensional composition works like so:

> Intensional Functional Application (IFA):
> if α is a branching node and $\{\gamma, \beta\}$ the set of its daughters, then, for any index i, if $[[\beta]]^i$ is a function whose domain contains $\lambda i'.[[\gamma]]^{i'}$, then $[[\alpha]]^i = [[\beta]]^i(\lambda i'.[[\gamma]]^{i'})$.[19]

Roughly, expressions ordinarily contribute compositional semantic values which are *extensions*, but sometimes a different form of meaning—*contents* in our terms—are contributed. When we have an expression like a propositional attitude report or a construction like 'in the fiction of', what follows it syntactically is an expression with an extensional meaning. But the report or construction denotes a function which is looking for a *content*, not an extension, leading to a type mismatch. We rectify this by appealing to IFA instead of standard function application in these instances. Still, IFA doesn't violate compositionality; the meaning of the composite remains a function of the meaning of the parts, even with the type shift. What differs is *which* function determines the composite meaning.

The important difference to note between CFA and IFA is that CFA is uniform across character composition, whereas IFA only applies when a certain sort of mismatch arises in an otherwise simpler story of composition. Whether an analogue of IFA will also prove necessary at the level of character composition is not something that we will address here, since this will depend on how we define the characters of the modals and related terms, and we hope to avoid making any explicit commitments on that front. That said given that CFA already takes account of contexts and indexes, we can see no reason why an elegant model with no need for an analogue of IFA shouldn't be available. Offering that model will have to wait for another occasion, however.

So much for our initial shot at a compositional theory of character. To further clarify the view, it will help to compare it to Braun (1994)'s earlier theory of composed character. Then we will offer several refinements.

[19]Cf. Heim & Kratzer (1998, p. 308). Note that, for Heim and Kratzer, these semantic values are considered relative to a world and assignment function, rather than an index. We have modified IFA to keep it in line with terminology we've been using here.

5 Comparison to Braun's View

Like ourselves, Braun (1994) sets out to offer a compositional theory of character. Braun's theory looks rather different from our own, however, in no small part because he is working in a structured propositions framework whereas we have been working in the sort of lambda notation typical of most recent formal semantic theories. At the level of content, this difference is more apparent than real. Consider, for example, how Braun would represent the content of our (17), or 'I am tall,' as uttered by Nat:

(20) ⟨*Nat*, *being tall*⟩

The proposition represented by this 2-tuple is then stipulated to be true if the property of being tall applies to Nat (in the world of utterance). So something like function application (or, really, predication) is still here; it's just that the formalism doesn't wear this on its sleeve.

When representing complex characters, Braun represents them as:

(21) ⟨{{**I**}},{{**being tall**}}⟩

Intuitively, Braun's proposal is to piggyback our model of structured character on the existing model of structured content. Instead of putting properties and objects into our ordered tuples, now however, we insert the characters of individual expressions. As Braun explains, what something like (21) represents is an entity which, once we saturate each character with a context, 'determines' a structured content (Ibid., p. 197). So, effectively, what (21) amounts to is a set of instructions for moving from a context to a content.

Several problems arise at this point. First, it is not at all clear that the formalism Braun relies on is actually well-defined. The problem stems from, of all things, the comma in (21). How are we to interpret this? Is it like the comma in (20), such that the 2-tuple is true if whatever is in the first slot bears whatever property is in the second slot? That seems wrong, since then (21) would seem to be either true, false, or undefined—and, plausibly, rules of use are none of those things. Alternatively, we might think of the comma in (21) as representing a function from contexts to the sort or predication represented by the comma in (20).[20] But if this is what Braun has in mind, he offers no indication of this in the text.

While this might seem pedantic, it's driven by a genuine concern. What we wanted was a model of, or way of formally representing, a *rule of use*

[20]In that case, perhaps it would be better transcribed as {{,}}.

or *meaning rule* for the sentence 'I am tall.' Instead, what we've gotten is a formal structure which we are told represents this rule, but which itself isn't fully defined. At best, this formal structure seems to tell us that the content of 'I am tall' is composed by, first, finding the content of 'I' in the context, and then composing that with the content of 'am tall' in the context. But if something like Kaplan's theory of character and content is correct, then this just looks to be a restatement of the obvious.[21]

Second, we began (following Stalnaker) with the hope that rules of use might help us to understand the sorts of information one can glean from the utterance of a sentence even in the absence of any real knowledge of the context. We wanted a way of systematically deriving, for instance, the fact that an utterance of (17) can be used to assert that the speaker is tall. Our preferred model of character wears this information on its sleeve, by representing the character of 'I am tall' as a function from contexts and indexes to either truth or falsity, depending on whether or not the speaker in the relevant context is tall. Braun's model, in contrast, represents such information only opaquely. One can infer this information from Braun's representation of the complex character in conjunction with (i) an understanding of what that character says about the associated content and (ii) an understanding of what contents themselves represent. This hardly strikes us as the most natural and informative way of putting things.

Third, and probably most importantly: our model of character naturally predicts that certain expressions don't fit together in the right way to form a complex character. So, for example, one cannot compose 'Nat' and 'Julia' according to either CFA or CFA*, since the extension of one is not a function whose domain contains the extension of the other. As far as we can tell, however, nothing in Braun's theory prevents us from generating the complex character:

(22) $\langle\{\{\mathbf{Nat}\}\},\{\{\mathbf{Julia}\}\}\rangle$

Granted, there is either no structured content corresponding to (22), or else that structured content is not well-formed. But for all that Braun tells us about composing characters, (22) is perfectly well-formed. We submit there is

[21]Perhaps this is what King & Stanley (2005) had in mind when they claimed that compositional theories of character would only ever be 'trivially' compositional. As noted in fn. 9, they do seem have had Braun's theory primarily in mind when making this claim. On the other hand, they also never define what they mean by 'trivial'. So it is difficult to be sure whether we are simply agreeing with them here.

no rule of use for using the string ‘Nat Julia’ (importantly, this is distinct from the well-formed string ‘Nat, Julia’), so we take this to be a bad prediction. In fact, we take it that one of the main reasons one might want a compositional theory of meaning of any sort is to rule out as ill-formed or meaningless cases like this one. So we take it that our theory, with its underlying type constraints, represents a significant advance over Braun’s proposal.

6 The Problem of Multiple Demonstratives

Following Kaplan, we started with indexicals like ‘I’ and ‘now’ as our core cases. So we offered a model of the rules of meaning that maps each instance of the same term, as they occur within one and the same sentence, to the same content. That’s fine for terms like these, so long as we keep away from oddball cases involving multiple speakers or extremely long utterances.

But such idealization seems much less compelling with respect to the demonstratives ‘this’ and ‘that’, which commonly occur in the course of one and the same sentence while referring to different things. Consider, for instance, an utterance of (23):

(23) I’ll take that, that, and that.

When one utters (23) in a La Bagel Delight, one usually isn’t asking for the same bagel thrice-over. Perhaps even more strangely, since our above model only contains one demonstratum d, it would seem to imply that there is a substantive sense in which the meaning rules for ‘this’ and ‘that’ are identical: both are used to refer to the demonstratum of the context. But that hardly seems right. Something, then, needs to change.

One option would be to follow a different line of Braun’s thought, developed in his (1996). In light of concerns like these, Braun suggests that demonstratives should be treated as having three levels of meaning, not just two. At the level of linguistic meaning, demonstratives like ‘that’ are posited to refer to whatever is demonstrated by an accompanying demonstration. In effect, Braun posits that terms like ‘this’ and ‘that’ are *incomplete*; they must be supplemented by a demonstration in order to have a character. That character, in turn, determines their extension, relative to a context (Ibid., pp. 155–56). So the linguistic meaning of a demonstrative tells us to look for an accompanying demonstration, and then a character is assigned, in context, to the pair of the demonstrative and that accompanying demonstration.

We find Braun's suggestion unsatisfactory for two main reasons: first, it remains the case that, at the level of linguistic meaning, 'this' and 'that' mean the same thing. Both simply refer to whatever the speaker happens to be demonstrating at the moment when they are used. Second, it is simply not the case that every felicitous use of a demonstrative to refer is accompanied by a demonstration or gesture.

Of course, one might try adjusting the notion of a demonstration so as to allow that something like a referential intention—even one not made manifest by an accompanying gesture—can also play this role. This is basically what Kaplan suggested in 'Afterthoughts' (1989a). This strategy leaves our first concern unresolved, however. What's more, it would also mean that characters and contents are assigned to pairs of expressions and referential intentions, rather than just to expressions. And while it might seem plausible enough that gestures should count as a part of the lexicon, it seems far less plausible that aspects of the speaker's mental state should be counted as such. One might try to finesse the issue by replacing the referential intentions with subscripts, but that looks like a formal fudge rather than any sort of explanation of what sort of meaning rules govern the use of demonstratives.

We prefer a different way out of this quandary: recall that characters were supposed to be rules for using an expression. Our proposal is to hold to this understanding of what we are trying to model and jettison what we take to be a bad model. This stands in contrast to Braun's strategy of holding fixed the model and adjusting how it is to be interpreted. In a word, we think that Kaplan went wrong by starting with the indexicals—terms which plausibly *do* pick out a single object in a context, more or less on their own—and then taking the sorts of rules which work for these to be representative of what meaning rules are like in general. In contrast, we propose to begin with what we take to be a plausible gloss on the rules of use for the simple demonstratives 'this' and 'that': 'that' can be used to refer to any distal object in the context; 'this', its dual, can be used to refer to any proximal object.[22] How might we formally represent such a rule? One salient option would be for 'this' and 'that' to be mapped to sets of objects as their extensions, rather than to any single object. Intuitively, this will be the set of objects the speaker could be referring to with the use of the term.

Representing this in our formalism can be done in the following manner. Let us say that demonstratives are associated with 'fuzzy characters', like the

[22] See Elbourne (2008) for a theory of contents roughly along these lines.

fuzzy credences we find in formal epistemology.[23] First, we need to modify our model of the context so that, instead of a demonstratum d, it instead contains a set of distal objects, *Dist*, and a set of proximal objects, *Prox*. Now we modify our earlier (11) as follows:

(24) $\{\{\textbf{that}\}\} = \{\lambda c.\lambda i.f(Dist_c) \mid f$ is a choice function on sets of individuals$\}$

A 'choice function' is just a function from a set to a member of that set. So each f takes us from the set of distal objects to a particular distal object. This yields a character of 'that' is now fuzzy in the sense that it picks out no single function from context to object, but rather a set of functions from a context to each distal object in that context. Intuitively, what this set represents is the options the speaker has when uttering 'that' in a context, the possible referents she could be picking out.

What remains is to offer composition rules for fuzzy characters like these. These will run more or less as one would expect. When one hits a fuzzy character, one cannot simply compose under the context and index as we did above. The problem is that now we are dealing with a *set* of underlying semantic values, a set of $\langle e\rangle$'s or $\langle e,t\rangle$'s rather than a single object or function of that type. What we need is a way to compose with sets.

We can derive such a way by stipulating that, when one is composing with a fuzzy element, one composes only at the non-fuzzy level. It will make things slightly simpler at this point to treat every character as fuzzy in the sense of being a set of functions, even if only a singleton set. Effectively, we do this by type-raising every character, treating a character of type T as being now of type $\langle T,t\rangle$. We'll still treat these as sets for readability. Now we can introduce a slightly more complex version of our earlier CFA:

Character Functional Application* (CFA*):
if α is a branching node and $\{\Gamma,\Delta\}$ the set of its daughters, then, for any context c and index i, if for each $\delta \in \{\{\Delta\}\}$, $\lambda c.\lambda i.\delta(c)(i)$ is a function whose domain contains $\lambda c.\lambda i.\gamma(c)(i)$ for each $\gamma \in \{\{\Gamma\}\}$, then $\{\{\alpha\}\} = \{\lambda c.\lambda i.\{\{\delta\}\}^{c,i}(\{\{\gamma\}\}^{c,i}) \mid \delta \in \{\{\Delta\}\}$ and $\gamma \in \{\{\Gamma\}\}\}$.[24]

[23] Cf. Levi (1980) and van Frasssen (1990). Thanks to Julien Dutant for discussion here.

[24] Note that this way of doing things returns undefined if there's a type mismatch between any members of Γ and Δ, thus recording the idea that no expression receives distinct characters whose specification under context and index are of distinct types.

In plain English, CFA* just says that fuzzy-ness bleeds out. As soon as one composes a fuzzy character with anything, the result is a fuzzy character. When one composes a fuzzy character with a non-fuzzy one, one ends up with as many possible ways the sentence might be used as the number of possibilities that fuzzy character represented. When two fuzzy characters interact, one ends up with the same number of possibilities as the cross-product of the two sets.

Is this what we should hope for in a theory of complex characters? We submit that the answer is 'Yes.' Consider again our (23). In uttering this in La Bagel Delight, there are as many combinations of things you could be asking for as there are combinations of three distal objects in the context. Granted, the listener is very likely to assume that the speaker is referring to three different things—and that all three are bagels—but that plausibly looks to be a matter of pragmatic reasoning rather than something baked into the rules for using demonstratives. Referring to the same bagel three times over doesn't seem impossible, or in violation of one or another rule of language; rather, it's just weird. Since that's compatible with everything we've said here, we take our theory to have made some reasonable predictions.

In contrast with Braun then, our proposal in terms of fuzzy characters allows us to retain just two levels of meaning while offering an informative model of the linguistic rules for demonstratives. That model allows for a straightforward characterization of the difference in meaning between 'this' and 'that', and is in no way dependent on the highly dubitable claim that any felicitous use of a demonstrative must be accompanied by a demonstration. Granted, the rules of use to which we appeal look rather different from those which Kaplan considered. Kaplan's paradigm semantic rules mapped us from a context to a single content. We are more willing to think that the rules governing certain terms map out a range of ways those terms can be used in a context—leaving it up to the speaker (or perhaps something else) to decide to which use a term has actually been put in a given context. None of this, as we just above, undermines the compositionality of these rules.[25]

[25]Note that, having thus modified our system, it now bears a certain formal similarity to the one sketched in Harris (2019). In particular, both systems maintain compositionality via systematic type shifting—though each implements this type shifting in a very different way. We eschew a full comparison here for two main reasons: first, Harris is clear that he is offering a model of semantic content, not character; and, second, Harris' model is intended to play a specific theoretical role in a 'pragmatist' approach to meaning, like Relevance

7 Presuppositions Integrated

Now let us consider a topic which has gone largely overlooked in the literature on Kaplanian character: presupposition. Again, a comparison to Braun's system will prove instructive.

Start by considering the sentence:

(25) I stopped smoking.

We can recover quite a bit of information from an assertive utterance of (25) without knowing anything about who 'I' refers to. In particular, we can recover that, whoever the speaker is, they used to smoke. But how can we recover this, in any principled way, from a representation of the character of 'I stopped smoking' as ⟨{{**I**}},{{**stopped smoking**}}⟩? This representation tells us nothing about how the presupposition carried by the predicate of the target sentence ends up being a presupposition of the whole sentence.

Of course, there's a natural addendum we could try: suppose the presuppositions of any part of a complex character 'project' out to the whole. This will explain how the relevant presupposition gets associated with (25). The trouble with this suggestion is two-fold, however.

First, it looks unmotivated on Braun's way of thinking about things. Complex characters are effectively a set of instructions for deriving sentential contents: first, we derive the contents of each individual expression, then we put these together. There would seem to be little to no room here for some element of the character to somehow 'pop out' and constrain the circumstances in which the sentence can be appropriately used, the circumstances in which it will turn out to have a well-formed content.

Second, and perhaps more importantly, this stipulated projection simply isn't accurate. Not every subsentential presupposition projects in this way. Famously, certain constructions like attitude verbs will 'block' presuppositions. So, whereas an utterance of (25) plausibly presupposed that the speaker smokes, an utterance of (26) almost certainly does not:

(26) Nat believes that I stopped smoking.

Theory. These divergent explanatory aims make direct comparison difficult, and essentially impossible without an exceedingly long digression.

CFA* also bears some formal similarity to work in formal linguistics on the composition of ambiguous sentences and question/answer structure. Again, we eschew a full comparison as the explanatory aims of these projects are very different. Still, the formal similarity is interesting. Compare, for instance, Poesio (1994) and Novel & Romero (2010).

What Braun would actually need to posit is not just that presuppositions project out, but also that certain other elements of a sentence can (somehow) block that projection. The problem is that there are no obvious tools in Braun's system that look fit to help us explain this blocking. He might try adding some, of course, but we cannot see any natural way to do so in a manner that strikes us as either informative or explanatory.

Can our system do better? Yes it can. We start by treating presuppositions as partial functions, as we take to be more or less standard in formal semantics.[26] So {{**stopped smoking**}} will be something like: $\lambda c.\lambda i.\lambda x.x$ *used to smoke at* w_i: *x stopped smoking at* w_i. This function is partial in the sense that for values of x which never smoked, the function is undefined. In other words, we take the rule for using 'stopped smoking' to be something like: to be applied truthfully to individuals who used to smoke and have stopped and falsely to individuals who used to smoke but have not stopped. This rule says nothing at all about individuals who never smoked.

With this character in hand, we can now appeal to CFA to derive the complex character of 'I stopped smoking'. First, we check to see whether the domain of $\lambda c.\lambda i.\lambda x$ *x used to smoke at* w_i: *x stopped smoking at* w_i contains the value of 'I' relative to any context and index. It does, though it presumably fails to map any number of these to an output. But, since the function is partial, that's just what we wanted! Now we compose as above, yielding:

(27) $\lambda c.\lambda i.a_c$ *used to smoke at* w_i: a_c *stopped smoking at* w_i

What this is is a partial function from contexts and indexes to truth-values where the truth-value is only defined in the case where the speaker at the context used to smoke at the world of the index. This would seem to be exactly what we wanted in order to represent the rule for using (25): it can be used to assert something true in cases where the speaker used to smoke and has stopped, something false in cases where the speaker used to smoke but has not stopped, and nothing at all when the speaker never smoked.

Again, one could equally well think of building up this complex character in a stepwise manner, by instantiating 'I' with a dummy value which we *assume* used to smoke. We then compose, since all functions are then defined, and 'undummy' the value for 'I', releasing as well the assumption that 'I' used to smoke. The result will be the same as above.

[26]Cf. Elbourne (2005) and Schoubye (2013).

What we've done here, as in the last few sections, isn't at all radical. We've relied on a general picture of how to compose rules, familiar from logic, to generate complex rules from simpler ones. In this particular case, this picture offers us a neat explanation of how presuppositions are carried from subsentential expressions to the whole: when we want to figure out the partial function corresponding to 'I stopped smoking', we start by looking at the defined values of 'stopped smoking' and work backwards. Just as when we want to figure out what the result of squaring the output of some partial function on the reals would be, we look at a generic defined value, square it, then release the presumption that it's defined when we abstract back out.

We could model this more syntactically, but there's no need given our aims here. The point we want to stress is that, once we remind ourselves that characters are rules of use and complex characters correspondingly more complex rules of use, an *explanatory* account of why whole sentences carry the presuppositions of their parts (at least in the cases where they do) falls out almost immediately: complex characters aren't just instructions for how to compose complex contents. Rather, they are instructions for *using* a sentence in a context—instructions which necessarily reflect the rules for using each part. Many such rules will have effects on the truth-conditions of the sentences to which they contribute, but some rules will have effects ranging beyond what is reflected in the strict truth-conditions.

Admittedly, presupposition is a complex phenomenon and there will be plenty of cases, like (26), where the presupposition fails to project out as we might have initially expected. While we will refrain from offering any specific theory of how such blocking works, it is not hard to see what, intuitively, is going on. Certain elements of the lexicon, e.g. attitude verbs, come with rules of use that indicate that whatever presuppositional content is embedded under them should be discarded. Formally, this might be modeled by introducing an operator (say, $\varnothing$) which maps any restriction on assignments of values to a trivial restriction (i.e., $x = x$). Hard questions remain regarding whether all presuppositions are alike in how they interact with blockers like attitude verbs, let alone whether all attitude verbs are alike in how they block presuppositions.[27] We will not attempt to address these questions here, as our aim is merely to provide a framework within which the presuppositional effects on character, not just content, can be productively explored.

[27] See Beaver & Geurts (2014) for helpful a helpful overview of these issues.

8 Logicality and Character

A final virtue of our approach is that it can help us to better understand the connection between logical truths, analyticity, and necessity. Getting this connection clearly in view will require a slight detour, so bear with us.

From Kaplan himself (1989b) and Zalta (1988), we know that some 'logical truths' in languages containing indexicals—'I am here now,' for instance—aren't necessary. This observation effectively severs one of the time-honored connections between logicality, necessity, and analyticity.[28] As with many time-honored views, however, we take there to be reason to suspect that there was *something* interesting and worthwhile that these theorists were aiming to capture.

It's relatively easy to articulate this something. If a sentence is a logical truth, even in a language with indexicals, then the rules for the logical expressions occurring in it suffice to guarantee that it's true. Knowing the rules for those expressions should then suffice to show that the logical truth is true. So it should be analytic (true in virtue of meaning). But as analytic truths are supposed to be necessary, at least in myth and legend, and 'I am here now' clearly isn't, something must have gone wrong. The connection between logical truth and analyticity just sketched is compelling, so it would be ideal if we could show that there was at least a sense in which it was right.

This can, in fact, be done. Woods (2016) shows that we can use standard mathematical criteria for logicality—invariance of meaning under transformations of the domain—to characterize two types of logicality. One type, applying to contents, captures how the meaning of a logical expression is immune to change in index. We can use this property to explain how logical truths are true no matter where they're assessed (that is, at every index, holding the context fixed). The other type, applying to characters, captures how the meaning of a logical expression is immune to change in context. We can use this property to show how logical truths are true no matter where they're uttered (that is, at at every context of utterance).

Logical truths in the fullest sense are invariant in both ways; they display no sensitivity at all to empirical facts encoded in either the context or index. There are also non-trivial members of each class of lesser logical truths. For instance, let's stipulate that σ is an operator identical to $\neg$ if the context of utterance contains no wombats, and $\neg\neg$ otherwise. Then $\sigma(\forall x(x = x))$

[28]Though see Hanson (2006) for a way to repair the gap.

has a non-constant character. Which function from index to truth-value it picks out depends on the context. But the particular function it picks out in the actual context is a constant function from any index to the true or a constant function from any index to the false. Its content, but not its character, is therefore invariant.

'I am here now,' on the other hand, picks out the same function in every context: the function from an agent of the context to the true. But that particular function isn't a constant function from index to truth-value, since, for some indices, it's simply not true that I'm here now. It's this kind of content-invariance that gives rise to the gap between analyticity and logical truth mentioned above.

These different forms of logicality track different types of sensitivities to the world, and both are well worth preserving. Additionally, each type of logicality connects up with a recognized form of necessity—truth at any index, on the one hand, or truth at any context of utterance, on the other.[29] Character invariance, in particular, neatly tracks our ability to 'suss out' the logical truthiness of a logical truth given a grasp of its meaning, as it tracks our ability to see that the meaning of what we've said makes it true (unlike mere content invariance where facts about where we're speaking may well play a role in making what we say true). That is, it tracks our ability to *know* that sentences are analytic in virtue of knowing its meaning.

What's the connection of this material to our topic? Well, in order to claim that 'I'm here now' can be seen to be true on the basis of its meaning alone, we've got to be in a position to *calculate* that meaning. Things which are logically true in virtue of meaning aren't sensitive to any particular context, so we need a context-free way of calculating the meaning of 'I'm here now.' That is, we need to be able to recognize that what we're saying is true at the context of utterance in virtue of meaning, even when we don't know which index is relevant or which context we're in (cf. §2).

What we need, in other words, is a method for calculating {{**I am here now**}}—the relevant semantic rule—on the basis of the atomic characters which make it up. And, in fact, we've provided just that above. A method of generating compositional characters thus helps us to see why (character-invariant) logical truths are analytic, once we've properly under-

[29]Woods also shows that composition of content out of character invariant materials where context makes no difference implies necessity in the second sense—i.e. that the natural account of character composition satisfies Funct(μ) and satisfaction of Funct(μ) together with character invariance implies truth at the actual world of any context.

stood the relevant notion of logical truth. If we couldn't grasp complex characters directly, we'd be hard pressed to explain how we can know that 'I'm here now' is true any time it is uttered, given only our grasp of the characters of 'I', 'here', and 'now'. After all, we might not know what their content actually is in any given instance, let alone in any arbitrary one.

Perhaps one will be tempted to respond that all we really need to know is that utterances of sentences like these are true no matter what their content is. We agree. As pointed out repeatedly above, saying this is just to say that we can grasp a complex semantic rule on the basis of simpler semantic rules. Which context we 'dummy' with in calculating the complex rule doesn't matter. But that's just the formal analogue of the claim that we can grasp complex characters. So our account of complex characters props up the connection between logical truths and analyticity, salvaging the important aspect of the traditional connection.

Note that our refined treatment of demonstratives yields an additional nice result here. Typically, in a Kaplanian framework, 'This is that' will come out as a logical truth in the same way that 'I am here now' does. This is because, on the standard Kaplanian theory, a given context contains only a single demonstratum. Braun's theory avoids this result by positing a necessary connection between demonstrative reference and a physical demonstration, which has the effect of making 'This is that' non-logical by fiat; the meaning of 'This is that', in the relevant sense, now depends on worldly facts involving this demonstration. Our view, in contrast, treats 'This is that' as a non-logical truth in exactly the right way. It's no part of the character that we've assigned to 'this' and 'that' that any choices of distal and proxmial demonstratum will be the same, so its character is variant, so it's not a logical truth. This will be true even when, as a matter of fact, this and that happen to necessarily (in a context) pick out only single entity—the character will remain variant, even if the content is constant. So it's no logical truth. This, we take it, is a good result.

We've banned so-called monsters—context shifting operators—from our semantics, following Kaplan. This restriction simplifies matters considerably, as adding a sense of logicality which allows operators that shift context is a bit more complicated.[30] The relevant comparison here is with modalized bits

[30]There is also the question of how to generate a logic that can account for non-operational context-shifts—that is, ordinary shifts in context over time. We will refrain from making any commitments here, but see Radulescu (2012, 2015) for such a system.

of language such as languages containing modals—operators which shift the index. Developing a notion of logicality for alethic modals, requires specifying which bits of index structure, such as accessibility relations between worlds, are preserved when testing for sensitivity (formally, how much structure we build into the transformations of the set of indices we treat as representing different index possibilities). This issue is vexed and somewhat undecided even for bits of language like alethic modals—that is, languages we've got a fairly good grasp of.[31] In the case of monsters, we've got much less of a grasp of what bits of contextual structure would need to be preserved in order to get a reasonable notion of logicality for monsters up and running.

Nevertheless, there's every reason to think that we could extend Woods' account to accommodate logicality for monsters. Once we have an appropriate account of the lexical meaning of whichever monsters there are, we can adapt the treatment of logicality for modal expressions, such as the use of invariance under bisimulations in Benthem & Bonnay (2008), in order to capture the logicality for monstrous expressions. As our goal is primarily to advertise and sketch how compositional characters are both possible and useful, we'll leave this particular application of the theory to future work.

9 Conclusion

Above we've done a few things. First, we've outlined what it would be for characters to be compositional in the sense that's relevant for a minimally adequate semantic theory. Second, we've developed an model of compositional characters for a non-trivial fragment of English which satisfies these tenets of compositionality. In the course of our discussion, we offered what we take to be an improved understanding of the linguistic meanings of demonstratives, the ways that presuppositions can be incorporated into character-based theories of meaning, and the notions of logical truth and analyticity.

In closing, we would like to return to our initial motivation for offering out theory: the desire to fill the lacuna in Stalnaker's account of how we can reason about utterances of sentences like (1) even when we have only imperfect knowledge of the context of utterance. We take ourselves to have made serious strides towards this end, but not to have finished the task. It will help to say something about what we have done, and what still remains.

[31]See MacFarlane (2000) and Benthem & Bonnay (2008) for relevant discussion.

In furnishing a theory of how we can come to know the rule of use for a sentence on the basis of our knowledge of the rules of use for its parts, we take ourselves to have offered an explanation of how one can associate a propositional concept with an arbitrary well-formed sentence. While the notations are superficially different, our complex characters can be used to generate sets of propositions by instantiating any context-sensitive elements (like a_c) with each member of the domain—or, in plainer English, with each candidate-value. So these characters represent information in the same sense that Stalnaker's propositional functions do. This information is *partial* relative to the information represented by what the speaker is saying; that is, on recovering it, the listener can only cut down her uncertainty about the world to a more limited extent than if she had successfully recovered what the speaker asserted.[32] Accordingly, one might conceive of the sort of system that we have provided, and which Stalnaker's style of reasoning required, as a theory of *partial understanding.* It characterizes the sort of understanding that a listener can achieve merely in virtue of engaging her syntactic and semantic capacities, without engaging in any general-purpose reasoning.[33] This yields a sort of content which is substantially less precise than truth-conditions, but we view that as a feature of the view rather than a bug.[34]

Still, much remains undone. First, we need to see whether the theory really can be extended to a larger fragment of natural language, one which includes attitude verbs, modals, etc. Second, to complete Stalnaker's picture, we would need to say more about the *logic* of complex characters, about their entailment structure. For Stalnaker, this is relatively straightforward: complex characters (or propositional concepts) are sets of propositions, which are themselves conceived of as sets of worlds. Entailment is just truth at any world in the set. But this means that the character of (1) entails every mathematical truth. Should we hope to avoid this sort of result, we'll need

[32]See Shannon (1948) for the relevant notion of information. We're tacitly assuming, as Shannon does, that the listener trusts the speaker to be speaking truthfully.

[33]Granted, the listener will have to bring certain background assumptions—about what language is being spoken, etc.—to bear that will likely involve general-purpose reasoning. We think of these assumptions as *enabling* a modular (or at least potentially modular) process of interpretation rather than interfering with or undermining it.

[34]See Borg (2004) for arguments to the effect that we should think of our semantic system as *both* modular and as outputting truth-conditions. We would note that it is perfectly compatible with our view that the semantic module is *also* capable of outputting truth-conditions, in addition to complex characters, perhaps given additional inputs. We are skeptical that it does, but we will have to defer arguing the point to another occasion.

to define a different sort of logic of character, one capable of preserving what we might intuitively call the *aboutness* of natural language sentences like (1). That is very unlikely to prove a trivial task.

In spite of these outstanding challenges, we hope to have made at least a modest contribution here. King & Stanley (2005) famously claimed that complex characters could only ever hope to be 'trivially' compositional, whatever exactly that means (p. 123). We have shown that nothing of the sort is true, largely by offering as a proof-of-concept a non-trivial compositional theory of character for a small but interesting fragment of English. Perhaps this theory will ultimately prove unhelpful in trying to understand our knowledge of language, our knowledge of meaning. But we take it that, at present, we have every reason to believe the contrary.[35]

References

Beaver, David I. & Bart Geurts. 2014. Presupposition. In Edward N. Zalta (ed.), *The Stanford encyclopedia of philosophy*, Winter 2014 edn.

Benthem, Johan Van & Denis Bonnay. 2008. Modal logic and invariance. *Journal of Applied Non-Classical Logics* 18(2-3). 153–173.

Borg, Emma. 2004. *Minimal semantics*. Oxford: Oxford University Press.

Braun, David. 1994. Structured characters and complex demonstratives. *Philosophical Studies* 74. 193–214.

Braun, David. 1996. Demonstratives and their linguistic meanings. *Noûs* 30. 145–173.

Cohen, Jonathan & Eliot Michaelson. 2013. Indexicality and the answering machine paradox. *Philosophy Compass* 8(6). 580–592.

Cumming, Samuel. 2008. Variabilism. *Philosophical Review* 117(4). 525–554.

Elbourne, Paul. 2005. *Situations and individuals*. Cambridge: MIT Press.

[35]For discusssion and feedback, thanks to the brothers Charlow, Sam Cumming, Julien Dutant, Jessica Keiser, Matt Mandelkern, Ethan Nowak, Brian Rabern, and Anders Schoubye. Thanks as well to or audiences at the Iceland Meaning Workshop, the University of Leeds, and King's College London.

Elbourne, Paul. 2008. Demonstratives as individual concepts. *Linguistics and Philosophy* 31(4). 409–466.

van Frasssen, Bas. 1990. Figures in a probability landscape. In Michael Dunn & Anil Gupta (eds.), *Truth or consequences*, 345–356. Dordrecht: Springer.

Glick, Ephraim. 2017. Cross-contextual semantics. Manuscript.

Hanson, William H. 2006. Actuality, necessity, and logical truth. *Philosophical Studies* 130(3). 437–459.

Harris, Daniel. 2019. Semantics without semantic content. Forthcoming in *Mind and Language*.

Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Oxford: Blackwell Publishers.

Kaplan, David. 1989a. Afterthoughts. In Joseph Almog, John Perry & Howard Wettstein (eds.), *Themes from Kaplan*, 565–614. Oxford: Oxford University Press.

Kaplan, David. 1989b. Demonstratives. In Joseph Almog, John Perry & Howard Wettstein (eds.), *Themes from Kaplan*, 481–563. Oxford: Oxford University Press.

Kennedy, Christopher. 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* 30(1). 1–45.

Kennedy, Christopher & Louise McNally. 2010. Color, context, and compositionality. *Synthese* 174(1). 79–98.

King, Jeffrey C. & Jason Stanley. 2005. Semantics, pragmatics, and the role of semantic content. In Zoltán Szabó (ed.), *Semantics versus pragmatics*, 111–164. Oxford: Oxford University Press.

Levi, Isaac. 1980. *The enterprise of knowledge*. Cambridge, Massachusetts: MIT Press.

Lewis, David. 1980. Index, context, and content. In S. Kanger & S. Öhman (eds.), *Philosophy and grammar*, 79–100. Dordrecht: Reidel.

MacFarlane, John. 2000. *What does it mean to say that logic is formal?* Pittsburgh: University of Pittsburgh dissertation.

Matthews, Robert. 2006. Knowledge of language and linguistic competence. *Philosophical Issues* 16. 200–220.

Novel, Marc & Maribel Romero. 2010. Movement, variables and hamblin alternatives. In Viola Schmitt & Sarah Zobel (eds.), *Proceedings of Sinn und Bedeutung*, vol. 14, 322–338. Vienna.

Pagin, Peter & Dag Westerståhl. 2011. Compositionality. In C. Maienborn, K. von Heusinger & P. Portner (eds.), *Semantics: An international handbook of natural language meaning*, 96–123. Elsevier.

Poesio, Massimo. 1994. Semantic ambiguity and perceived ambiguity. In Kees van Deemter & Stanley Peters (eds.), *Semantic ambiguity and underspecification*, 159–201. Stanford: CSLI Press.

Rabern, Brian. 2012. Against the identification of assertoric content with compositional value. *Synthese* 189(1). 75–96.

Rabern, Brian. 2013. Monsters in kaplan's logic of demonstratives. *Philosophical Studies* 164(2). 393–404.

Radulescu, Alexandru. 2012. *The logic of indexicals*. Los Angeles: University of California, Los Angeles dissertation.

Radulescu, Alexandru. 2015. The logic of indexicals. *Synthese* 192(6). 1839–1860.

Schoubye, Anders J. 2013. Ghosts, murderers, and the semantics of descriptions. *Noûs* 47(3). 496–533.

Shannon, Claude E. 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27. 379–423, 623–656.

Stalnaker, Robert. 1978. Assertion. In Peter Cole (ed.), *Syntax and semantics 9: Pragmatics*, 315–332. Academic Press.

Westerståhl, Dag. 2012. Compositionality in Kaplan style semantics. In W. Hinzen, E. Machery & M. Werning (eds.), *The Oxford handbook of compositionality*, 192–219. Oxford University Press.

Woods, Jack. 2016. Characterizing invariance. *Ergo* 3(30). 778–807.

Zalta, Ed. 1988. *Intensional logic and the metaphysics of intentionality*. Cambridge, Massachusetts: MIT Press.