Contextualism Challenged:

Three Puzzles Concerning Indexicality and Context-Sensitivity

A dissertation submitted for the degree of
Doctor of Philosophy
in Philosophy

by

Sebastian Petzolt

St Hilda's College

University of Oxford

Contextualism Challenged:

Three Puzzles Concerning Indexicality and Context-Sensitivity

A dissertation submitted for the degree of
Doctor of Philosophy
in Philosophy

Sebastian Petzolt
St Hilda's College
University of Oxford
Trinity Term 2014

ABSTRACT

This thesis contains three papers, *Indexicality and Co-Reference*, *Subtle Indexicals: Kaplanian Semantics and the Answering Machine Paradox*, and *Compounds in Context*. Each paper presents a set of linguistic data that challenges our understanding of context-sensitivity. The first two papers consider certain utterances of indexical-containing sentences and show that Kaplan's theory of indexicals cannot assign the correct content to the utterances in question. The third paper rebuts an argument according to which a certain version of Contextualism is incoherent with the Principle of Compositionality.

Indexicality and Co-Reference considers sentences in which a given indexical occurs multiply, such as 'The man I see now is younger than the man I see now'. While there are true utterances of this sentence, Kaplan's account implies that all of the sentence's utterances are false. Although Kaplan's theory can be revised so as to avoid the problem, the revised theory is substantially less attractive than the original theory.

Subtle Indexicals presents and ultimately strengthens the so-called Answering Machine Paradox. Roughly, while Kaplan's theory implies that utterances of 'now' and 'here' always refer to the time and place of the utterance, there is linguistic evidence to the contrary, including instances of the so-called historic present: 'It is 14 July 1789. Now the partisans storm the Bastille.' Several solutions to the paradox are considered, but it is argued that they are unsuccessful.

Compounds in Context takes a more positive stance. The paper notes that noun-compounds like 'carbon filter' and 'fish van' are intuitively context-sensitive. Then an argument is presented according to which Contextualism about noun-compounds is inconsistent with the Principle of Compositionality. The argument is rebutted and a Contextualist semantic theory for noun-compounds is developed.

In sum, each paper starts from linguistic intuitions, shows that taking these intuitions at face value leads to a puzzle about context-sensitivity, and asks whether a semantic theory can be developed that solves the puzzle.

Word count: 64,249

Contents

Overview 1

Indexicality and Co-reference 8

1. The Kaplanian theory 9
2. The Multiple Occurrence Problem 19
3. Further instances of the MOP 26
4. The Multiple Contexts Analysis 38
5. The Token-Reflexive Theory 42
6. The CST: contexts 60
7. The CST: semantics 62
8. The CST: utterance-content 67
9. The CST: Occurrence 69
10. Validity 74
11. Synonymy 76
12. Conclusion 85

Subtle Indexicals: Kaplanian Semantics and the Answering Machine Paradox 88

1. The Kaplanian theory 88
2. The Puzzle 91
3. Four potential solutions to the AMP 95
4. Corazza's Anaphoric Theory of Indexicals 108
5. EXPRESSION and the Humpty-Dumpty Problem 113
6. PROPOSITION: Some details 114
7. PROPOSITION: finding suitable antecedents 117
8. PROPOSITION: lack of suitable antecedents 121
9. Strengthening the paradox 126

Compounds in Context 139

1. The Compound-Argument against Compositionality 140
2. The syntactic status of ‘carbon filter’: word or phrase? 145
3. The syntactic status of ‘carbon filter’: constituents 155
4. The context-sensitivity of ‘carbon filter’ 166
5. The meanings of ‘carbon’ and ‘filter’: preliminaries 173
6. A semantic theory for ‘carbon’ 182
7. Objections 194
8. Conclusion 202

Overview

Oh Context*: velut luna, statu variabilis.*
Nunc per ludum dorsum nudum fero tui sceleris.

To say that an expression is *context-sensitive* is to say, roughly, that the expression has different contents on different occasions of use. More formally put, context-sensitive expressions do not have contents *simpliciter*, but only relative to points of semantic evaluation, called 'contexts'. Hence when *e* is a context-sensitive expression, one cannot speak of *the* content of *e*, but only of the content that *e* has *relative to* a particular context *c*.

This thesis contains three papers: *Indexicality and Co-Reference*, *Subtle Indexicals: Kaplanian Semantics and the Answering Machine Paradox*, and *Compounds in Context*. Each paper presents a set of linguistic data that challenges our understanding of context-sensitivity. The first two papers consider certain utterances of indexical-containing sentences and show that Kaplan's theory of indexicals cannot assign the correct content to the utterances in question. They conclude that Kaplan's theory does not provide adequate semantics for indexicals. While this is a fairly technical result, the third paper discusses a more conceptual challenge: it rebuts an argument according to which a certain version of Contextualism is incoherent with the Principle of Compositionality.

In sum, each paper starts from linguistic intuitions, shows that taking these intuitions at face value leads to a puzzle about context-sensitivity, and asks whether a semantic theory can be developed that solves the puzzle.

Since both *Indexicality and Co-Reference* and *Subtle Indexicals* discuss indexicals, an overlap is unavoidable. In particular, both theories outline Kaplan's theory of

indexicals. However, *Indexicality and Co-Reference* provides a much more detailed presentation, and *Subtle Indexicals* occasionally refers to *Indexicality and Co-Reference* to avoid unnecessary repetitions.

Although *Indexicality and Co-Reference* and *Subtle Indexicals* argue that Kaplan's theory of indexicals is inadequate, *Compounds in Context* presupposes a broadly Kaplanian framework of context-sensitivity. In particular, the paper assumes that expressions have contents relative to contexts, that sentences can be true and false relative to contexts, and that utterances occur in contexts. This does not contradict the results from *Indexicality and Co-Reference* and *Subtle Indexicals*, however: The latter two papers argue that Kaplan's theory does not provide adequate semantics for *indexicals*. Yet in and of itself, this does not show that one cannot spell out claims about context-sensitivity in terms of broadly Kaplanian notions and concepts.

I consider several solutions for each of the puzzles I raise. However, I cannot consider all possible solutions. In particular, I do not consider relativist theories, according to which propositions do not have truth-values simpliciter. I also don't consider any *unarticulated-constituents*-theories. Finally, I assume that every sentence-utterance expresses (at most) one proposition, i.e. I ignore *multiple-propositions*-theories.

The remainder of this introduction outlines each of the three papers in some more detail.

Indexicality and Co-Reference considers sentences in which a given indexical occurs multiply, such as 'The man I see now is younger than the man I see now'. While there are apparently true utterances of this sentence, it is proven that Kaplan's account implies that all utterances of the sentence at hand are false: Each sentence-utterance occurs in *exactly one* context, according to Kaplan, and the utterance's content is the

content that the uttered sentence has at the context in question. In turn, each indexical has *exactly one* content at any given context. Hence when the above sentence is evaluated at a given context, both occurrences of 'now' are assigned the same content. Therefore, the sentence expresses a proposition of the form of (U) at every context, according to Kaplan.

(U) that the man S sees at *t* is younger than the man S sees at *t*.

All propositions of this form being false, Kaplan's theory cannot account for the fact that there are true utterances of the above sample sentence.

I show that the so-called Token-Reflexive Theory can solve the problem at hand. This theory does not assign contents to expressions in context, but rather assigns contents to word-and sentence-utterances directly. In particular, given any utterance of the above sample sentence, the two relevant utterances of 'now' are interpreted *separately*, so that they can refer to different times. When they do refer to different times, the sentence's utterance expresses a content that is not of the form of (U), according to the Token-Reflexive Theory.

Since the Token-Reflexive Theory does not assign contents to expressions in context, it is a rather radical alternative to Kaplan's theory of indexicals. Further, proponents of Kaplan's theory frequently object that the Token-Reflexive Theory is inferior because it cannot provide an adequate definition of validity. (Kaplan's theory provides such a definition.) I reply to this objection on behalf of the Token-Reflexive Theory. However, I then object that the theory cannot provide an adequate definition of synonymy. As Kaplan's theory provides such a definition, I conclude that the Token-Reflexive Theory is indeed inferior to Kaplan's theory.

I next ask whether a revised version of Kaplan's theory can solve the problem at hand, while still providing an adequate definition of synonymy. In particular, I consider

a version of Kaplan's theory that allows sentence-utterances to occur in multiple contexts. E.g. when the above sample sentence is uttered, the relevant utterance of 'The man I see now' occurs in one context, while the utterance of 'is younger than the man I see now' occurs in a different context. Since 'now' may have different contents at the contexts in question, this guarantees that some of the sentence's utterances express a content that is not of the form of (U).

While the proposed revision allows the Kaplanian to solve the problem at hand, it is shown that the resulting theory can no longer provide an adequate definition of synonymy. (An adequate definition of validity can still be provided, however.) I conclude, therefore, that the revised theory is not an adequate theory of indexicals.

Subtle Indexicals presents and ultimately strengthens a fairly well-known paradox, viz. the so-called Answering Machine Paradox. Roughly, while Kaplan's theory of indexicals implies that utterances of 'now' and 'here' always refer to the time and place of the utterance, there is linguistic evidence to the contrary. E.g. instances of the so-called historic present provide evidence of this kind: 'It is 14 July 1789. Now the partisans storm the Bastille.' Contrary to Kaplan's theory, the relevant utterance of 'now' intuitively refers to 14 July 1789, even if the utterance was produced on a different day.

Several solutions to the paradox are considered, but it is argued that they are ultimately unsuccessful. According to one solution, utterances of the above sentence pragmatically communicate a statement about 14 July 1789, while semantically expressing a statement about the time of the speaking. It is shown that standard accounts of conversational implicature do not warrant this tenet. According to another solution, utterances of 'now' always refer to the time that the speaker intends as referent, so that 'now' refers to the time of the speaking only if the speaker intends to speak about this

time. It is shown that in ordinary conversations, ‘now’ refers to the time of the speaking, regardless of the speaker’s intentions. According to a third solution, the above utterance of ‘now’ is anaphoric on ‘14 July 1789’ and inherits this time as referent. It is shown that this tenet cannot account for all the data. E.g. if someone’s will states ‘I am now no longer with you’, the relevant utterance of ‘now’ intuitively refers to the time at which the will is read, even if the deceased did not utter an expression that denotes this time, i.e. even if there is no expression that the utterance of ‘now’ could be anaphoric on.

The final section of the paper presents a set of previously unnoticed data. As the data shows, whether a given utterance of ‘now’ can refer to a time in the past is a rather subtle matter. E.g. in ‘On 14 July 1789, the partisans now storm the Bastille’, the utterance of ‘now’ cannot be read as referring to 14 July 1789 (unless it was produced on this day). By contrast, it can be so read in e.g. ‘On 14 July 1789, the French Revolution reaches a peak as the partisans now storm the Bastille.’ No theory of indexicals I know of can account for this subtlety, which considerably strengthens the Answering Machine Paradox.

In sum, *Indexicality and Co-Reference* and *Subtle Indexicals* argue for a rather negative result: although Kaplan’s theory is widely accepted among philosophers of language, and although it is arguably the best available semantic theory of indexicals, it is flawed.

Compounds in Context takes a more positive stance. Leaving behind indexicals, the paper considers whether noun-compounds like ‘carbon filter’ and ‘fish van’ may be context-sensitive. Apparently they are: e.g. ‘carbon filter’ can denote a device for filtering carbon, but also a device that is made of carbon and used for filtering e.g. smoke. However, it is arguable that Contextualism about noun-compounds is

inconsistent with the Principle of Compositionality. According to Compositionality, the content of any complex expression *e* at context *c* is determined by the contents of the constituents of *e* at *c*, together with *e*'s structure. Yet, whether 'carbon filter' denotes, at *c*, a device for filtering carbon or a filtering device that is made of carbon, 'carbon' apparently denotes carbon at *c* and 'filter' denotes filters at *c*. Hence the contents of 'carbon' and 'filter' do not co-vary with the content of 'carbon filter', which is to say that the content of 'carbon filter' at *c* is not determined by the contents of its constituents at *c* and its structure.

The paper considers three replies. According to the first, 'carbon filter' is not a complex expression, appearances to the contrary notwithstanding, so that Compositionality does not apply to 'carbon filter'. It is argued that the reply fails because 'carbon filter' behaves like a complex expression. E.g. just like one can insert an expression into the phrase 'large table' – as in 'large wooden table' – one can insert an expression into 'carbon filter': e.g. 'carbon airlift filter' is perfectly felicitous.

The second reply rejects that 'carbon filter' is context-sensitive, by contrasting 'carbon filter' with indexicals. In particular, that there are true utterances of 'It is Monday today' does not imply that it is Monday today. By contrast, if John doesn't own a carbon filter, it seems that there can be no true utterances of 'John owns a carbon filter'. While it is acknowledged that this test brings out an interesting contrast, the test is rejected in the end: Only indexicals and demonstratives pass the test. Hence unless one is willing to accept that only indexicals and demonstratives are context-sensitive, one cannot use the test to establish that 'carbon filter' is context-insensitive.

The third reply develops a Contextualist theory for 'carbon'. Roughly, 'carbon' has the same content as 'carbon-made' at contexts at which 'carbon filter' denotes a filter that is made of carbon. It has the same content as 'carbon-purifying' at contexts at which 'carbon filter' denotes a device for filtering carbon. If this is correct, the content

of ‘carbon’ co-varies with the content of ‘carbon filter’, whence the context-sensitivity of ‘carbon filter’ is reconciled with Compositionality. After replying to several objections to the Contextualist theory at hand, I conclude that it is correct, and that Contextualism about noun-compounds is consistent with Compositionality.

Indexicality and Co-reference

In this paper I shall present and discuss a new problem for the standard Kaplanian account of indexicals. According to this account, indexicals and sentences that contain indexicals do not have contents simpliciter, but only relative to *contexts*. Further, the content of an *utterance* of a given sentence is the content that the uttered sentence has relative to the context in which the utterance occurs. Now consider the following case.

CLUBLAND. Peter and Frank are members of a private membership club in London. One day when in the club, they speak to another club member who claims: 'All club members who live in London regularly dine at the club.' Peter knows that this is false and states: 'Some of the members who live in London don't dine at the club.' To prove his point, he adds:

(1) Frank lives here and Frank never dines here.

Intuitively, Peter's utterance of (1) expresses that Frank lives *in London* and never dines *at the club*. Yet the Kaplanian account does not yield this result. If *c* is the context in which Peter's utterance occurs, then the account implies that the content of Peter's utterance of (1) is the content that (1) has relative to *c*. Further, what content (1) has at *c* depends on what contents the constituents of (1) have at *c*, and in particular on what content 'here' has at *c*: if London is the content of 'here' at *c*, the content of (1) at *c* is the proposition that Frank lives in London and never dines in London; if the club is the content of 'here' at *c*, the content of (1) at *c* is the proposition that Frank lives in the club and never dines in the club. More generally, the content of (1) at *c*, and hence the content of Peter's utterance, is a *uniform* proposition of the form of (U), according to the Kaplanian account.

(U) that Frank lives in *l* and never dines in *l*.

Yet as noted, Peter's utterance of (1) does not express a proposition of the form of (U). Thus, the Kaplanian account fails to assign the correct content to Peter's utterance of (1).

The case at hand involves 'here', but one can create cases like it for most indexicals, as will be shown below. The cases all follow the same pattern: a sentence is chosen in which the target indexical occurs multiply, a particular utterance of the sentence is considered, and it is shown that the Kaplanian account makes a false prediction about the content of the sentence-utterance. Call this the *Multiple Occurrence Problem*, or *MOP*. After a summary of the Kaplanian account, I shall explain the problem more fully and shall discuss it in detail.

PART I: THE KAPLANIAN THEORY AND THE MULTIPLE OCCURRENCE PROBLEM

1. The Kaplanian theory

The theory I am about to present is in essence Kaplan's theory of pure indexicals, but deviates from his original theory in some details, mostly for the sake of simplicity. The deviations have no bearing on the discussion in this paper.

1.1 Indexicals and their contents

To start, Kaplan acknowledges two kinds of indexicals: pure indexicals and true demonstratives. To identify the expressions in question, he does not use a definition, but simply lists the expressions that he takes to be indexicals. In particular, he lists 'I', 'here', 'now', 'today', 'yesterday', 'tomorrow', 'actual', and 'present' as pure indexicals, while listing 'this', 'that', 'he', 'she', and 'it' as true demonstratives. One

can consider adjusting Kaplan's list[1], and one can try to define what pure indexicals and demonstratives are, but I shall attempt neither.

In his *Demonstratives*[2], Kaplan presents a semantic theory for pure indexicals, and a second, separate theory for true demonstratives. The present paper is only concerned with the former theory, but towards the end of this section I shall briefly explain how his two theories differ from one another. Otherwise I shall largely ignore demonstratives. Since pure indexicals are the only indexicals I consider in detail, I shall henceforth use the term 'indexical' to mean the expressions that Kaplan calls 'pure indexicals'.

Turning, then, to Kaplan's theory of (pure) indexicals, the first claim that Kaplan makes is that indexicals are *context-sensitive*. He spells this idea out by saying that indexicals do not have contents *simpliciter*, but only relative to entities that he calls *contexts*. Relative to any given context, each indexical has exactly one content, so that there are functions from contexts to contents, one function for each indexical. Kaplan calls the functions *characters*. I shall return to characters below. For now, it suffices to note that indexicals have contents only relative to contexts.

Indexicals aren't the only expressions that have contents relative to context: indeed, Kaplan holds that every word can be evaluated for content at any given context. If a word is context-*in*sensitive, it has the same content at all contexts. In that sense, context-insensitive words have contents *simpliciter*. If there are context-sensitive words other than indexicals, then these words vary in content from context to context.

Sentences, too, are assigned contents relative to contexts, in accordance with the following principle:

[1] See e.g. Cappelen and Lepore 2005 for such an attempt.
[2] Kaplan 1989a.

SENTENCE-CONTENT. For all sentences σ and contexts *c*, the content of σ at *c* is determined by σ's compositional structure and the contents that σ's constituents have at *c*.

Thus to determine what content σ has at *c*, one must first determine what contents the constituents of σ have at *c*. These contents, together with σ's compositional structure, then yield the content of σ at *c*. To illustrate, consider sentence (2).

(2) Smith now courts Mary.

Let *c* be a particular context, and suppose that June 2013, say, is the content of 'now' at *c*. Assume further that 'Smith', 'courts', and 'Mary' denote Will Smith, the *courting*-relation, and Queen Mary at *c*. Then (2) expresses at *c* that Will Smith courts Queen Mary in June 2013.

Since the notion of context is central to Kaplan's theory, the next sub-section spells this notion out in detail.

1.2 Contexts

According to Kaplan, contexts are ordered *n*-tuples. More particularly, I shall assume that Kaplanian contexts are ordered quadruples, containing an agent, a time interval, a region of space, and a possible world, respectively called the *speaker*, *time*, *location* and *world* of the context. For the most part of the paper, the worlds of the contexts can be ignored: unless stated otherwise, the world of any context that I mention is the actual world.

Kaplan distinguishes *proper* from *improper* contexts. Any ordered quadruple that contains an agent, a time, a location, and a world is a context. For a quadruple to be a *proper* context, the quadruple's agent must, in the quadruple's world, be located in the

quadruple's location, at the quadruple's time. Kaplan further stipulates that expressions can only be evaluated at proper contexts. I shall accept this stipulation, but nothing depends on it.

Note that Kaplan's notion of context is a *theoretical* notion and distinct from various more informal notions of context. According to these latter notions, a context is roughly the particular situation in which, or the background against which, a conversation takes place. By contrast, Kaplanian contexts are abstract, set-theoretic entities.

It matters for at least two reasons whether one uses a formal or rather an informal notion of context. Firstly, ordered *n*-tuples can contain any time and location one likes, including e.g. very short timespans during which no (actual or possible) human could utter anything. Hence on the technical notion of 'context', there are contexts that include extremely short time spans and one can evaluate expressions relative to these contexts. By contrast, no conversation takes place during any 'real-life' situation which last for only such an extremely short time span. Hence on the informal notion of 'context', there is no context which last for only an extremely short time span. Thus which notion of context one uses influences what 'contexts' there are.

Which notion of context one uses matters for a second reason. To see it, consider the view that 'know' is context-sensitive. A proponent of this view might claim: 'exactly one epistemic standard is operative in each context' [3]. This is fine as a claim about ordered *n*-tuples as one can, of course, stipulate that each of the *n*-tuples one calls 'contexts' contains exactly one epistemic standard (in addition to a speaker, a time, a place, and a world). By contrast, if the claim is understood as meaning that there is exactly one epistemic standard operative in each conversation, the claim is rather

[3] Very roughly, an epistemic standard is a criterion by which we judge claims of the form 'S knows that *p*'; see e.g. DeRose 2009 for further details.

contentious: it is by no means clear that each conversation, or even each utterance of a sentence, comes with exactly one criterion for judging the epistemic state of agents.

Note that I shall henceforth mean Kaplanian contexts when I say 'context'.

1.3 Characters

As stated, each indexical has exactly one content relative to each context, so that each indexical can be associated with a content-determining function from contexts to contents. Kaplan calls these functions *characters*. The content of indexical *i* relative to context *c* is then the value that *i*'s character yields for *c*, and the claim that indexicals are context-sensitive becomes tantamount to the claim that indexicals have non-constant characters.

It is widely agreed that the character of 'I' maps each context *c* to the agent of *c*. It is furthermore fairly common to think that the characters of 'now' and 'here' function as described in NOW and HERE[4].

NOW. For all contexts *c*, the character of 'now' maps *c* to the time of *c*.

HERE. For all contexts *c*, the character of 'here' maps *c* to the location of *c*.

Below I shall present a reason for thinking that the characters of 'now' and 'here' are more complex than NOW and HERE suggest. Yet for now, I assume that NOW and HERE are correct.

The characters of 'today', 'yesterday', and 'tomorrow' are somewhat more difficult to characterise. Let the day of *c* be the day on which the time of *c* takes place[5]. Then one might think that the three characters respectively map each context *c* to the day of *c*, to the day before the day of *c*, and to the day after the day of *c*. However, consider e.g. 'In

[4] Kaplan (1989a, p. 564) endorses HERE; e.g. Predelli (1998) endorses both NOW and HERE.
[5] If there are contexts the times of which are longer than a day, matters become even more complicated.

the Middle Ages, many people died from smallpox. Today smallpox is no longer a threat.' The second sentence can plausibly be interpreted as expressing that smallpox is no longer a threat *nowadays*, or *in our current era*. That the sentence can be thus interpreted suggests that there is a context at which our current era is the content of 'today'. If this is correct, then it is not true that the character of 'today' maps each context *c* to the day of *c*. I shall return to this in sub-section 2.3, but shall assume for now that the characters of 'today', 'yesterday', and 'tomorrow' are as described above.

1.4 Validity and synonymy

Now that the notions of context and character have been introduced, we can present one of the major advantages of the Kaplanian theory: it can provide adequate definitions of several semantically important notions, such as validity and synonymy.

Intuitively, e.g. 'If John is here now, then John is here now' is valid, while 'If John is here today, then John is here tomorrow' is invalid. Further, 'John is here now' and 'Now John is here' are intuitively synonymous. A good semantic theory of indexicals should capture these intuitive judgements, and the Kaplanian theory can capture them. Starting with validity, the Kaplanian definition is as follows[6].

> VAL$_K$. For all sentences σ, σ is valid iff: for all contexts *c*, the proposition that σ expresses at *c* is true in the world of *c*.

At any context *c*, 'If John is here now, then John is here now' expresses that John is in *l* at *t* if John is in *l* at *t* – where *l* and *t* are the location and time of *c*. Since this proposition is necessary, it is true in the world of *c*. Hence VAL$_K$ correctly judges the sentence to be valid. By contrast, 'If John is here today, then John is here tomorrow'

[6] This is a simplified version of Kaplan's definition of validity, but for present purposes, the simplified version will suffice. For Kaplan's original definition, see his 1989a, p. 547.

expresses at any context c that John is in l on d_2 if John is in l on d_1 – where d_1 is the day of c, and where d_2 is the day after d_1. Since this proposition may well be false in the world of c, the sentence is correctly judged as invalid.

As contexts are defined, the agent of any proper context c is located in the location of c at the time of c in the world of c. Given VAL_K, and given that expressions may only be evaluated at proper contexts, this means that 'I am here now' is valid. While Kaplan held that this is the correct result, other authors hold that 'I am here now' is invalid[7]. If one prefers the latter view, the best option is to keep VAL_K as it stands, but allow expressions to be evaluated at improper contexts.

The Kaplanian can also define logical consequence: a sentence σ is a consequence of a set Γ of sentences iff for all contexts c, the content of σ at c is true in the world of c if the contents of the sentences in Γ at c are true in the world of c. Further, a set Γ of sentences is consistent iff there is a context c such that all the sentences in Γ express a proposition at c that is true in the world of c. Γ is inconsistent iff it is not consistent.

Next consider the notion of synonymy. Pre-theoretically, synonymy is sameness of meaning. Now, the Kaplanian theory recognises two kinds of 'meaning', viz. character and content. Clearly, that two expressions have the same content relative to a particular context doesn't make the expressions synonymous. E.g. one can find a context relative to which 'now' and 'today' have the same content, but 'now' and 'today' are clearly not synonymous. Hence synonymy is defined as sameness of character: "To supply a synonym for a word or phrase is to find another with the same *character*."[8]

Traditionally, synonymous expressions are thought to share all semantic properties. A major advantage of the proposed definition is that it respects this aspect of synonymy. For instance, consider the property of expressing a true content at c. If two sentences

[7] See e.g. Predelli 2004.
[8] Kaplan 1989a, p. 521, emphasis in the original.

have the same character, both sentences have the said property if one does. Hence given the current definition of synonymy, it is never the case that only one of two synonymous sentences instantiates the property in question.

Another advantage of the definition is that it applies to sentences from different languages. E.g. the definition correctly yields that 'John is here now' is synonymous with the German sentence 'John ist jetzt hier'.

In sum, the Kaplanian theory provides definitions of validity and synonymy that are at least prima facie adequate. Later in this paper, I shall present alternatives to the Kaplanian theory. One way of judging these theories will be by asking whether they, too, can provide adequate definitions of validity and synonymy.

1.5 Utterances

The previous sections explained how Kaplan's theory assigns contents to indexicals and sentences at contexts. However, *utterances* also have contents. So the next step is to explain how Kaplan's theory can assign contents to utterances. Strictly, that is a deviation from Kaplan's original account: in his *Demonstratives*, Kaplan does not consider utterances, and his theory officially doesn't assign contents to sentence-utterances[9]. He only considers sentence-context pairs, which *represent* utterances, according to Kaplan. But it is now common to extend Kaplan's theory so that it can assign contents to utterances also[10]. In what follows, I shall focus on utterances of sentences.

To start, the notion of context is connected with the notion of utterance, by saying that utterances *occur* in contexts. Commonly, the context in which a given utterance

[9] I say that the theory *officially* doesn't assign contents to sentence-utterances because Kaplan does occasionally speak about the contents of utterances.
[10] E.g. Predelli (1998), Sidelle (1991), and Perry (2003) propose extensions of this kind.

occurs in called the utterance's *context of origin*. In turn, the notion of occurrence can be spelled out thus[11].

OCCUR. For all utterances *u*, *u* occurs in <*S*, *t*, *l*, *w*> iff: S uttered *u* during time *t* in world *w*, and *l* is the smallest location that S occupies at *t*.

By the time *during* which S utters *u* I mean the time interval that starts when S starts uttering *u* and ends when S has finished uttering *u*.

To illustrate OCCUR, suppose Jane utters (2) = 'Smith now courts Mary' at 10 am, and suppose that it takes her two seconds to utter this sentence. Then the context in which her sentence-utterance occurs has Jane as speaker. It has as time the interval that starts at 10 am and ends at 10:00:02. The location of the context is the smallest location that Jane occupies during that interval.

With OCCUR in place, the Kaplanian can assign contents to sentence-utterances[12]. This is done by means of the following principle[13, 14]:

UTTERANCE-CONTENT: For all utterances *u* and sentences σ, if *u* is an utterance of σ, then the content of *u* is the content that σ has relative to the context in which *u* occurs.

To illustrate, assume again that Jane utters (2). Suppose further that her utterance occurs in context *c*, and that the content of (2) at *c* is the proposition that Smith courts Mary at *t*. Then UTTERANCE-CONTENT implies that this proposition is the content of Jane's utterance.

---

[11] Perry (2003, p. 375) endorses a version of this principle.
[12] I'll assume, with Kaplan, that for each utterance *u*, there is a unique expression that *u* is an utterance of.
[13] Kaplan makes a statement that comes very close to UTTERANCE-CONTENT: "The Content of a sentence in a context is […] the proposition the sentence would express if uttered in that context". See Kaplan 1989a, p. 546. Predelli (1998, p. 107) and Sidelle (1991, p. 525) also endorse versions of UTTERANCE-CONTENT.
[14] Some authors hold that an utterance of a demonstrative has a content only if the utterance is accompanied by a demonstration (see below). If this is correct, then UTTERANCE-CONTENT does not apply to utterances of sentences that have demonstratives as constituents. Since I won't consider demonstratives in any detail, this complication is irrelevant for present purposes.

Note an important feature of UTTERANCE-CONTENT: what content an utterance *u* of sentence σ has depends on what content the constituents of σ have at the context in which *u* occurs, but does not depend on what contents the sub-utterances of *u* have. To illustrate, suppose again that *c* is the context in which Jane's utterance of (2) occurs. What content Jane's utterance of (2) has then depends on what contents the *words* 'Smith', 'now', 'courts', and 'Mary' have at *c*, but not on what content Jane's *utterances* of these words have. Nor does it depend on what contents the words have at the contexts in which their utterances occur.

Note further that UTTERANCE-CONTENT presupposes that each sentence-utterance occurs in *exactly one* context. For, in assigning a content to *u,* the principle speaks of *the* context in which *u* occurs. Thus by endorsing UTTERANCE-CONTENT, the Kaplanian commits to UNIQUENESS:

> UNIQUENESS. For all sentence-utterances *u*, there is exactly one context *c* such that *u* occurs in *c*.

With UTTERANCE-CONTENT and UNIQUENESS in place, I am in a position to develop the MOP. But before I turn to this, I'd like to explain very briefly how Kaplan's theory of indexicals differs from his theory of demonstratives.

1.6 Indexicals versus Demonstratives

On Kaplan's theory, indexicals are *automatic*: What content a given indexical-utterance has – if any[15] – is determined only by the context in which the utterance occurs. In particular, the utterance's content is determined independently of any intentions that the speaker may have: If the speaker has no reference-fixing intentions

[15] It's a good question whether indexical-utterances can *fail* to refer: If one utters an indexical, the utterance occurs in some context or other. In turn, every context contains by definition a speaker, a time, and a location. Hence it seems that utterances of 'I', 'now', and 'here' cannot fail to refer. On the other hand, 'you' may be an indexical and utterances of 'you' can clearly fail to refer.

when producing her utterance, the utterance can still acquire a content. On the other hand, the utterance will acquire whatever content the context supplies, even if the speaker intends the utterance to refer to another individual, time, or location. In sum, as some authors put it, indexicals and indexical-containing sentences have their contents *just in virtue of being uttered*[16, 17]. Due to this fact, Kaplan speaks of *pure* indexicals.

By contrast, demonstratives are not automatic, according to Kaplan, which is why he develops a separate semantic theory for demonstratives. If one utters a demonstrative, one's utterance won't acquire a content unless the utterance is accompanied by an extra-linguistic act, a 'demonstration'. Pointing gestures are the paradigm example of demonstrations, but most philosophers think that not all demonstrations are pointing gestures. E.g. Kaplan holds in his later work that directing intentions are a kind of demonstration[18]. But whatever demonstrations are, utterances of demonstratives must be accompanied by *them* in order to acquire a referent, according to Kaplan.

Some authors hold *pace* Kaplan that indexicals are not automatic, that utterances of indexicals and indexical-containing sentences must also be accompanied by an extra-linguistic act to acquire a referent. Yet whether this is correct is irrelevant for my argument: *as Kaplan's account stands*, indexicals are automatic. And it is Kaplan's account as it stands that is the topic of this paper.

2. The Multiple Occurrence Problem

The present section shows that the sketched Kaplanian account faces a serious problem, viz. the MOP. In developing the MOP, I shall focus on CLUBLAND. Further instances of the problem will be presented in section three.

[16] Cf. e.g. Braun 1996, p 1.
[17] Holding that indexicals are automatic is consistent with holding that they have non-indexical uses. E.g. it is often held that 'here' is used demonstratively when one points at a map and says 'Last Monday, I was here'. For details, see e.g. Sidelle 1991, p.528, and the second paper in this dissertation, sec. 2.2 and 3.3.
[18] See his 1989b, esp. pp. 582 – 584.

2.1 The compositional structure of (1)

The first step is to establish a claim about the compositional structure of sentence (1). Let me repeat (1) for convenience. Also, recall that propositions of the form of (U) are called *uniform* propositions.

(1) Frank lives here and Frank never dines here.

(U) that Frank lives in *l* and never dines in *l*.

Next consider the following claim.

(Uni) For all contexts *c*, (1) expresses a uniform proposition at *c*.

Now, although the Kaplanian account is consistent with the negation of (Uni), the Kaplanian cannot plausibly reject (Uni). To see this, consider (3) and (4), which are the two conjuncts of (1).

(3) Frank lives here.

(4) Frank never dines here.

By SENTENCE-CONTENT, the content of (3) at a given context *c* is determined by the compositional structure of (3), together with the contents that the constituents of (3) have at *c*; similarly for (4). Now, it is *very* plausible that the compositional structures of (3) and (4) are such that COM-3 and COM-4 are true.

COM-3. For all contexts *c*, if location *l* is the content of 'here' at *c*, then (3) expresses at *c* that Frank lives in *l*.

COM-4. For all contexts *c*, if location *l* is the content of 'here' at *c*, then (4) expresses at *c* that Frank never dines in *l*.

By contrast, it would be *absurd* to think that (3) and (4) express e.g. P_1 and P_2 at all contexts at which England, say, is the content of 'here'.

P_1 that Frank lives in London.

P_2 that Frank never dines at the club.

Further, consider AND, and note that it is very plausible, too[19].

AND. For all contexts *c*, if sentence s_1 expresses that *p* at *c* while sentence s_2 expresses that *q* at *c*, then 's_1 and s_2' expresses at *c* that *p* and *q*.

Finally, AND, COM-3, and COM-4 together yield (Uni): Let *c* be a particular context and let *l* be the content that 'here' has at *c*. Then (3) and (4) express at *c* that Frank lives in *l*, and that Frank never dines in *l*, by COM-3 and COM-4. By AND, (1) hence expresses at *c* that Frank lives in *l* and never dines in *l*, and that is a uniform proposition. Since *c* was chosen arbitrarily, (Uni) follows.

In sum, SENTENCE-CONTENT, AND, COM-3, and COM-4 together imply (Uni). Thus since SENTENCE-CONTENT is central to the Kaplanian account, and since it would be *ad hoc* for the Kaplanian to reject one of the other three principles, it would be very implausible for the Kaplanian to reject (Uni).

2.2 The major argument

With (Uni) in place, return to CLUBLAND: Peter and Frank speak to another club member, who claims that all club members who live in London regularly dine at the club. Peter disagrees and utters (1).

Next, let us spell out CLUBLAND a little further. In particular, assume that Frank lives in London and that he does not live at the club. Assume further that he never dines at the club and that he often dines in London; in particular, he frequently dines at home. If the case is spelt out in this way, Peter's utterance of (1) is intuitively *true*. Thus, the

[19] Kaplan (1989a, p. 545) endorses a closely related principle: $\models_{cftw} \varphi \wedge \psi$ iff $\models_{cftw} \varphi$ and $\models_{cftw} \psi$.

Kaplanian account assigns the correct content to Peter's sentence-utterance only if it assigns a true proposition to his utterance.

By UTTERANCE-CONTENT, the content of Peter's utterance is the proposition that (1) expresses at the context in which Peter's utterance occurs. (Uni) in turn implies that (1) expresses a uniform proposition at that context. Hence Peter's utterance of (1) expresses a uniform proposition, according to the Kaplanian account. However, if CLUBLAND is spelled out as described above, all relevant uniform propositions are *false*. In particular, P_3 to P_6 are false.

P_3 that Frank lives in London and never dines in London.

P_4 that Frank lives in the club and never dines at the club.

P_5 that Frank lives at his home and never dines at his home.

P_6 that Frank lives in England and never dines in England.

Thus the Kaplanian account assigns a false proposition to Peter's utterance of (1). Consequently, the account fails to assign the correct content to Peter's sentence-utterance.

2.3 Occurrence

For later purposes, I'd like to revisit OCCUR, NOW, and HERE at this point. OCCUR is repeated for convenience.

> OCCUR. For all utterances u, u occurs in $<S, t, l, w>$ iff: S uttered u during time t in world w, and l is the smallest location that S occupies at t.

Recall that the time *during* which S utters u is the time interval that starts when S starts uttering u and ends when S has finished uttering u. Recall further that HERE and NOW state that the characters of 'now' and 'here' map each context c to the time of c

and to location of *c*. As I shall now explain, it is not very plausible to endorse OCCUR together with NOW and HERE: the Kaplanian must revise either OCCUR or NOW and HERE. Yet, importantly, neither revision will help the Kaplanian solve the MOP. Let me explain these points in turn.

To see why it is not plausible to endorse both OCCUR and HERE, suppose that while in London, I am asked about London's population; I reply: 'Over eight million people live here'. Intuitively, my utterance is true. Yet if the Kaplanian endorses both OCCUR and HERE, her account instead assigns a false content to my utterance: Let c_L be the context in which my sentence-utterance occurs. By OCCUR, the location of c_L is the smallest location I occupy when uttering the sentence. By HERE, this location is furthermore the content of 'here' at c_L. Hence by SENTENCE-CONTENT and UTTERANCE-CONTENT, my sentence-utterance expresses that over eight million people live in smallest location that I occupy when uttering the sentence, and that proposition is obviously false. Hence, the Kaplanian cannot plausibly endorse both OCCUR and HERE.

To see that the Kaplanian cannot plausibly accept OCCUR and NOW, suppose that Jane utters (2) = 'Smith now courts Mary'. Immediately afterwards Tom utters 'Jane said that Smith now courts Mary'. Then if both OCCUR and NOW are true, Jane's utterance expresses that Smith courts Mary at the time during which Jane speaks, while Tom's utterance expresses that Jane said that Smith courts Mary at the time during which Tom speaks. Yet since the times during which Jane and Tom speak are disjoint, Jane *didn't* say that Smith courts Mary at the time during which Tom speaks. Hence, given OCCUR and NOW, Tom's utterance is not an accurate speech report of Jane's utterance, contrary to our intuitions. Thus the Kaplanian cannot plausibly endorse both OCCUR and NOW[20].

[20] Admittedly, the argument that OCCUR and NOW cannot be combined is not as straightforward as the argument that OCCUR and HERE cannot be combined: to spell the argument out in detail, one would have

The arguments I just presented establish what I shall call the FLEXIBILITY RESULTS: one can use 'here' and 'now' not only to make statements about the smallest location one occupies when uttering 'here', and the time during which one utters 'now', but also to make statements about larger locations and longer time spans.

Note that the FLEXIBILITY RESULTS only show that one cannot *combine* OCCUR with NOW and HERE. They show neither that OCCUR is false, nor that NOW and HERE are false. Hence the Kaplanian could choose to retain OCCUR and revise NOW and HERE, or she could choose to revise OCCUR and retain NOW and HERE. I shall suppose that she chooses the first option, but nothing depends on this.

The Kaplanian can revise NOW and HERE in at least two different ways: firstly, she could hold that *salience* plays a role in fixing the contents of 'now' and 'here'; secondly, she could hold that *speaker intentions* play such a role. To illustrate, return to my utterance of 'Over eight million people live here', and recall that c_L is the context in which my utterance occurs. The first proposal would be that 'here' refers to London at c_L because London is maximally salient to me (to the speaker of c_L) at the time of c_L in the world of c_L[21]. The second proposal would be that 'here' refers to London at c_L because at the time of c_L and in the world of c_L, I (the speaker of c_L) intend that the utterance of 'here' refers to London. Both proposals imply, correctly, that my utterance of 'Over eight million people live here' expresses that over eight million people live in London.

For present purposes, the Kaplanian needn't settle for either proposal: it is enough that she acknowledges that the characters of 'now' and 'here' are more complex than

---

to consider the semantics of speech reports more closely. For instance, one would need to consider the view that Cappelen and Lepore (2005) defend, on which there is more than one accurate way to report what Jane said. Still, the argument shows that combining OCCUR with NOW is problematic, so that it is worth considering replacing either OCCUR or NOW.

[21] This is one of the points at which it is important that contexts contain worlds: without making reference to a particular world, one cannot meaningfully state that this-and-this location is salient to the speaker of c at the time of c.

NOW and HERE suggest. Yet it will be useful for the following discussions to spell out one of the proposals, and to suppose that the Kaplanian endorses it. I will suppose that the Kaplanian endorses the intention-based proposal, but nothing depends on this. Spelled out in detail, then, the intention-based proposal is this.

NOW*. For all contexts *c*, the character of 'now' maps *c* to the time of *c*, *unless* there is a unique time *t* such that:

(a) *t* includes the time of *c*, and

(b) during the time of *c* and in the world of *c*, the speaker of *c* produces one or more utterances of 'now' and intends the utterance(s) to refer to *t*;

in that case, the character of 'now' maps *c* to *t*.

HERE*. For all contexts *c*, the character of 'here' maps *c* to the location of *c*, *unless* there is a unique location *l* such that:

(a) *l* includes the location of *c*, and

(b) at the time of *c* and in the world of *c*, the speaker of *c* produces one or more utterances of 'here' and intends the utterance(s) to refer to *l*;

in that case, the character of 'here' maps *c* to *l*.

Given HERE* and NOW*, 'here' and 'now' are no longer fully automatic, but are still clearly distinct from demonstratives: If a speaker has suitable reference-fixing intentions when uttering 'here' or 'now', the intended referent – and not the referent the context supplies – will be the content of the relevant utterance. However, if the speaker has no (or conflicting) intentions, her utterance can still acquire a referent from context,

just as before; that latter feature sets 'here' and 'now' apart from demonstratives[22]. To coin a phrase, then, HERE* and NOW* make 'here' and 'now' *semi*-automatic.

Following the example of NOW* and HERE*, one can also provide more accurate descriptions of the characters of 'today', 'yesterday', and 'tomorrow'. E.g. the character of 'today' can be described by saying that it maps each context *c* to the day of *c*, unless the speaker of *c* produces an utterance of 'today' at the time of *c* and in the world of *c*, and intends that utterance to refer to a longer time span. That way, one can allow 'today' to have e.g. our current era as content at some contexts.

To conclude this sub-section, note that revising NOW and HERE does *not* help the Kaplanian solve the MOP; revising OCCUR wouldn't help either. For, the arguments from the previous two sub-sections do not rely on any assumptions about the particular context in which Peter's utterance of (1) occurs, or on any assumption about the content that 'here' has at that context. According to the Kaplanian account, (1) expresses a uniform proposition at *all* contexts, and hence at the context in which Peter's utterance occurs – *whatever* that context may be and *whatever* content 'here' may have at it. Thus one cannot solve the MOP by revising OCCUR, NOW, or HERE.

Before I turn to ways in which one can potentially solve the MOP, I'd like to present further instances of the problem.

3. Further instances of the MOP

In this section, I shall present instances of the MOP for indexicals other than 'here'. That one can construct these instances shows that the MOP is not specific to 'here', but rather a general challenge to theories of indexicals.

[22] Also note that HERE* and NOW* use context to *constrain* the possible referents of 'here' and 'now': if a speaker utters 'now' during time *t*, her utterance needn't refer to *t*, but NOW* guarantees that it refers to a time that *includes t*; similarly for HERE* and 'here'.

For each instance of the MOP, the argument is the same as above: a sentence is presented in which the target indexical occurs multiply, and a particular utterance of the sentence is considered. While the utterance is intuitively true, it is shown that the Kaplanian assigns a false proposition as content to the utterance.

It would be tedious to spell out the argument for each instance of the MOP in full detail. So I shall merely present the puzzling sentence-utterances, and state what content the utterance intuitively expresses, and what content the Kaplanian account assigns to the utterance.

3.1 Cases for 'now' and 'today'

THE EXHIBITION. A museum docent is explaining several different paintings to a group of visitors. Two art experts see this, but cannot hear what the docent is saying. They wonder if there is any connection between the paintings that are being explained, and one of them suggests that they may all be from the same period. His colleague disagrees and utters (5):

(5) The painting he is pointing at now is much older than the painting he is pointing at now.

Given appropriate assumptions about the ages of the two paintings, the utterance of (5) is intuitively true. However, the Kaplanian account implies that the utterance expresses a proposition of the form of P_7, and all propositions of this form are false.

P_7 that the painting the docent demonstrates at *t* is much older than the painting he demonstrates at *t*.

UNCLE SAM. On 14 May 2013, two students are talking about their research. The one asks the other: 'How is it going with your sociology paper?' The other student, Anna, replies that it is going well and utters (6).

(6) Earlier today I realised why 20% of Americans volunteer for the Army today.

To spell the case out a little, suppose that 20% of Americans volunteer for the Army nowadays, and that Anna realised this on 14 May 2013. Suppose further that there is no single day on which 20% of Americans volunteer for the Army. Then Anna's utterance is intuitively true, as it intuitively expresses that Anna realised earlier *on 14 May 2013* why 20% of Americans volunteer for the Army *nowadays*, or *in our current era*. By contrast, all propositions of the form of P_8 are false, due to the fact that 'realise' is factive.

P_8 that earlier on *d*, Anna realised why 20% of Americans volunteer for the Army on *d*.

As the Kaplanian account implies that Anna's utterance expresses a proposition of the form of P_8, it assigns the wrong content to Anna's utterance[23].

3.2 Cases for 'you' and 'we'

Let me next consider 'you'. For two reasons, the case of 'you' is slightly more complex than the previous cases. Firstly, opinions differ on whether 'you' is an indexical or a demonstrative. Kaplan states that 'you' is among the expressions he is *interested in*[24] without saying explicitly whether he takes 'you' to be a pure indexical or

[23] The following proposition is true: that Anna realised earlier *in our current era* why 20% of Americans volunteer for the Army in our current era. Yet this is not the content that Anna's utterance intuitively expresses: since she says '*Earlier* today…' we do not read her first utterance of 'today' as referring to our current era. To bring this out even more clearly, suppose that Anna replies to her fellow student's question with: 'Yesterday, I thought that my paper was going nowhere. But earlier today I realised why 20% of Americans volunteer for the Army today.' If Anna says this, it is even harder to read her first utterance of 'today' as refers to our current era. Thus, the Kaplanian cannot plausibly solve UNCLE SAM by assigning the abovementioned proposition to Anna's utterance of (6).

[24] Cf. Kaplan 1989, p. 489.

a demonstrative. Braun, for instance, lists 'you' among the demonstratives[25], while Egan treats 'you' as an indexical[26].

The case of 'you' is more complicated secondly because some argue that the singular second-person pronoun 'you' is distinct from, and merely homonymous with the plural second-person pronoun 'you'[27]. If this is correct, then sentences that contain each pronoun only once will not give rise to the MOP as the Kaplanian can assign different contents to the two pronouns at any given context.

If 'you' is a demonstrative, then so be it: my aim is to show that one can construct instances of the MOP for all, or most, indexicals. Hence, if 'you' is a demonstrative rather than an indexical, I can simply ignore 'you'[28]. But to make things a bit more interesting, let's assume for now that 'you' is an indexical.

Further, one needn't assume that the singular second-person pronoun 'you' is identical with the plural second-person pronoun 'you' in order to be able to construct instances of the MOP that involve 'you'. To see this, suppose that the two pronouns are merely homonymous and write 'you_{sg}' and 'you_{pl}' to distinguish between the two. Next, note that (7) and (8) are still well-formed sentences in English:

(7) The car you_{sg} own is faster than the car you_{sg} own.

(8) The game you_{pl} are playing is older than the game you_{pl} are playing.

[25] Cf. Braun 1996, p. 146.
[26] Cf. Egan 2009.
[27] The argument relies on fact that in languages like German and French, the two pronouns are not homonymous. I cannot discuss the argument in detail, but suspect that it is not as straightforward as sometimes assumed. German and French contain *three* non-homonymous second-person pronouns – two are informal and one is formal. Hence, since English contains *at most* two homonymous second-person pronouns, the argument can't be: if there are languages that contain *n* non-homonymous second-person pronouns, then English contains *n* homonymous second-person pronouns. On the other hand, claiming that English contains *n* homonymous second-person pronouns *only if* there is a language that contains *n* non-homonymous second-person pronouns, does not yield that English contains two homonymous second-person pronouns.
[28] The reader may be wondering whether one can construct instances of the MOP that involve demonstratives. I shall return to this question at the end of this section.

Now, if there are any true utterances of (7) and (8) – and it seems that there are – then they are instances of the MOP. For the Kaplanian account will respectively assign propositions of the forms of P_9 and P_{10} to these utterances, and all propositions of these forms are false.

P_9 that the car that person *p* owns is faster than the car that *p* owns.

P_{10} that the game that group *g* is playing is older than the game that *g* is playing.

If the two pronouns are not merely homonymous, but identical, one can construct further instances of the MOP. The following sentence is an example.

(9) Since you are Mr Smith, I assume that you are the Smith family.

It seems that one can utter this sentence truthfully, e.g. at a meet-and-greet event. Yet the Kaplanian account will imply that the utterance expresses a false proposition.

If 'you' is an indexical, perhaps 'we' is an indexical also. If it is, one can construct an instance of the MOP that involves 'we'.

CHRISTMAS TIME. Jake and Anna normally spend Christmas with another couple, Tom and Lena. This year, however, each couple will spend the holidays on their own. In an email to Tom and Lena, Jake and Anna write:

(10) We regret that we won't spend the holidays together this year.

Intuitively, Jake and Anna's utterance of (10) expresses that *Jake and Anna* regret that *Jake, Anna, Tom, and Lena* won't spend the holidays together this year. This proposition is true, but the Kaplanian account cannot assign it to Jake and Anna's utterance of (10). Instead it assigns a proposition of the form of P_{11} to the utterance.

P_{11} that group *g* regrets that group *g* won't spend the holidays together.

3.3 Cases for ‘I’

In this sub-section, I shall present two instances of the MOP for ‘I’. As will be seen, constructing these instances is much harder than it was to construct the previous instances. Further, the Kaplanian may be able to solve both instances. Yet, importantly, the solutions that the Kaplanian can give to the instances for ‘I’ cannot be extended to yield a solution to the MOP in general.

3.2.1 Multiple speakers

AMNESIA. Bill and Bob wake up in an unfamiliar room, not remembering who they are. To make their predicament worse, they have both lost their ability to speak. They have a piece of paper, however, and two pens. A while later, they find out that one of them is the governor of California while the other is the governor of Idaho, but they haven’t yet found out who is who. Bill now takes pen and paper and writes: ‘If I am the governor of California,’. Before he can write anything else, Bob writes: ‘then I am the governor of Idaho’. Hence the piece of paper now shows an inscription that looks thus:

If I am the governor of California, then I am the governor of Idaho.

Suppose for a moment that the inscription on Bill and Bob’s paper is a proper sentence-utterance, and not just an ink-mark that looks like an utterance. If this is so, the inscription is more particularly an utterance of sentence (11) = ‘If I am the governor of California, then I am the governor of Idaho’. In turn, if Bill and Bob’s inscription is an utterance of (11), it is also an instance of the MOP: their utterance expresses that Bob is the governor of Idaho if Bill is the governor of California. Yet the Kaplanian account does not assign this proposition to (11) at any context, whence it does not assign the correct content to the given sentence-utterance.

It is by no means obvious that Bill and Bob's inscription is a sentence-utterance. To decide whether it is, one would need to know more generally what it is for an inscription *i* to be an utterance of expression *e*, what it is for *i* to be a *token* of a particular *type*. Yet this is a difficult and controversial question. According to one view, expression types are a kind of *pattern*, and *i* is an utterance of *e* if and only if *i* instantiates pattern *e*. To instantiate *e*, *i* must *look* a certain way[29]. According to another view, *i* is an utterance of *e* if and only if the speaker *intends* to utter *e* when producing *i*. What *i* looks like and whether *i* resembles other utterances of *e* is irrelevant on this view[30]. Finally, Cappelen holds that it is a matter of *convention* whether *i* counts as an utterance of *e*[31]. Thus, it is by no means clear what it is for *i* to be an utterance of *e*, and hence not easy to decide whether the inscription on Bill and Bob's paper is an utterance of (11).

If the pattern-view is correct, Bill and Bob's inscription is an utterance of (11): the inscription looks the same as any ordinary written utterance of (11), so the inscription presumably does instantiate the relevant pattern.

If the intentional view is correct, matters are more complicated: Bill presumably does not intend to utter (11), but rather intends to utter 'If I am the governor of California then *you are* the governor of Idaho'. Bob also doesn't intend to write down *all* of (11), but he plausibly intends to *turn* what Bill has written into an utterance of (11). In that sense, he does intend to produce an utterance of (11). Whether this implies, on the intentional view, that the inscription is an utterance of (11) is not clear, however.

Turning to Cappelen's view, the question is whether there is a convention to count inscriptions like Bill and Bob's as sentence-utterances. More generally, the question is whether we have linguistic conventions that apply to inscriptions that are produced by

[29] This view is sometimes ascribed to Goodman and Quine (1947).
[30] Incidentally, Kaplan (1990) defends this view.
[31] See Cappelen 1999.

multiple speakers. Perhaps we do: Imagine a language student completing a cloze test. The exercise book reads 'Schnee…' and the student writes 'ist weiß'. (Snow is white.) Then we do treat the completed test as a sentence-utterance. E.g. it would be natural for the teacher to point at the completed test and say 'Please translate *that sentence*' – meaning the sentence that is inscribed in the exercise book. This suggests – but certainly doesn't prove – that we count some inscriptions which were produced by multiple speakers as sentence-utterances. Further, Cappelen claims "that the classification of entities as sign tokens is about […] how the classified entities can be used linguistically"[32]: if an inscription functions like a sentence-utterance "for all linguistic purposes"[33], then it is a sentence-utterance. And Bill and Bob's inscription at least appears to function like a sentence-utterance: it appears to express a proposition. Thus the least we can say is that a version of Cappelen's view is imaginable on which Bill and Bob's inscription counts as a sentence-utterance.

In sum, each of the three views is at least *compatible* with the claim that Bill and Bob's inscription is an utterance of (11). On the other hand, neither the intentional view nor Cappelen's view *implies* that the inscription is an utterance of (11). Hence in response to AMNESIA, the Kaplanian could adopt either of these latter views and try to spell them out so that they yield that the inscription is not an utterance of (11). Thus while AMNESIA is a challenge to the Kaplanian, she may be able to meet the challenge.

There is one important caveat: If the Kaplanian adopts the intentional view or Cappelen's view, she can solve *one instance* of the MOP, viz. AMNESIA. Yet neither view will provide a solution to the MOP *in general*: on all three views, Peter produces a sentence-utterance when he says: 'Frank lives here and Frank never dines here'.

[32] Cappelen 1999, p. 96.
[33] Cappelen 1999, p. 96.

3.2.2 Descriptive uses of 'I'

Let me present one more puzzle case involving 'I'. Suppose that Obama utters the following sentence:

(12) Traditionally, I live in the White House.

On its most natural reading, Obama's utterance expresses a general proposition – that *the US President* traditionally lives in the White House – rather than a singular proposition – that *Obama* traditionally lives in the White House. Hence at least prima facie, Obama's utterance of 'I' does not refer to Obama, but has roughly the same semantic value as 'the US President': prima facie, Obama's utterance of 'I' is not directly referential but has a descriptive content[34].

Grant for a moment that Obama's utterance of (12) indeed expresses the said general proposition, and that there are utterances of 'I' that are not directly referential. Now consider the following case.

THE PRESIDENT. It is December 2009 and Obama is about to celebrate his first Christmas as President. Assume that the US President is traditionally served lobster on Christmas Day: every year since 1789, lobster is served on Christmas Day to the man who is the President that year. Assume further that Obama has so far always eaten duck on Christmas Day; there is no tradition to serve lobster to Obama on Christmas Day. Finally suppose Obama utters (13).

(13) Although I had duck last year, I am traditionally served lobster on Christmas Day.

Now, if Obama's utterance of (12) expresses a general proposition, then it is plausible that Obama's utterance of (13) expresses P_{12}:

[34] Nunberg (1993) considers the contrast between referential and descriptive uses of 'I' in much more detail. Sentence (12) is inspired by Nunberg's examples.

P_{12} that the US President is traditionally served lobster on Christmas Day, although Obama had duck on Christmas Day in 2008.

As the case was spelled out, P_{12} is true. Yet the Kaplanian account does not assign P_{12} as content to Obama's utterance of (13). Instead, it assigns a false proposition, P_{13}, to his utterance:

P_{13} that Obama is traditionally served lobster on Christmas Day, although Obama had duck on Christmas Day in 2008.

It won't help the Kaplanian to adjust the character of 'I': she could claim that at some contexts *c*, 'I' expresses roughly the same as 'the (current) US President'. Then she can assign P_{14} as content to Obama's utterance, but P_{14} is also false.

P_{14} that the (current) US President is traditionally served lobster on Christmas Day although the (current) US President had duck on Christmas Day in 2008.

Hence the problem is the same as before: The sentence-utterance intuitively expresses a true proposition. Yet the Kaplanian account can only assign false propositions to the utterance. Thus, THE PRESIDENT is an instance of the MOP involving 'I'.

Let's consider how the Kaplanian could respond to THE PRESIDENT. As I pointed out, Obama's utterance of (12) *intuitively* expresses a general proposition. One can take this intuition at face value and hold that the utterance indeed has a general proposition as semantic content. Then one cannot plausibly deny that Obama's utterance of (13) expresses P_{12}. However, one can also hold that the intuition about (12) is misleading, and that 'I' is always directly referential. Obama's utterance of (12) semantically expresses a singular proposition on this view, and the general proposition is merely pragmatically communicated. If this view is correct, the Kaplanian account assigns the correct content to Obama's utterance of (13). For, on the view at hand, P_{13} is the

semantic content of Obama's utterance of (13). P_{12} is merely pragmatically communicated by the utterance.

In sum, the Kaplanian could respond to THE PRESIDENT by arguing that 'I' is always directly referential. King (2006) argues for this conclusion, but due to limitations of space, I cannot consider his arguments here.

The caveat I issued regarding AMNESIA applies to THE PRESIDENT also: if the Kaplanian can defend the claim that 'I' is always directly referential, she can solve *one instance* of the MOP, but the claim in question cannot be extended in any way to yield a solution to the MOP *in general*: that 'I' is always directly referential does not imply e.g. that Peter's utterance of (1) semantically expresses a uniform proposition, nor does it allow the Kaplanian to assign a non-uniform proposition to Peter's sentence-utterance. That 'I' is always directly referential may mean that 'here' is also always directly referential, but that is irrelevant: there was never any doubt that both of Peter's utterances of 'here' are directly referential.

To summarise, one can construct instances of the MOP for 'I', but it is not clear how much of a challenge they are for the Kaplanian: for all that's been said, the Kaplanian can solve the MOP's instances for 'I'.

3.3 Demonstratives again

Above I mentioned that Kaplan presents two theories: one is intended as a theory of indexicals, the other as a theory of demonstratives. And while the two theories resemble each other in some respects, they are substantially different. Now, in his *Demonstratives and Their Linguistic Meanings* (Braun 1996), Braun considers Kaplan's theory of demonstratives and raises a problem for it which is closely related to the MOP:

Kaplan's theory of demonstratives cannot assign the correct content to true utterances of sentences like 'That is not that'.

I would like to make three points about the relation between Braun's problem and the MOP. Firstly, although the two problems are clearly related, Braun's problem is not a problem for Kaplan's theory of indexicals, and the MOP is not a problem for Kaplan's theory of demonstratives. The reason is simply that the former theory does not apply to sentences like 'That is not that'; the latter theory does not apply to sentences like (1). In other words, the question of whether Kaplan's theory of indexicals (demonstratives) assigns the correct content to utterances of e.g. 'That is not that' (to utterances of e.g. (1)) does not even *arise*, and one therefore cannot argue that the theory assigns the wrong content to the utterances in question.

Secondly, since Kaplan's two theories are different, they have different resources at their disposal when it comes to solving the problems that they face. For instance, Braun's preferred solution to the problem he raises heavily relies on the idea that utterances of demonstratives are associated with demonstrations[35]. This means that it is not at all clear that Braun's solution can be converted into a solution to the MOP and incorporated into Kaplan's theory of indexicals. For, according to Kaplan's theory of indexicals, utterances of indexicals are not always associated with demonstrations.

Thirdly, one could respond to the MOP by rejecting the distinction that Kaplan draws between indexicals and demonstratives: one could abandon his theory of indexicals completely, and hold that his theory of demonstratives is correct as a theory of the semantics of indexicals, too. Indexicals are then no longer automatic, and the MOP becomes a problem for the theory of demonstratives, as that theory then applies to sentences like (1). The hope would be that the theory of demonstratives can solve the MOP, i.e. that the theory can assign the correct content e.g. to Peter's utterance of (1).

[35] See Braun 1996, sections five to seven.

Yet, in the present paper, I shall not pursue this line of thought: in the remaining sections, I shall rather investigate whether one can develop a solution to the MOP that espouses that indexicals are (semi-)automatic.

PART II: TWO POTENTIAL SOLUTIONS TO THE MOP

In this part of the paper, I shall survey two potential solutions to the MOP. They have in common that they replace UTTERANCE-CONTENT with an analysis of utterance-content that does not imply or presuppose UNIQUENESS. Section four presents a fairly straightforward analysis of this kind. I will call it the *Multiple Contexts Analysis*, or MCA. The MCA looks promising at first as a solution to the MOP, but fails in the end. Section five surveys the so-called Token-Reflexive Theory of indexicals. While this theory can solve the MOP, the Kaplanian has several reasons for rejecting the theory. A third solution to the MOP will be discussed in detail in the final part of the paper.

4. The Multiple Contexts Analysis

4.1 The MCA and the MOP

The MCA is intended as an alternative to UTTERANCE-CONTENT. It consists of two parts, the first of which assigns contents to word-utterances.

> MCA (W). For all word-utterances *u*, the content of *u* is the content that word *w* has at the context in which *u* occurs.

Note that this principle uses the notion of occurrence. As before, this notion is analysed using OCCUR.

Next, if *u* is an utterance of sentence σ, let us speak of the *constituent utterances* of *u* to mean those sub-utterances of *u* that are an utterance of one of the constituent words

of σ. E.g. if John utters 'Frank is tall', the constituent utterances of his sentence-utterance are his utterances of 'Frank', 'is', and 'tall'. His utterance of the VP 'is tall' is not a constituent utterance of his sentence-utterance.

The second part of the MCA assigns contents to sentence-utterances: Given an utterance *u* of a sentence σ, one first evaluates the constituent utterances of *u* for content, using MCA (W). One then refers to the compositional structure of σ and combines the contents of the constituent utterances of *u* to form the content of *u*. This can be summarised thus:

> MCA (S). For all utterances *u* and sentences σ, if *u* is an utterance of σ, then the content of *u* is determined by σ's compositional structure and the contents of *u*'s constituent utterances.

One can describe the difference between UTTERANCE-CONTENT and the MCA as follows. Given a particular sentence-utterance *u*, UTTERANCE-CONTENT assigns a content to *u without* first assigning contents to *u*'s constituent utterances. By contrast, the MCA *first* assigns contents to *u*'s constituent utterances, and then builds the content of *u* out of those contents.

Suppose the Kaplanian replaces UTTERANCE-CONTENT with the MCA. Then her account assigns the correct content to Peter's utterance of (1). To see this, suppose that 'here' is the only context-sensitive expression in (1). Let c_1 and c_2 be the contexts in which Peter's utterances of 'here' occur[36], and suppose that London is the content of 'here' at c_1 while the club is the content of 'here' at c_2[37]. By MCA (W), London and the club are hence the contents of Peter's utterances of 'here'. Given MCA (S), it follows that Peter's utterance of (1) expresses that Frank lives in London and never dines at the

[36] Peter's two utterances of 'here' occur in different contexts because the times during which Peter produces the utterances are disjoint.

[37] Given HERE*, this will be so provided that Peter intends his utterances of 'here' to refer to London and the club. That Peter has these intentions is certainly very plausible.

club. This is not only a non-uniform proposition; it is also the content that Peter's utterance of (1) intuitively expresses. Thus, with the MCA in place, the Kaplanian can assign the correct content to Peter's utterance of (1) and can hence solve CLUBLAND.

Speaking more generally, the MCA solves each instance of the MOP by treating each utterance of the multiply occurring indexical *separately*: it considers each utterance of the relevant indexical in the context in which it occurs, and can hence assign different contents to different utterances. All of these contents then 'go into' the content of the sentence-utterance. Therefore, using the MCA, one can assign the correct contents to the puzzling sentence-utterances.

Note that the MCA sidesteps UNIQUENESS: MCA (W) and MCA (S) don't presuppose UNIQUENESS, but as far as the MCA is concerned, it is *irrelevant* whether each sentence-utterance occurs in exactly one context: to assign a content to a sentence-utterance, the MCA only considers the contexts in which the sentence-utterance's constituent utterances occur. It therefore doesn't matter whether the sentence-utterance itself occurs in exactly one context or in several different contexts.

4.2 The MCA and sentences in context

Let me raise a problem for the MCA. The current proposal is to replace UTTERANCE-CONTENT with the MCA, while leaving the rest of the Kaplanian theory in place. Hence if the Kaplanian follows this proposal, her theory will contain both the MCA and SENTENCE-CONTENT. However, one cannot plausibly combine the two principles in question. To see this, return to (1). SENTENCE-CONTENT implies (Uni)[38], so the Kaplanian is committed to the claim that (1) expresses a uniform proposition at every context. At the same time, the MCA predicts that some utterances of (1), like Peter's, do

[38] More precisely, SENTENCE-CONTENT *together with* AND, COM-3, and COM-4 implies (Uni). But as explained above, the Kaplanian cannot plausibly reject any of the latter three claims. Thus as long as the Kaplanian endorses SENTENCE-CONTENT, she is committed to (Uni).

not express a uniform proposition. Thus if the Kaplanian endorses both the MCA and SENTENCE-CONTENT, she is stuck with the fact that some *utterances* of (1) do not express a non-uniform proposition, although (1) expresses a uniform proposition at every *context*. In other words, she is stuck with two different theories: one is an utterance-content-theory, the other is a sentence-in-context-theory, and the two just don't line up.

Since the two theories don't line up, it is rather unclear what the *purpose* of the sentence-in-context-theory is. Contexts are supposed to represent circumstances of utterance, and sentence-context pairs are supposed to represent utterances. Yet if there is an utterance of (1) that expresses a content which (1) doesn't express at any context, then contexts no longer represent circumstances of utterance, and sentence-context pairs no longer represent utterances. So it is unclear what the purpose of a theory is that evaluates sentences relative to contexts. Thus, as long as the Kaplanian endorses SENTENCE-CONTENT, she cannot plausibly endorse the MCA.

One might reply that contexts still represent circumstances of utterance, and that sentence-context pairs still represent utterances – except that there are utterances and circumstances of utterance that are not represented by any sentence-context pair or any context. Accepting this still doesn't answer the question of what the purpose of the sentence-in-context-theory is, and why the Kaplanian would endorses it: if there are circumstances of utterance that aren't represented by any context, the Kaplanian should refine her notion of context so that all circumstances of utterance are again represented by some context; and she should replace SENTENCE-CONTENT and develop a new sentence-in-context-theory which assigns a non-uniform proposition to (1) relative to at least one of the new, refined contexts. In the third part of the paper, I shall present a theory that embraces this idea. Before I turn to presenting that theory, I'd like to consider the Token-Reflexive Theory.

5. The Token-Reflexive Theory

The Kaplanian account – in its original version and in its MCA-version – assigns contents both to sentence-utterances and to sentences in contexts. By contrast, the Token-Reflexive Theory only assigns contents to (word- and sentence-)utterances. It doesn't assign contents to sentences in contexts at all. Therefore, the Token-Reflexive Theory is not a modification to Kaplan's theory, but rather a radical alternative.

5.1 The Token-Reflexive Theory and the MOP

It is beyond the scope of this paper to describe the Token-Reflexive Theory in full detail[39]. I shall therefore merely present the claims that the theory makes about (1), and show that these claims allow the Token-Reflexive Theory to assign the correct content to Peter's utterance of (1). To start, the theory assigns contents to utterances of indexicals. In particular, utterances of 'I', 'now' and 'here' are treated thus[40]:

TRT_I. For all utterances *u* of 'I', the content of *u* is the person who utters *u*.

TRT_{NOW}. For all utterances *u* of 'now', if *S* utters *u* during *t*, then the content of *u* is *t*, *unless* there is a unique time *t** that includes *t* and is intended by S as the referent of *u*. In that case, the content of *u* is *t**.

TRT_{HERE}. For all utterances *u* of 'here', if *S* utters *u* during *t*, then the content of *u* is the smallest location *l* that *S* occupied during *t*, *unless* there is a unique location *l** that includes *l* and is intended by S as the referent of *u*. In that case, the content of *u* is *l**.

[39] See e.g. García-Carpintero (1998) for a detailed description.

[40] It is sometimes assumed that the Token-Reflexive Theory endorses the following principle for 'now'.

TRT_{Now}*. For all utterances *u* of 'now', the content of *u* is the time during which *u* is uttered.

E.g. Predelli (2006) describes a version of the theory that comprises TRT_{Now}*. Yet TRT_{Now}* does not respect the FLEXIBILITY RESULTS: TRT_{Now}* implies that one can use 'now' only to make statements about the time during which one utters 'now', and this is false as we saw. Therefore, I think that the best version of the Token-Reflexive Theory will endorse TRT_{Now} rather than TRT_{Now}*.

Given TRT_{Here}, utterances of (1) can be assigned a content as follows:

$\text{TRT}_{(1)}$. For all utterances *u* of (1), *u* expresses that Frank lives in l_1 and never dines in l_2, where l_1 and l_2 are the contents of the utterances of 'here' that are part of *u*.

Thus instead of endorsing a general principle like UTTERANCE-CONTENT, the Token-Reflexive Theory puts forth a separate principle for each indexical, and for each sentence in the language. Note that the notion of occurrence is not used in these principles. So the Token-Reflexive Theory has no need for statements like OCCUR.

To apply $\text{TRT}_{(1)}$ to Peter's utterance of (1), grant that Peter intends his utterances of 'here' to refer to London and the club. By TRT_{Here}, London and the club are then the contents of his two utterances of 'here'. $\text{TRT}_{(1)}$ hence implies that Peter's utterance of (1) expresses that Frank lives in London and never dines at the club. Thus, the Token-Reflexive Theory can assign the correct content to Peter's utterance of (1) and can hence solve CLUBLAND. The remaining instances of the MOP are solved in the same way.

5.2 Against token-reflexivity

As explained, the Token-Reflexive Theory does not assign contents to sentences in contexts. Yet many theorists in the Kaplanian tradition hold that this is a substantial flaw in the theory: they argue that a satisfying semantic theory of indexicals should assign contents both to utterances and to sentences in contexts. Those who subscribe to this argument conclude that the Token-Reflexive Theory is inferior to the Kaplanian

theory[41]. García-Carpintero, among others, tries to defend the Token-Reflexive Theory against these arguments[42].

A detailed comparison of the Kaplanian theory with the Token-Reflexive Theory provides enough material for a separate paper, so I cannot attempt such a comparison here. I shall restrict myself to presenting three objections to the Token-Reflexive Theory. The first two objections are often raised in the literature, but I don't find them entirely convincing. As far as I know, the third objection is not discussed in the literature, but I believe it to be more successful than the other two objections.

5.2.1 Impossible utterances

The first objection has to do with the interpretation of extremely long sentences. Unfortunately, it is not easy to make out what the argument is supposed to be that drives the objection. Here is Braun's formulation of the objection:

> "Consider a sentence consisting of 'I exist now' conjoined with itself ten trillion times; and consider a context which has me as its agent and which lasts less than one minute (i.e., the time assigned to 'now' by the context is less than one minute). [The Kaplanian theory] can assign a content to this sentence in this context […]. But there is no (actual or possible) utterance of this sentence to which [the Token-Reflexive Theory] can assign this content. For I have never actually uttered this sentence; and it is not even possible for me to utter this sentence within the time frame of the context." (Braun 1996, p. 152)

Let LONG be the sentence that Braun considers and let c_B be a context that has Braun as agent and the first second of 2014 as time. Then Braun's first claim is that the Kaplanian theory can assign a content to LONG at a context like c_B. To be more specific, PRO is plausibly the content of LONG at c_B:

PRO: the proposition that is the conjunction of the following proposition with itself ten trillion times: that Braun exists during the first second of 2014[43].

[41] E.g. Braun (1996) argues for this. More recently, Predelli (2006) has argued for the same conclusion.
[42] See his 1998.

Braun's second claim is that it is impossible for him to utter LONG during a second-long time interval, so that PRO isn't the content of any (actual or possible) utterance of LONG. I'm not sure that it is impossible for Braun to utter LONG during a second-long time interval[44], but let's grant that it is.

Note that the Token-Reflexive Theory is consistent with both of Braun's claims, so we don't yet have an objection to the theory. But consider the following argument.

(i) LONG has a content relative to c_B.

(ii) The Token-Reflexive Theory doesn't assign a content to LONG at c_B.

(iii) Thus, the Token-Reflexive Theory misses something about the semantics of LONG.

I don't think that this argument is a convincing objection to the Token-Reflexive Theory, nor do I think that the argument is the objection that Braun has in mind. I don't find the argument convincing because of premise (i). If (i) is supposed to mean that *according to the Kaplanian theory*, LONG has a content at c_B, then (i) is true. Yet then (i) and (ii) no longer imply (iii): that one semantic theory implies that *p* doesn't mean in and of itself that semantic theories that don't imply *p* are flawed. On the other hand, if (i) is saying that there is some *pre-theoretic*, intuitive reason to think that LONG has a content relative to c_B, then (i) is false. For, we simply don't have pre-theoretic intuitions about *theoretic* entities like c_B[45]. Either way, the argument is not convincing.

I don't think that the above argument is the objection that Braun has in mind, firstly because one doesn't need to consider a sentence like LONG to run the argument: the

---

[43] We can be neutral about whether the conjunction of a proposition *p* with itself is distinct from *p*.

[44] Perhaps it's possible to build a giant computer screen that can display all of LONG at once. It may also be possible to program a computer to display LONG on the screen whenever this-and-this button is pressed. And if pressing the button counts as uttering LONG, then it *is* possible for Braun to utter LONG during a second-long time interval.

[45] To illustrate this, recall that Kaplan holds that expressions can only be evaluated at proper contexts. (*c* is proper iff the speaker of *c* is located in the location of *c* at the time of *c* in the world of *c*.) Now, there is just no *intuition* about whether this is correct, no intuition about whether sentences have contents relative to improper contexts. It is a *fact* that there are functions that map improper contexts to propositions, and it is also a *fact* that there are functions that have only proper contexts in their domain. But there is no linguistic intuition about whether the 'true character' of a given sentence is a function of the former or the latter kind.

Token-Reflexive Theory also doesn't assign a content to e.g. 'I exist now' at c_B. So if Braun had the above argument in mind, why did he bother considering LONG? Secondly, Braun claims that it is impossible to utter LONG during a second-long time interval, but this claim does not feature in the argument: if Braun had the above argument in mind, why did he bother making a claim about the utterances of LONG? Hence, I think that Braun has a more complex objection in mind. Let me try to develop such an objection.

First, given an indexical-containing sentence σ, recall that there is no such thing as *the* content of σ. To put it the other way around, each indexical-containing sentence is 'semantically associated' with multiple contents and has multiple 'admissible interpretations'. In some appropriate sense of 'can', each indexical-containing sentence 'can' express different contents. E.g. (2) is associated with the content that Smith courts Mary during the first second of 2014, and also with the content that Smith courts Mary during the last minute of 2005, but not with the content that snow is white.

Next, the idea that each indexical-containing sentence has multiple admissible interpretations is spelt out differently by the Kaplanian theory and the Token-Reflexive Theory. The Kaplanian theory says: proposition *p* is an admissible interpretation of sentence σ if and only if *p* is in the range of the character of σ. By contrast, the Token-Reflexive Theory says: proposition *p* is an admissible interpretation of sentence σ if and only if *p* is expressed by at least one (actual or possible) utterance of σ. On both accounts, e.g. that Smith courts Mary during the first second of 2014 is an admissible interpretation of (2); that snow is white isn't an admissible interpretation of (2) on either account. Prima facie, these results are correct.

Now return to LONG and PRO. Since PRO is in the range of the character of LONG, the Kaplanian theory's analysis yields that PRO is an admissible interpretation of

LONG. Yet since there is no actual or possible utterance of LONG that expresses PRO, the Token-Reflexive Theory's analysis yields that PRO is not an admissible interpretation of LONG. Thus, one might argue against the Token-Reflexive Theory as follows:

(i) Judged pre-theoretically, PRO is an admissible interpretation of LONG.

(ii) The Kaplanian theory yields that PRO is an admissible interpretation of LONG.

(iii) The Token-Reflexive Theory yields that PRO is *not* an admissible interpretation of LONG.

(iv) Thus, the Token-Reflexive Theory is flawed.

The argument at hand is more likely than the previous argument to be the objection that Braun has in mind. Firstly, one can't run the argument with 'I exist now'. Secondly both of Braun's claims feature in the argument, if only indirectly: his first claim is needed to establish (ii) and his second claim is needed to establish (iii). The question is whether the argument is convincing and more particularly whether (i) is true.

Premise (i) presupposes that we can make a pre-theoretic judgement about the admissible interpretations of LONG. Let's consider what such a judgement might be based on. If there is an actual utterance of a given sentence σ that expresses *p*, then this is evidence that *p* is an admissible interpretation of σ. Similarly, if one is able to describe a (counterfactual) scenario in which an utterance of σ would express *p*, then this is also evidence that *p* is an admissible interpretation of σ. However, the claim that PRO is an admissible interpretation of LONG couldn't be supported by either kind of evidence: by assumption, no (actual or possible) utterance of LONG expresses PRO.

We do have a pre-theoretic understanding of indexicality, and it follows from this understanding that indexical-containing sentences have multiple admissible interpretations: we understand that there is no such thing as *the* content of σ, for any

indexical-containing sentence σ, which arguably means that every indexical-containing sentence is associated with multiple contents. So, pre-theoretic considerations imply *that* indexical-containing sentences have multiple admissible interpretations. However, this doesn't mean that pre-theoretic considerations suffice to *identify* all admissible interpretations of every indexical-containing sentence. Thus the onus is on the proponent of (i) to show that our pre-theoretic understanding of indexicality yields that PRO is an admissible interpretation of LONG. For all that's been said, the proponent of the Token-Reflexive Theory can simply reject (i) and deny that there is any pre-theoretic reason to think that PRO is an admissible interpretation of LONG. Thus if the above argument is indeed Braun's objection, then I conclude that Braun's objection is not convincing.

In sum, although some authors in the Kaplanian tradition argue that unutterably long sentences give rise to an objection to the Token-Reflexive Theory, neither version of the objection that was considered is convincing.

5.2.2 Validity

To introduce the next objection, recall from section one that a good semantic theory should provide an adequate definition of validity. Now, proponents of the Token-Reflexive Theory agree with the proponents of the Kaplanian theory that e.g. 'If John is happy now, then John is happy now' is valid. It is further agreed that e.g. 'An utterance exists now' is not valid. However, proponents of the Kaplanian theory claim that the Token-Reflexive Theory cannot provide a definition of validity that yields these results. That is the second objection to the Token-Reflexive Theory.

Recall how the Kaplanian defines validity, and note that this definition yields, correctly, that (14) is valid and that (15) is invalid:

VAL_{K}. For all sentences σ, σ is valid iff: for all contexts *c*, the proposition that σ expresses at *c* is true in the world of *c*.

(14) If John is happy now, then John is happy now.

(15) An utterance exists now.

Turning to the Token-Reflexive Theory, suppose that the following definition of validity is added to the theory:

VAL_{TRT}*. For all sentences σ, σ is valid iff: for all (actual and possible) utterances *u* of σ, the proposition that *u* expresses is true in the world in which *u* occurs.

Predelli suggests that this definition is the most "obvious"[46] choice for the proponent of the Token-Reflexive Theory. Further, I think that it is a definition like VAL_{TRT}* that authors have in mind when objecting to the Token-Reflexive Theory. And VAL_{TRT}* is indeed objectionable, as it does not yield the correct results for e.g. (14) and (15). Starting with (14), simply note that the Token-Reflexive Theory implies that there are utterances of (14) that express a proposition of the form of P_{15}.

P_{15} that John is happy at t_2 if John is happy at t_1.

This means that there are false utterances of (14), whence VAL_{TRT}* falsely implies that (14) is not valid.

As for (15), TRT_{Now} implies that the content of any utterance of 'now' is a time that includes the time at which the utterance was produced. Further, any utterance of (15) expresses that an utterance exists at *t*, where *t* is the content of the relevant utterance of 'now'. Hence every utterance of (15) expresses a true proposition, whence VAL_{TRT}* falsely implies that (15) is valid.

[46] Predelli 2006, p. 13.

Thus, the Kaplanian concludes, the Token-Reflexive Theory is inferior to the Kaplanian theory because the Token-Reflexive Theory's definition of sentence-validity yields false results about sentences like (14) and (15).

The proponent of the Token-Reflexive Theory needn't endorse VAL_{TRT}*: I believe she can give a definition of validity that yields the correct results for both (14) and (15). To develop this definition, return to TRT_{Now}. This principle implies that an utterance of 'now' can refer to a time that is *longer* than the time that it took to produce the utterance. E.g. if it takes one second to produce a particular utterance of 'now', the utterance can still refer to a minute-long time span, provided that the speaker has the appropriate intentions. Consequently, it is possible to produce *co-referential* utterances of 'now': e.g. one can utter 'If Smith is in London now, then Jones is in London now, too' in such a way that both utterances of 'now' refer to the day of the utterance, say. Similarly for all other indexicals: given an indexical *i*, it is possible to produce two or more co-referential utterances of *i*[47]. Given this, let us make the following definition.

> CO-REF. For all sentence-utterances *u*, *u* is *indexically co-referential* iff for all indexicals *i*, *u* was produced in such a way that all utterances of *i* that are part of *u* (if any) are co-referential.

To illustrate CO-REF, suppose I say 'Either John is here now or Jane is here now.' For my sentence-utterance to be indexically co-referential, I must produce it in such a way that my two utterances of 'here' co-refer, and further in such a way that my two utterances of 'now' co-refer. (Of course, it is not required that my utterances of 'here' co-refer with my utterances of 'now'.) To achieve this, I must have the appropriate intentions: there must be a time *t* which includes the times during which I produce my

[47] I haven't stated the Token-Reflexive Theory's rules for 'today', 'tomorrow', and 'yesterday', but it should be obvious that these rules will imply that it is possible to produce co-referential utterances of 'today', of 'tomorrow', and of 'yesterday'.

utterances of 'now', and which I intend to be the referent of the utterances; similarly for 'here'.

Further, as CO-REF is intended, all utterances of any indexical-free sentence are indexically co-referential. (E.g. every utterance of 'John is English' is indexically co-referential.) Similarly, if a sentence contains no indexical more than once, all its utterances are indexically co-referential. (E.g. every utterance of 'I am here now' is indexically co-referential.)

Using CO-REF, the proponent of the Token-Reflexive Theory can provide a definition of validity that implies that (14) is valid and that (15) is invalid[48]:

VAL$_{TRT}$: For all sentences σ, σ is valid iff: for all (actual and possible) utterances *u* of σ, if *u* is indexically co-referential, then *u* expresses a necessarily true proposition.

To see that VAL$_{TRT}$ correctly implies that (14) is valid, simply note that all indexically co-referential utterances of (14) express a proposition of the form of P_{16}.

P_{16} that John is happy at *t* if John is happy at *t*.

Since any proposition of the form of P_{16} is necessarily true, it follows that all indexically co-referential utterances of (14) express a necessarily true proposition. Thus, VAL$_{TRT}$ correctly implies that (14) is valid.

As for (15), since this sentence contains 'now' only once, all its utterances are trivially indexically co-referential. Yet no utterance of (15) expresses a necessarily true proposition. Thus, VAL$_{TRT}$ correctly implies that (15) is not valid.

Note that VAL$_{TRT}$ is neutral about why a given sentence-utterance is indexically co-referential. Suppose e.g. that, contrary to TRT$_{Now}$, speaker intentions play no role in

[48] Dorr (forthcoming) and Richard (1993, p. 144) advocate defining sentence-validity in terms of necessarily true propositions. Both authors remain neutral on what the relevant sense of necessity is and I shall follow them. Field (2008, cap. 2) argues that sentence-validity should not be defined in terms of necessary truth.

determining the contents of utterances of 'now'. In particular, suppose that an utterance of 'now' refers to the time during which it was produced unless a longer time is *salient* to the speaker. That won't change the fact that (14) *has* indexically co-referential utterances, and that these utterances express necessarily true propositions. Hence VAL_{TRT} implies that (14) is valid, regardless of whether TRT_{Now} is true.

To conclude this section, consider sentences of the form of 'If *p* then actually *p*' and note that these sentences pose a challenge to VAL_{TRT}: no utterance of a sentence of this form expresses a necessarily true proposition, whence VAL_{TRT} implies that the sentences are not valid. However, one might think that the sentences are in fact valid, which would mean that VAL_{TRT} yields the wrong result.

While a more detailed discussion would be in order, I believe that the proponent of VAL_{TRT} can deny that the sentences in question are valid. The thought would be the following: there can be no doubt that sentences of the form of 'If *p* then actually *p*' are 'special' in some way. However, being valid is just one way of being 'special'. Hence one can allow that the sentences are 'special', and still deny that they are valid. For instance, one can ascribe to them a special *epistemic* status: the sentences are 'special' in that all their indexically co-referential utterances express a content that we *know*, possibly even a content that we know *a priori*[49].

Due to limitations of space, I cannot discuss this line of thought any further. The least we can say is this: when it comes to defining validity, the Token-Reflexive Theory fairs much better than proponents of the Kaplanian theory suggest: the theory can provide a definition of validity that yields the correct results for sentences like (14) and (15); and it is at least arguable that the definition also yields the right results for sentences of the form of 'If *p* then actually *p*'. The next section considers if the Token-Reflexive Theory can provide an adequate definition of synonymy.

[49] E.g. Davies and Humberstone (1980) argue for this latter conclusion.

5.2.3 Synonymy

Pre-theoretically, synonymy is sameness of meaning. As explained, the Kaplanian captures this by defining synonymy as sameness of character – character being one of the levels of 'meaning' that the Kaplanian acknowledges. Yet since the Token-Reflexive Theory does not recognise characters, the theory's proponent must provide an alternative definition of synonymy. As I shall now argue, it is not clear that she can provide such a definition.

5.2.3.a Content-rules

According to the Token-Reflexive Theory, utterances or tokens are the primary bearers of semantic properties. Still, the theory "incorporate[s] the point that […] linguistic conventions […] are associated with repeatable entities, types and not tokens"[50]. Put slightly differently, the proponent of the Token-Reflexive Theory acknowledges that expression *types* have a kind of 'meaning'. More particularly, she holds that the meaning of an expression type *e* is given by the rule that assigns a content to each utterance of *e*. E.g. TRT_I gives the meaning of 'I': *the meaning of 'I' is such that* each utterance of 'I' refers to the person who produced the utterance. Similarly, TRT_{Now}, TRT_{Here}, and $TRT_{(1)}$ give the meanings of 'now', 'here', and (1). With this in mind, let us consider how synonymy could be defined within the Token-Reflexive Theory.

Let us speak of a *content-rule* to mean rules like TRT_I, TRT_{Now}, TRT_{Here}, and $TRT_{(1)}$. Since according to the Token-Reflexive Theory, content-rules give the meaning of the expression type with which they are associated, it is natural to think that the proponent

[50] García-Carpintero 1998, p. 534.

of the Token-Reflexive Theory can define synonymy in terms of content-rules. In particular, the following definition suggests itself.

SYN_{TRT}*. For all expressions e_1 and e_2, e_1 and e_2 are synonymous iff the content-rule of e_1 is also the content-rule of e_2.

While SYN_{TRT}* may seem obvious, it is inadequate. For, many complex expressions are, or can be, associated with more than one rule, so that *the* content-rule of *e* does not always exist. For instance, both R_1 and R_2 are suitable content-rules for (16), in that they assign the same content to any given utterance of (16).

(16) John is now a policeman.

R_1. For all utterances *u* of 'John is now a policeman', *u* expresses that John is a policeman at *t*, where *t* is the content of the utterance of 'now' that is part of *u*.

R_2. For all utterances *u* of 'John is now a policeman', *u* expresses that John is a policeman at *t*, where *t* is the content of the third word-utterance that is part of *u*.

Since R_1 and R_2 assign the same content to any given utterance of (16), it will be rather difficult to maintain that only one of the rules is associated with (16): that R_1 (R_2) is bogus while R_2 (R_1) is the only 'real' content-rule of (16) will be difficult to motivate. In all likelihood, the proponent of the Token-Reflexive Theory will have to admit that R_1 and R_2 are equally correct content-rules for (16), and more generally that at least some complex expressions are associated with several content-rules at once.

Seeing that some complex expressions are associated with several content-rules, the proponent of the Token-Reflexive Theory could try to define synonymy thus:

SYN_{TRT}. For all expressions e_1 and e_2, e_1 and e_2 are synonymous iff e_1 and e_2 are associated with the same set of content-rules.

To see why SYN_{TRT} fails, consider (17). R_3 and R_4 are plausibly two of the content-rules that are associated with this sentence.

(17) Now John is a policeman.

R_3. For all utterances *u* of 'Now John is a policeman', *u* expresses that John is a policeman at *t*, where *t* is the content of the utterance of 'now' that is part of *u*.

R_4. For all utterances *u* of 'Now John is a policeman', *u* expresses that John is a policeman at *t*, where *t* is the content of the first word-utterance that is part of *u*.

Now, (16) and (17) are intuitively synonymous[51]. However, SYN_{TRT} does not yield this result as the sentences are apparently not associated with the same rules. For one thing the rules mention different sentences: R_1 and R_2 mention (16) while R_3 and R_4 mention (17). And even if we ignore this fact, R_2 cannot plausibly be associated with (17), while R_4 cannot plausibly be associated with (16). Thus (16) and (17) are not associated with the same set of content-rules, whence SYN_{TRT} implies, falsely, that the two sentences are not synonymous.

5.2.3.b 'Giving the meaning'

That (16) and (17) are not associated with the same set of rules gives rise to a second, somewhat more abstract problem. Originally, the Token-Reflexive Theory claimed that *the* content-rule of expression *e* 'gives the meaning' of *e*. That's what motivated defining synonymy in terms of content-rules. Yet since some expressions are associated with several rules, this claim must be revised. There are two options: one can either say that *each* of *e*'s rules *by itself* 'gives the meaning' of *e*; or one can say that *all* of *e*'s

[51] Note that the sentences have the same character according to the Kaplanian account, so that they are synonymous by the Kaplanian's definition. Thus if the proponent of the Token-Reflexive Theory were to deny that (16) and (17) are synonymous, the Kaplanian could argue that her definition of synonymy is better simply because it fits our intuitive understanding of synonymy better.

rules *together* 'give the meaning' of *e*, so that each of *e*'s rules is only a partial description, as it were, of *e*'s meaning. Let's consider these options in turn.

Following the first option, each of R_1 and R_2 'gives the meaning' of (16). Now, since (16) and (17) are synonymous and hence have the same meaning, any rule that 'gives the meaning' of (16) thereby also 'gives the meaning' of (17). However, R_2 does not 'give the meaning' of (17). For, even if we ignore the fact that R_2 mentions (16) rather than (17), R_2 is not a suitable content-rule for (17) and cannot be associated with (17). Thus, it is not true that each of (16)'s rules 'gives the meaning' of (16).

The second option fairs no better. Suppose for a moment that R_1 and R_2 are the only rules associated with (16). Then R_1 and R_2 together 'give the meaning' of (16). Yet they don't 'give the meaning' of (17) as R_2 is not even a partial description of (17)'s meaning. Hence the second option fails also.

In sum, the fact that some complex expressions are associated with more than one content-rule calls into question the Token-Reflexive Theory's tenet that content-rules 'give the meaning' of the expressions with which they are associated. In turn, this suggests – pace García-Carpintero – that the Token-Reflexive Theory *doesn't* "incorporate the point that […] linguistic conventions […] are associated with repeatable entities, types and not tokens"[52]. If this is correct, it is a flaw of the theory.

5.2.3.c Indexical co-reference and translations

As a last option, the proponent of the Token-Reflexive Theory could try to define synonymy in terms of indexical co-reference. First, let us expand CO-REF thus:

CO-REF*. For all sentence-utterances u_1 and u_2, u_1 and u_2 are *jointly* indexically co-referential iff: for all indexicals *i*, u_1 and u_2 were produced in such a way that

[52] García-Carpintero 1998, p. 534.

any utterance of i that is part of u_1 (if any) is co-referential with all utterances of i that are part of u_2 (if any).

To illustrate, consider 'John is in London now' and 'Jane is in London now'. To produce jointly indexically co-referential utterances of these sentences is to utter the sentences in such a way that the two relevant utterances of 'now' co-refer. Note that u_1 and u_2 are jointly indexically co-referential only if both u_1 and u_2 are indexically co-referential.

Now, using CO-REF*, the proponent of the Token-Reflexive Theory can advance the following definition of synonymy.

SYN$_{TRT}$#. For all expressions e_1 and e_2, e_1 and e_2 are synonymous iff: any two jointly indexically co-referential utterances of e_1 and e_2 express the same content.

This definition correctly implies that (16) and (17) are synonymous. Let u_1 and u_2 be utterances of (16) and (17) respectively, and suppose that u_1 and u_2 are jointly indexically co-referential. Suppose also that t is the content of the utterance of 'now' that is part of u_1. Then u_1 expresses that John is a policeman at t. Since u_1 and u_2 are jointly indexically co-referential, t is also the content of the utterance of 'now' that is part of u_2, so that u_2 also expresses that John is a policeman at t. Hence the two utterances have the same content. Since u_1 and u_2 were chosen arbitrarily, it follows more generally that any two jointly indexically co-referential utterances of (16) and (17) have the same content. Hence SYN$_{TRT}$* implies that (16) and (17) are synonymous, as desired.

Although SYN$_{TRT}$# yields that (16) and (17) are synonymous, it is not an adequate definition of synonymy. To see this, note firstly that CO-REF* implies the following: if sentences σ_1 and σ_2 have no indexical in common, any utterance of σ_1 is jointly

indexically co-referential with any utterance of σ_2[53]. Secondly, the *translation* of a sentence into another language is synonymous with the original sentence. E.g. (18), which is the translation of (16) into German, is synonymous with (16).

(18) John ist jetzt Polizist.

Now, since (16) and (18) do not have any indexical in common, any utterance of (16) is jointly indexically co-referential with any utterance of (18). Since (16) and (18) are synonymous, SYN_{TRT}# hence implies that any given utterance of (16) expresses the same content as any given utterance of (18). This implication is clearly false, however: for instance, there is an utterance of (16) that expresses that John is a policeman at t_1, and an utterance of (18) that expresses that John is a policeman at t_2. Thus, pairs like (16) and (18) show that SYN_{TRT}# is not an adequate definition of synonymy.

That SYN_{TRT}, SYN_{TRT}*, and SYN_{TRT}# are inadequate as definitions of synonymy does not show that the proponent of the Token-Reflexive Theory cannot provide an adequate definition at all. Still, that the above definitions fail shows that the proponent of the Token-Reflexive Theory finds it difficult to adequately define synonymy: The theory acknowledges two types of 'meaning', viz. content and content-rules. Yet synonymy cannot be defined as sameness of content, and apparently also not as sameness of content-rule(s). Thus, there is no type of 'meaning' left for synonymous sentences to have in common. Unless the proponent of the Token-Reflexive Theory can overcome this problem, the Kaplanian theory has an advantage over the Token-Reflexive Theory.

[53] Proof: suppose that σ_1 and σ_2 have no indexical in common. E.g. suppose that i_1 is the only indexical in σ_1, while i_2 is the only indexical in σ_2. Then no part of u_1 is an utterance of i_2, and no part of u_2 is an utterance of i_1. Hence, all utterances of i_2 that are part of u_1 co-refer with all utterances of i_2 that are part of u_2; all utterances of i_1 that are part of u_2 co-refer with all utterances of i_1 that are part of u_1. Consequently, u_1 and u_2 are jointly indexically co-referential.

5.3 The Token-Reflexive Theory: conclusion

The objection from validity and the long-sentence-objection are far less convincing than they are sometimes made out to be. The objection from synonymy is considerably more serious, however. Since the Kaplanian theory provides an adequate definition of synonymy, this is a reason for preferring the Kaplanian account. On the other hand, the Token-Reflexive Theory can assign the correct content to Peter's utterance of (1), and can more generally solve the MOP. So in this regard, the Token-Reflexive Theory fares better than the Kaplanian theory. Thus the question is whether the Kaplanian can modify her theory in such a way that it solves the MOP, while still providing an adequate definition of synonymy. The final part of this paper considers such a modification. Like the original Kaplanian theory and unlike the Token-Reflexive Theory, it assigns contents to both sentence-utterances and sentences in contexts.

PART III. THE CONTEXT SHIFT THEORY

In this part of the paper, I shall present and discuss the *Context Shift Theory*, or CST, which is a heavily revised version of the original Kaplanian theory. To summarise the revision, recall that the original Kaplanian theory comprises UTTERANCE-CONTENT, OCCUR, and SENTENCE-CONTENT. By contrast, the CST rejects UTTERANCE-CONTENT and OCCUR and replaces them with principles which do not imply or presuppose UNIQUENESS. This allows the CST to hold that Peter's utterance of (1) occurs in more than one context, which will, in turn, allow the CST to assign a non-uniform proposition as content to Peter's utterance. Yet in connection with the MCA, we saw that this is not enough: a satisfying solution to the MOP further entails that there is a context at which (1) expresses a non-uniform proposition. Since SENTENCE-CONTENT does not entail this,

the CST rejects this principle also and replaces it with an assignment of contents to sentences in contexts that assigns a non-uniform proposition to (1) at some contexts.

Finally, the CST also revises the Kaplanian notion of context. This revision will be explained in the next section. Section seven then puts this new notion of context to work and explains how the CST assigns contents to sentences in contexts, i.e. it presents the CST's semantics. Section eight presents the principle with which the CST replaces UTTERANCE-CONTENT, and section nine presents the CST's analysis of the notion of occurrence. Sections ten and eleven discuss the CST's strength and weaknesses.

Let me make a bibliographical point before I begin. In his *Demonstratives and Their Linguistic Meanings* (Braun 1996), Braun outlines a theory that is very similar to the CST. The major difference between the theories is that Braun's is a theory of demonstratives, while the CST is a theory of (pure) indexicals. Since the CST resembles Braun's theory in many ways, it would be tedious to cite Braun's paper whenever there is an overlap. Note, however, that Braun focusses on providing a semantic theory for sentences in context. He does not provide an analysis of the notion of occurrence.

6. The CST: contexts

As the reader will recall, a Kaplanian context is an ordered quadruple, containing an agent, a time interval, a region of space, and a possible world. So a Kaplanian context looks like this:

$<a_1, t_1, l_1, w_1>$

According to the CST, contexts are still ordered quadruples, but instead of containing a *single* agent, time, location, and world, they contain a *sequence* of agents, a sequence of times, a sequence of locations, and a sequence of worlds. I shall speak respectively of

the agent-, time-, location-, and world-sequence of *c*. The contexts of the CST thus look like this:

$$<<a_1, a_2, a_3, \ldots>, <t_1, t_2, t_3, \ldots>, <l_1, l_2, l_3, \ldots>, <w_1, w_2, w_3, \ldots>>$$

Sequences may contain elements repeatedly: e.g. $<a_1, a_2, a_2, a_1>$ is an admissible agent-sequence, as is $<a_1, a_1, a_1, a_1>$. While one could allow contexts to contain infinitely long sequences, it will be stipulated that all sequences are of finite length: all *sentences* are of finite length, and as will become clear later, this means that contexts with finitely long sequences are all the CST's semantics require. For convenience, it will also be stipulated that all four sequences of a given context have the same length; but nothing depends on this.

Crucially, if *c* is a context, the *successor* of *c*, *c*+, is the ordered quadruple that contains the same sequences as *c* (in the same order), except that the first member of each sequence has been 'cut off'. So the successor of the above CST-context is:

$$<<a_2, a_3, \ldots>, <t_2, t_3, \ldots>, <l_2, l_3, \ldots>, <w_2, w_3, \ldots>>$$

There is one exception: in case *c* is a context the sequences of which contain only one member each, *c*+ = *c*. In sum, each context has a successor, which is itself again a context.

The *second* successor of *c* is the successor of the successor of *c*, i.e. *c*++. The third successor of *c* is *c*+++, and so forth.

If the first element of the agent-sequence of *c* is identical to the first member of the agent-sequence of *c*+, then I shall say that the two contexts *a-resemble* each other. Similarly, *c* can *t*-resemble, *l*-resemble, and *w*-resemble *c*+. If *c* resembles *c*+ with regard to all four sequences, I shall simply say that *c* resembles *c*+.

By *the* speaker of *c* I shall mean the *first* member of *c*'s agent-sequence. (The second member of *c*'s agent-sequence is the speaker of *c*+.) Similarly, the time, location, and world of *c* are the first members of *c*'s time-, location-, and world-sequence.

As was the case with the original Kaplanian theory, worlds can largely be ignored: unless stated otherwise, the world-sequence of any context that I mention contains only the actual world.

The original Kaplanian theory endorsed NOW* and HERE*, and these two principles will stay in place: the content of 'now' at a given context *c* is the time of *c* (i.e. first member of *c*'s time-sequence), unless at the time of *c* and in the world of *c*, the speaker of *c* intends a longer time to be the content of 'now'; similarly for the content of 'here'. The content of 'I' at *c* is the speaker of *c*. Note that an indexical has the same content relative to *c* and *c*+ if and only if *c* resembles *c*+ with regard to the appropriate sequence.

The next section introduces the CST's semantics, and shows how the theory utilises the new notion of context to solve the MOP.

7. The CST: semantics

It would be too complex to formulate the CST's semantics in full generality. So I shall restrict attention to (1) and (2), and shall only present the claims the CST wishes to make about the semantics of these sentences. What claims the CST wishes to make about the semantics of other sentences can be extrapolated from this fairly straightforwardly.

7.1 Sentences in context

To start, indexicals play a central role in the CST's semantics: they are *context-shifters*. In turn, a context-shifter is, roughly, an expression that influences how contents are assigned to sentences in contexts. To explain this further, recall first how the original Kaplanian theory assigns contents to sentences in contexts: to assign a content to σ at *c*, the theory evaluates all the constituent words of σ at the *same* context, viz. at *c*, at the context at which σ is to be evaluated. In general, the CST assigns contents to sentences in the exact same manner. However, if sentence σ contains a context-shifter, then the CST treats σ differently: the CST then evaluates only *some* of σ's constituents at *c*, while evaluating the remaining constituents at the successor of *c*. Let me elaborate.

Let σ be a sentence and suppose one wants to evaluate σ at a given context *c*. Then one must evaluate the constituent words of σ in a certain *order*, according to the CST: one starts with the first constituent word of σ and evaluates it relative to *c*. At what context the second constituent word of σ must be evaluated depends on whether the first constituent word is a context-shifter: if it is *not* a context-shifter, then the second constituent word is evaluated at *c*; yet if the first constituent word *is* a context-shifter, the second constituent word is evaluated at *c*+. The third constituent word is evaluated either relative to the context at which the second constituent word was evaluated or relative to the successor of this context, depending on whether the second constituent word is a context-shifter. And so on[54]. The results of these evaluations, together with σ's compositional structure, then yield the content that σ has relative to *c*.

To illustrate the described evaluation procedure, consider again (2) = 'Smith now courts Mary'. To evaluate (2) at a given context *c*, one begins by evaluating the first

[54] Hence the words are evaluated in the order in which they occur in the sentence's surface form. Towards the end, I shall consider a version of the CST that proceeds in a different order. This has certain advantages, but for now, I'll focus on the former approach as it is easier to describe.

constituent word of (2) at *c*, i.e. one begins by evaluating 'Smith' at *c*[55]. Next, whether 'now' is evaluated at *c* or at *c*+ depends on whether 'Smith' is a context-shifter. To keep matter simple, I shall assume for now that only indexicals are the context-shifter. Hence 'Smith' is not a context-shifter, and 'now' must be evaluated at *c*. By contrast, 'courts' is evaluated at *c*+, because 'now' is a context-shifter. Finally, 'Mary' is evaluated at *c*+ also since 'courts' isn't a context-shifter. Thus to complete the evaluation of (2) at *c*, it only remains to ask what contents 'Smith' and 'now' have at *c*, and what contents 'courts' and 'Mary' have at *c*+.

The content of 'now' at *c* is given by NOW*. As for the remaining words, let's assume that 'now' is the only context-sensitive expression in (2). Then letting $[e]_c$ be the content that expression *e* has relative to *c*, we can formulate the following content-clauses for 'Smith', 'courts', and 'Mary'.

For all contexts *c*,

$[\text{Smith}]_c$ = Smith; $[\text{Mary}]_c$ = Mary; $[\text{courts}]_c$ = the *courting*-relation.

Given these content-clauses and given that the compositional structure of (2) is uncontentious, the CST can formulate the following content-clause for (2).

CONTENT$_{(2)}$*. For all contexts *c*, $[\text{Smith now courts Mary}]_c$ is the proposition that the relation $[\text{courts}]_{c+}$ holds between the individual $[\text{Smith}]_c$, the time $[\text{now}]_c$ and the individual $[\text{Mary}]_{c+}$.

Since 'now' is (assumed to be) the only context-sensitive expression in (2), one can simplify this clause:

CONTENT$_{(2)}$. For all contexts *c*, $[\text{Smith now courts Mary}]_c$ is the proposition that Smith courts Mary at time $[\text{now}]_c$.

[55] I shall assume throughout that a sentence's word-order can be determined by looking at the sentence's surface form. While this assumption may not be plausible in general – because of phenomena like movement –, as far as I can see it is plausible for all the sample sentences that I consider.

Since 'now' is the only context-sensitive expression in (2), the context-shift has no effect on the semantic evaluation of (2). By contrast, (1) contains (at least) two context-sensitive expressions, so that context-shift will impact the semantic evaluation of (1).

Turning, then, to the evaluation of (1), note that the third and the final word in (1) are context-shifters. Hence when (1) is evaluated at a given context *c*, the first three words in (1) will be evaluated at *c*, and the remaining words will be evaluated at *c*+. Thus the CST endorses the following clause for (1)[56].

> CONTENT$_{(1)}$. For all contexts *c*, [Frank lives here and Frank never dines here]$_c$ is the proposition that Frank lives in location [here]$_c$ and never dines in location [here]$_{c+}$.

Now, CONTENT$_{(1)}$ goes a long way towards solving the MOP: let *c* be a particular context and suppose that London is the content of 'here' at *c* while the club is the content of 'here' at *c*+. Then CONTENT$_{(1)}$ implies that (1) expresses at *c* that Frank lives in London and never dines at the club. Hence CONTENT$_{(1)}$ implies that there is a context at which (1) expresses the said proposition. More generally, the CST provides an assignment of contents to sentences in contexts that allows (1) to express a non-uniform proposition relative to some contexts, as desired.

To conclude this sub-section, return briefly to the MCA. Although the MCA is an analysis of utterance-content and not a sentence-in-context theory, one can loosely compare the CST's semantics with the MCA: the MCA treated Peter's two utterances of 'here' separately, assigning them different contents. Similarly, the CST's semantics treats the two occurrences of 'here' separately and assigns them different contents. And as the MCA built the content of Peter's utterance of (1) out of the contents of the relevant utterances of the constituents of (1), the CST builds the content of (1) at *c* out

[56] I assume that 'here' is the only context-sensitive word in (1).

of the contents that are assigned to the constituents of (1). In that sense, the CST's semantics is the sentence-in-context equivalent to the MCA.

7.2 Compositionality

At this point, I'd like to return to a principle I used in constructing the MOP, viz. AND. Let me repeat the principle:

> AND. For all contexts c, if sentence s_1 expresses that p at c while sentence s_2 expresses that q at c, then 's_1 and s_2' expresses at c that p and q.

While this principle is part of the original Kaplanian theory, AND *fails* in the CST. To see this, let c be a particular context and suppose that l_1 is the content of 'here' at c while l_2 is the content of 'here' at c+. Then (3) = 'Frank lives here' expresses at c that Frank lives in l_1, while (4) = 'Frank never dines here' expresses at c that never dines in l_1. However, (1) does *not* express at c that Frank lives in l_1 and never dines in l_1. It expresses at c that Frank lives in l_1 and never dines in l_2. Thus AND fails in the CST.

That AND fails does not mean that the CST cannot formulate a compositional rule for conjunction. To formulate the rule, compare 'Smith now courts Mary' to 'I now court Jane'. While the former sentence contains only one indexical, the latter sentence contains two. Hence when the former sentence is evaluated at a given context c, its last word, 'Mary' is evaluated at c+. By contrast, when the latter sentence is evaluated at c, its last word, 'Jane', is evaluated at c++. Let's summarise this by saying that the sentences differ with regard to their *shift-result*: relative to c, the shift-result of 'Smith now courts Mary' is c+, and the shift-result of 'I now court Jane' is c++.

To generalise this, if a given sentence σ contains n indexicals, then the shift-result of σ at any given context c is the nth successor of c. (The 0th successor of c is c.) With this

definition in place, the CST can formulate the following compositional rule for conjunction:

> AND_{CST}. For all contexts c, if sentence s_1 expresses that p at c while sentence s_2 expresses that q at the shift-result of s_1 at c, then 's_1 and s_2' expresses at c that p and q.

To illustrate, let c be a particular context and suppose that l_1 is the content of 'here' at c while l_2 is the content of 'here' at c+. At c, (3) expresses that Frank lives in l_1. Furthermore, the shift-result of (3) at c is c+. In turn, (4) expresses at c+ that Frank never dines in l_2. Hence AND_{CST} correctly implies that the conjunction of (3) and (4) – i.e. (1) – expresses at c that Frank lives in l_1 and never dines in l_2.

Thus although AND fails within the CST, one can formulate a compositional rule for conjunction. More generally, AND_{CST} illustrates how compositional rules can be formulated within the CST: they will closely resemble the rules of the original Kaplanian theory, but make heavy use of the notion of shift-result.

This completes the presentation of the CST's assignment of contents to sentences in context. The next step is to explain how the CST analyses the notion of utterance-content.

8. The CST: utterance-content

Suppose that Jane utters (2) = 'Smith now courts Mary'. Note that 'Smith' is the first constituent of (2), and that Jane utters 'Smith' when uttering (2). Further, let c_S be the context in which her utterance of 'Smith' occurs. (The CST's analysis of occurrence will be presented in the next sub-section.) Then according to the CST, the content of Jane's sentence-utterance is the content that (2) has at c_S. More generally, the CST advances the following analysis of utterance-content for sentence-utterances:

UTTERANCE-CONTENT$_{CST}$. Let u be an utterance of sentence σ, and let u* be the first word-utterance that is part of u. Then the content of u is the content that σ has relative to the context in which u* occurs.

Note that UTTERANCE-CONTENT$_{CST}$ does not imply or presuppose UNIQUENESS. More to the point, as was the case with the MCA, it is *irrelevant* for the CST's purposes whether UNIQUENESS is true: when assigning a content to a sentence-utterance, UTTERANCE-CONTENT$_{CST}$ only considers the context in which the sentence-utterance's first sub-utterance occurs, but not the context(s) in which the sentence-utterance itself occurs.

Let's apply UTTERANCE-CONTENT$_{CST}$ to Peter's utterance of (1). The first word-utterance that Peter produces in uttering (1) is his first utterance of 'Frank'. Let u_{Frank} be that utterance. Further, let c_{Frank} be the context in which u_{Frank} occurs. By UTTERANCE-CONTENT$_{CST}$, the content of Peter's utterance of (1) is then the content that (1) has at c_{Frank}. By CONTENT$_{(1)}$, the content of (1) at c_{Frank} is the proposition that Frank lives in location $[\text{here}]_{c\text{Frank}}$ and never dines in location $[\text{here}]_{c\text{Frank}+}$. Thus the MOP is *almost* solved: *if* London and the club are the contents of 'here' at c_{Frank} and $c_{\text{Frank}}+$, then UTTERANCE-CONTENT$_{CST}$ and CONTENT$_{(1)}$ together imply that Peter's utterance of (1) expresses that Frank lives in London and never dines at the club. Hence what is missing now is an analysis of the notion of occurrence that implies (PC).

(PC) Utterance u_{Frank} occurs in a context c, such that $[\text{here}]_c$ is London and $[\text{here}]_{c+}$ is the club.

The next section shows how the CST could provide such an analysis.

9. The CST: Occurrence

I first show that the CST can provide an analysis of occurrence that implies (PC), and then suggest a refinement to the CST.

9.1 Peter's utterance of 'Frank'

The aim is to assign a context of origin to u_{Frank} in such a way that (PC) is true. In outline, the CST's proposal is as follows: When Peter utters (1), he utters 'here' twice. So the context in which u_{Frank} occurs is defined in terms of these two indexical-utterances, and more particularly in terms of the times during which Peter produces the utterances, and in terms of the locations that Peter occupies at those times.

In more detail, since (1) contains exactly two context-shifters, it is firstly stipulated that u_{Frank} occurs in a context with two-membered sequences. Secondly, if Peter produces his two utterances of 'here' during t_1 and t_2 respectively, then $<t_1, t_2>$ is the time-sequence of the context in which u_{Frank} occurs. Thirdly, if l_1 and l_2 are the smallest locations that Peter occupies when producing his utterances of 'here', then $<l_1, l_2>$ is the location-sequence of the context in question. Finally, the agent-sequence is <Peter, Peter>, and <@, @> is the world-sequence. Thus in sum, the proposal doesn't at all consider the time during which Peter produces u_{Frank}, or the location that he then occupies. The proposal only considers Peter's utterances of 'here'.

It is easy to see that the proposal yields (PC), at least if Peter has the appropriate intentions. Note simply that London and the club include the smallest locations that Peter occupies when producing his utterances of 'here'. Given the current analysis of occurrence, London and the club hence include the locations of the context in which u_{Frank} occurs and the successor of this context. By HERE*, this means that London and the club are the contents of 'here' at the said contexts, provided that Peter intends his

utterances of ‘here’ to refer to London and the club. Since it is plausible that Peter has these intentions, the current analysis of occurrence thus implies (PC), as desired, and allows the CST to complete its solution to the MOP.

In the remainder of this section, I shall explain how the current proposal can be extended and generalised, and I consider a revision to the proposal.

9.2 Peter’s remaining word-utterances

The current proposal only defines in what context u_{Frank} occurs. To solve CLUBLAND, this is enough. Yet it is interesting to note that the CST can also define in what contexts Peter’s remaining word-utterances occur, by using the notion of a context-shifter. The general idea is this: when a context-shifter is uttered in context c, a shift in context is triggered and the conversation continues in c+. So, since ‘Frank’ is not a context-shifter, u_{Frank} does not trigger a shift in context. Peter’s utterance of ‘lives’ consequently occurs in the same context in which u_{Frank} occurs. Peter’s first utterance of ‘here’ occurs in that context also as ‘lives’ is not a context-shifter either. Yet ‘here’ is a context-shifter, so Peter’s first utterance of ‘here’ triggers a shift in context. His utterance of ‘and’ hence occurs in the successor of the context in which his first utterance of ‘here’ occurs. His remaining word-utterances occur in that latter context, too, as ‘and’, ‘never’, and ‘dines’ aren’t context-shifters. Thus using the notion of a context-shifter, the CST can define in what contexts Peter’s individual word-utterances occur.

9.3 Generalisation

The proposal at hand can be generalised further: If S utters a sentence σ that contains n context-shifters, with $n \geq 1$, then the first word-utterance that S produces in uttering σ occurs in a context that contains only n-membered sequences. The agent-sequence of

that context contains S exactly *n* times. If S uttered the *n*th indexical in σ at *t*, then *t* is the *n*th member of the time-sequence. The location-sequence is constructed in the same way. The world-sequence contains only the world in which the sentence-utterance occurs.

The notion of a context-shifter is used to define in what contexts S's remaining word-utterances occur: let *u* and *u** be the *n*th and *n*+1th word-utterance that S produces in producing the sentence-utterance at hand. Then if *u* is an utterance of a context-shifter, *u** occurs in the successor of the context in which *u* occurs; otherwise *u* and *u** occur in the same context.

Note an important feature of the current analysis of occurrence: it is formulated in terms of *context-shifters*, rather than in terms of *indexicals*. Therefore, the analysis is neutral on whether indexicals are the only context-shifters. E.g. suppose that Jane utters (3) = 'Frank lives here'. Then the current analysis *only* implies that Jane's utterance of 'Frank' occurs in a context with *n*-membered sequences, where *n* is the number of context-shifters that (3) contains, and that Jane's utterances of 'lives' and 'here' occur in the same context iff 'lives' is not a context-shifter. When *combined* with the assumption that indexicals are the only context-shifters, the analysis implies moreover that Jane's utterance of 'Frank' occurs in a context with one-membered sequences, and that her utterances of 'lives' and 'here' occur in the same context. Yet if it turns out, say, that 'lives' is also a context-shifter, the analysis no longer implies this. Instead, it implies that Jane's utterance of 'Frank' occurs in a context with two-membered sequences, and that her utterances of 'lives' and 'here' occur in different contexts. In that sense, the current analysis is rather flexible. The next sub-section makes use of this fact.

9.4 Ubiquitous shift

So far it's been assumed that indexicals are the only context-shifters. Yet given the CST's current analysis of occurrence, it is preferable to assume that *all* words are context-shifters. I start by explaining why it is preferable to assume this, and then show that the assumption can easily be incorporated into the CST.

The context in which a given word-utterance occurs – together with the uttered word – is supposed to *represent* the utterance. Further, if u was produced during t, the best way to represent u is by means of a context that has t as time. Therefore, the following principle is a plausible constraint on the notion of occurrence.

CONTEXT-TIME. For all word-utterances u, if u was uttered during time t, then the time of the context in which u occurs is t.

Given the CST's analysis of occurrence, CONTEXT-TIME fails if indexicals are the only context-shifters. To see this, suppose Jane utters (3) = 'Frank lives here'. Further, let t be the time during which Jane utters 'here'. Then according to the current analysis of occurrence, t is the time of the context in which Jane's utterance of 'Frank' occurs. However, since Jane utters 'here' after uttering 'Frank', t is in fact distinct from the time during which Jane utters 'Frank'. Thus, the time of the context in which Jane's utterance of 'Frank' occurs is not the time during which Jane uttered 'Frank', contrary to CONTEXT-TIME.

Now suppose that all words are context-shifters, so that (3) contains three context-shifters. According to the current analysis of occurrence, this means that Jane's utterance of 'Frank' occurs in a context with three-membered sequences. More particularly, if t_1, t_2 and t_3 are the times during which Jane utters 'Frank', 'lives', and 'here', $< t_1, t_2, t_3>$ is the time-sequence of the context in which her utterance of 'Frank'

occurs. Consequently, the time during which Jane utters ‘Frank’ is the time of the context in which her utterance of ‘Frank’ occurs, as desired.

Further, if all words are context-shifters, Jane’s utterance of ‘Frank’ triggers a context-shift. Hence, her utterance of ‘lives’ occurs in the successor of the context in which her utterance of ‘Frank’ occurs. Since t_2 is the time of that context, the time during which Jane utters ‘lives’ is the time of the context in which her utterance of ‘lives’ occurs. Since her utterance of ‘lives’ triggers another context-shift, it is further the case that the time during which Jane utters ‘here’ is the time of the context in which her utterance of ‘here’ occurs. Hence all of Jane’s word-utterances satisfy CONTEXT-TIME. More generally, if all words are context-shifters, the current analysis of occurrence implies CONTEXT-TIME as desired. Thus, given the CST’s current analysis of occurrence, it is best to assume that *all* words are context-shifters.

If all words are context-shifters, some *details* of the CST must be changed, but the overall theory remains unchanged. More particularly, $\text{UTTERANCE-CONTENT}_{\text{CST}}$ need not be changed at all, and the assignment of contents to sentences in contexts also remains the same overall. Only the content-clauses for the individual sentences must be adjusted. E.g. to assign a content to (2) at *c*, one evaluates ‘Smith’ at *c*, ‘now’ at *c*+, ‘courts’ at *c*++ and ‘Mary’ at *c*+++. (Previously, ‘Smith’ and ‘now’ were evaluated at *c* and ‘courts’ and ‘Mary’ were evaluated at *c*+.) Similarly, to assign a content to (1) at *c*, one now evaluates the two occurrences of ‘here’ at the second and seventh successor of *c*, rather than at *c* and *c*+. The new content-clause for (1) still entails that there is a context at which (1) expresses a non-uniform proposition. Provided that Peter intends his two utterances of ‘here’ to refer to London and the club, the theory entails as before that his utterance of (1) expresses that Frank lives in London and never dines at the club.

To summarise this section, the CST can provide an analysis of occurrence which implies (PC). Using this analysis, the CST can complete its solution to the MOP.

10. Validity

When discussing the Token-Reflexive Theory, I said that a satisfying theory of indexicals should provide an adequate definition of sentence-validity, a definition that yields e.g. that (14) is valid and that (15) is not valid.

(14) If John is happy now then John is happy now.

(15) An utterance exists now.

In this section, I show that the CST can provide such an analysis. In showing this, I shall suppose that the CST assumes that all words are context-shifters.

To start, let c_1 and c_2 be contexts with ten-membered sequences. Suppose further that c_1's time-sequence contains ten *different* times, while all ten members c_2's time-sequence are *identical*. Then (14) expresses a proposition of the form of P_{15} at c_1 while expressing a proposition of the form of P_{16} at c_2.

P_{15} that John is happy at t_2 if John is happy at t_1.

P_{16} that John is happy at t if John is happy at t.

Note that any proposition of the form of P_{16} is necessarily true, and hence guaranteed to be true in the world of c_2. By contrast, propositions of the form of P_{15} are not necessarily true, and hence potentially false in the world of c_1. In defining sentence-validity, the CST will therefore focus on contexts like c_2, and say that a sentence is valid if and only if it expresses a true proposition at such a context.

To spell this out, let us introduce the notion of a *repetitive* context: a context c is repetitive if and only if: for every sequence s of c, all members of s are identical. To

illustrate, the first of the following contexts is repetitive, while the remaining two contexts are not:

$$<<a_1, a_1>, <t_1, t_1>, <l_1, l_1>, <w_1, w_1>>$$

$$<<a_1, a_2>, <t_1, t_2>, <l_1, l_2>, <w_1, w_2>>$$

$$<<a_1, a_2>, <t_1, t_1>, <l_1, l_1>, <w_1, w_1>>$$

Now note that (14) expresses a proposition of the form of P_{16}, and hence a necessarily true proposition, at all repetitive contexts. Thus the following definition yields, correctly, that (14) is valid.

VAL$_{CST}$. For all sentences σ, σ is valid iff: for all repetitive contexts *c*, the proposition that σ expresses at *c* is true in the world of *c*.

The definition further correctly yields e.g. that 'If John is here and I am here, then I am here' and 'I am here now or I am not here now' are valid. (15), and 'If John is here today, then John is here tomorrow', are correctly judged to be not valid. In sum, the CST can provide an adequate definition of validity.

There is an interesting point to note about VAL$_{CST}$: a sentence is valid according to VAL$_{CST}$ if and only if it is valid according to Kaplan's original definition of validity. To see this, suppose that sentence σ is valid according to VAL$_{CST}$. Then σ is true at all repetitive contexts[57]. Since all one-membered contexts are repetitive context, σ is therefore true at all one-membered contexts. Next, the CST assigns proposition *p* to σ at $<<a>, <t>, <l>, <w>>$, iff the Kaplanian theory assigns *p* to σ at $<a, t, l, w>$. Consequently, if σ is true at all one-membered contexts (according to the CST), then σ is true at all Kaplanian contexts (according to the Kaplanian theory). In turn, if σ is true at all Kaplanian contexts, it is valid according to Kaplan's definition. Thus if a sentence

[57] A sentence is true at *c* iff the proposition it expresses at *c* is true in the world of *c*.

is valid according to VAL_{CST}, it is also is valid according to Kaplan's definition. The converse can be established in the same way.

That Kaplan's theory and the CST agree on whether or not a given sentence is valid reinforces the point that the CST is a revised version of the Kaplanian theory rather than an alternative. Insofar as one thinks that the Kaplanian offers an adequate definition of validity, the result furthermore supports the CST, as it implies that VAL_{CST} is adequate also.

11. Synonymy

In this final section I shall show that the CST struggles to provide an accurate definition of synonymy. In presenting the problem, I will first assume that the CST holds that all words are context-shifters. I shall then consider whether the CST can solve the problem by going back to holding that only indexicals are context-shifters.

11.1 Synonymy and sameness of character

The original Kaplanian theory defines synonymy as sameness of character: "To supply a synonym for a word or phrase is to find another with the same *character*"[58]. Recall that this definition has two advantages: firstly, character being one of the two kinds of 'meaning' that the Kaplanian account recognises, the definition captures the pre-theoretic notion that synonymy is sameness of 'meaning'. Secondly, the definition guarantees that synonymous expressions have exactly the same semantic properties. With this in mind, consider how the CST might define synonymy.

One might think that the CST can give the same definition as the original Kaplanian account, and define that two expressions are synonymous iff they have the same

[58] Kaplan 1989a, p. 521, emphasis in the original.

character. To see why this approach fails, suppose that 'ophthalmologist' and 'eye doctor' are synonymous. If this is correct, then sentences (16) and (17) are also synonymous:

(16) John is an ophthalmologist and I am a dentist.

(17) John is an eye doctor and I am a dentist.

However, the CST assigns different characters to (16) and (17). 'I' is the sixth word in (16), but the seventh word in (17)[59]. Hence when (16) is evaluated at a context *c*, 'I' will be evaluated at the fifth successor of *c*. By contrast, when (17) is evaluated at a context *c*, 'I' will be evaluated at the sixth successor of *c*. Therefore, if Jane and Mary are the sixth and seventh member of the agent-sequence of a particular context *c*, (16) and (17) are not assigned the same content at *c*: (16) is assigned the proposition that John is an ophthalmologist and Mary a dentist, while (17) is assigned the proposition that John is an ophthalmologist and Jane a dentist. Hence according to the CST, there is a context at which (16) and (17) do not express the same proposition, which means the theory assigns different characters to (16) and (17). Thus, if the CST defines synonymy as sameness of character, the theory falsely implies that (16) and (17) are not synonymous.

The CST could again use the notion of a repetitive context and propose the following definition, which correctly implies that (16) and (17) are synonymous:

Syn_{CST}. Two sentences are synonymous iff: for all repetitive contexts *c*, the sentences have the same content at *c*.

While Syn_{CST} yields the correct verdict for (16) and (17), it is problematic as a definition of synonymy: as explained above, synonymous sentences are supposed to

[59] While some would argue that 'eye doctor' is a single word, the more common view on noun-noun-compounds is that they are phrases, consisting of two (or more) words, just like e.g. 'young doctor'. See e.g. Giegerich 2004 for a defence of this latter view. See also the third paper in this dissertation.

have the same 'meaning', and identical semantic properties. However, if (16) and (17) do not have the same character, then there is no kind of 'meaning' that the CST ascribes to both (16) and (17). Further, Syn_{CST} allows sentences with different semantic properties to be synonymous: e.g. there is going to be a context *c* such that (16), but not (17), instantiates the property of expressing a true proposition at *c*. Thus, Syn_{CST} is not an adequate definition of synonymy.

11.2 Synonymy and context-shift

As we saw, if the CST endorses that all words are context-shifters, the theory implies that (16) and (17) have different characters. Yet if paired with the assumption that indexicals are the only context-shifters, the CST implies that (16) and (17) have the same character[60]. Hence the CST could goes back to assuming that indexicals are the only context-shifters, and define synonymy as sameness of character after all.

As explained in section 9.4, if the CST holds that indexicals are the only context-shifters, its analysis of occurrence implies that CONTEXT-TIME is false. Yet this may be the smaller bullet to bite: arguably, providing an adequate definition of synonymy is more important than providing a flawless analysis of occurrence. Let's assume, therefore, that the CST abandons the claim that all words are context-shifters. The question is whether this makes it plausible for the CST to define synonymy as sameness of character.

Consider the following two sentences:

(18) I now court Mary.

(19) Now I court Mary.

[60] If indexicals are the only context-shifters, no context-shifter occurs before 'I' in (16) or (17). Consequently, both (16) and (17) will express at any context *c* that John is an ophthalmologist and $[I]_c$ is a dentist.

Intuitively, (18) and (19) are synonymous. What's more, the original Kaplanian theory implies that (18) and (19) are synonymous, as both sentences have the same character according to the original theory. Thus insofar as the CST is supposed to be a revised version of the original Kaplanian theory, and not an alternative, the CST should agree with the original theory on the synonymy of (18) and (19).

It is easy to see, however, that the CST does not assign the same character to (18) and (19), even if it is assumed that indexicals are the only context-shifters: Let *c* be a particular context and suppose that the agent- and time-sequences of *c* are <John, Jack> and $<t_1, t_2>$. Then (18) expresses at *c* that John courts Mary at t_2, while (19) expresses at *c* that Jack courts Mary at t_1. Hence there is a context at which (18) and (19) do not express the same proposition, which means that (18) and (19) have different characters. Thus, even if the CST assumes that indexicals are the only context-shifters, the theory cannot adequately define synonymy as sameness of character.

11.3 Refined context-shift

The CST could refine the notion of context-shift. To explain this, consider a context *c* and its successor *c*+.

$$c = <<a_1, a_2, a_3, \ldots>, <t_1, t_2, t_3, \ldots>, <l_1, l_2, l_3, \ldots>, <w_1, w_2, w_3, \ldots>>$$

$$c+ = <<a_2, a_3, \ldots>, <t_2, t_3, \ldots>, <l_2, l_3, \ldots>, <w_2, w_3, \ldots>>$$

Comparing *c* to *c*+, *each* of *c*'s four sequences differs from the corresponding sequence in *c*+: to move from *c* to *c*+, one 'cuts off' the first member of *each* of *c*'s sequences. Yet one could form successor contexts by altering only *one* of the sequences of the original context, leaving the other three intact. There are, of course, four ways to do this: the *agent*-successor of *c*, c^a, is formed by removing the first member of *c*'s agent-sequence, leaving *c*'s time-, location-, and world-sequence as they were. The

time-, *location*-, and *world*-successors of c, c^t, c^l, and c^w, are formed according to the same pattern. Hence the four successor contexts of the above context c look thus:

$$c^a = \langle\langle a_2, a_3, \ldots\rangle, \langle t_1, t_2, t_3, \ldots\rangle, \langle l_1, l_2, l_3, \ldots\rangle, \langle w_1, w_2, w_3, \ldots\rangle\rangle$$

$$c^t = \langle\langle a_1, a_2, a_3, \ldots\rangle, \langle t_2, t_3, \ldots\rangle, \langle l_1, l_2, l_3, \ldots\rangle, \langle w_1, w_2, w_3, \ldots\rangle\rangle$$

$$c^l = \langle\langle a_1, a_2, a_3, \ldots\rangle, \langle t_1, t_2, t_3, \ldots\rangle, \langle l_2, l_3, \ldots\rangle, \langle w_1, w_2, w_3, \ldots\rangle\rangle$$

$$c^w = \langle\langle a_1, a_2, a_3, \ldots\rangle, \langle t_1, t_2, t_3, \ldots\rangle, \langle l_1, l_2, l_3, \ldots\rangle, \langle w_2, w_3, \ldots\rangle\rangle$$

With the refined notion of context in place, the CST can also refine the notion of a context-shifter: occurrences of 'I' shift a given context c to c^a; occurrences of 'now', 'today', 'yesterday', and 'tomorrow' shift c to c^t; and occurrences of 'here' shift c to c^l. Thus the content-clauses for (18) and (19) are:

CONTENT$_{(18)}$. For all contexts c, [I now court Mary]$_c$ is the proposition that [I]$_c$ courts Mary at time [now]$_{c^a}$.

CONTENT$_{(19)}$. For all contexts c, [Now I court Mary]$_c$ is the proposition that [I]$_{c^t}$ courts Mary at time [now]$_c$.

These clauses guarantee that both sentences express the same contents at all contexts. For, for any context c, the content of 'I' at c is identical to the content of 'I' at c^t; and the content of 'now' at c is identical to the content of 'now' at c^a. Thus by revising the notion of context-shift, the CST can assign the same character to (18) and (19).

The refined version of the CST is superior to the simple version in at least two respects: Firstly, as just mentioned, it yields that sentences like (18) and (19) have the same character, as desired. Secondly, it simply seems more natural to suppose that occurrences of 'I' ('now' / 'here') only affect the context's agent-sequence (time- / location-sequence), rather than affecting all four sequences: after all, the CST was designed to handle sentence in which a given indexical occurs *multiply*, and

not to make the referent of one indexical dependent on occurrences of *other* indexicals.

Despite these advantages, the refined version of the CST still struggles to define synonymy. To see this, consider (20) and (21).

(20) Today John read yesterday's newspaper.

(21) John read yesterday's newspaper today.

If (18) and (19) are synonymous, (20) and (21) are plausibly also synonymous. (Again, the original Kaplanian theory implies that (20) and (21) are synonymous.) However, even given the refined notion of context-shift, the CST implies that (20) and (21) have different characters, which means that the CST still can't define synonymy as sameness of character. The reason is simply that both 'today' and 'yesterday' shift contexts to their time-successors: let c be a particular context and let $[\text{today}]_c$ = Friday and $[\text{yesterday}]_c$ = Thursday; further, let $[\text{today}]_{c^t}$ = Thursday and $[\text{yesterday}]_{c^t}$ = Wednesday. Then (20) expresses at c that John read Wednesday's newspaper on Friday, while (21) expresses at c that John read Thursday's newspaper on Thursday. Hence at c, (20) and (21) do not express the same content, whence they do not have the same character. Thus, the refined version of the CST still cannot define synonymy as sameness of character.

11.4 Order of evaluation

Let me consider one last solution, which proposes to change the CST's order of evaluation. Recall, when evaluating a given sentence, the CST proceeds left-to-right: e.g. when 'John read yesterday's newspaper today' is evaluated, the CST starts by evaluating the leftmost word in the sentence, and then goes on to evaluating the second word from the left; and so on. Hence since 'yesterday' occurs to the left of

'today' in (21), 'yesterday' is evaluated before 'today' when (21) is evaluated. By contrast, since 'yesterday' occurs to the right of 'today' in (20), 'yesterday' is evaluated after 'today' when (20) is evaluated. Consequently, the CST assigns different contents to the two sentences at some contexts.

Now, the current proposal proceeds differently and insists that 'today' must be evaluated before 'yesterday' in *both* (20) and (21). To spell this out, note that both (20) and (21) can be paraphrased thus:

(22) Today, the following happened: John read yesterday's newspaper.

Further, while (22) is only a paraphrase and not a proper syntactic or semantic analysis, (22) makes clear that 'today' in some sense *takes scope* over 'yesterday'. Similarly, 'now' seems to take scope over 'I' in both (18) and (19). Accordingly, the current proposal is that the CST's semantic evaluation proceeds in the order of scope, rather than left-to-right: Since 'today' takes scope over 'yesterday' in both (20) and (21), the evaluation of *either* sentence must evaluate 'today' before evaluating 'yesterday': whether (20) or (21) is evaluated at c, 'today' is evaluated at c and yesterday is evaluated at $c+$. Thus the constituent words of (20) and (21) are evaluated in the same order, whence both sentences express the same content at any given context and have the same character.

I see at least one problem with the current proposal: It is quite plausible that 'now' and 'today' respectively take scope over 'I' and 'yesterday' in (18) – (21); paraphrases like the one in (22) support this. Equally, it is plausible that neither occurrence of 'here' in (1) = 'Frank lives here and Frank never dines here' takes scope over the other. However, there are many sentences were it is not so clear whether one indexical takes scope over another.

For instance, suppose Lena asks Bill whether he can pick her up from the airport the next day. First Bill says that he can't pick her up because his wife told him that she will need help with the shopping. Later, however, it turns out that Bill's wife no longer needs help with the shopping; so Bill tells Lena he will pick her up. Yet then the situation changes again: Bill's wife now needs help in the garden. Bill says to Lena:

(23) I can't pick you up tomorrow after all because my wife now needs my help with something else.

If 'now' takes scope over 'I' in 'I now court Mary', it's plausible that 'tomorrow' takes scope over 'I' in (23). However, it is not so clear whether it also takes scope over 'now'. In particular, paraphrases won't help decide the matter as both of the following two paraphrases seem equally correct:

(24) Tomorrow, the following will be the case: I can't pick you up because my wife now needs my help with something else.

(25) Now that my wife needs my help with something else, the following is the case: tomorrow, I can't pick you up.

This isn't to say that no viable notion of scope can be defined, but giving such a definition will not be trivial.

There may be a second problem, but it is fairly technical: if indexicals have a scope, their scopes can potentially conflict with the scopes of other expressions, and in particular with those of quantifiers. To illustrate this, consider the following sentence.

(26) Everybody John met yesterday will buy today's newspaper sometime after tomorrow.

Let's start with the quantifiers. Apparently, 'everybody' takes scope over 'sometime': the sentence says that each of the people John met will buy a certain newspaper after a certain point in time, without saying that all the newspaper purchases

will be simultaneous. More formally, for every person *p* John met yesterday, there is a time *t* after tomorrow, such that *p* buys a certain newspaper at *t*.

Turning to the indexicals, it is plausible that '(sometime after) tomorrow' takes scope over 'today' – just like 'today' takes scope over 'yesterday' in 'John read yesterday's newspaper today'. However, note that the following paraphrase can*not* be used to assign scopes to the indexicals:

(27) Sometime after tomorrow, the following will be the case: Everybody John met yesterday will buy today's newspaper.

The reason is simply that (27) isn't an accurate paraphrase of (26), because of the difference in quantifier-scope: (27) says that there is a certain time at which all the people John met will buy a certain newspaper – something that (26) doesn't express.

As I said, the problem is fairly technical and (26) is certainly an odd sentence[61]. Still, there is a point to be made here: if one thinks that indexicals have a scope, one must ensure that their scopes match those of other expressions. It remains to be seen whether this is easy or hard to ensure.

In sum, making the CST's semantic evaluation sensitive to scope seems promising in general. However, for the proposal to be fully successful, much more needs to be said about how scope is going to be defined.

[61] Propositional-attitude verbs may be another source of potential conflict. Consider the following sentence:

John now believes that yesterday's newspaper was stolen.

Under its *de re* reading, but not under its *de dicto* reading, the sentence can be paraphrased thus:

Of yesterday's newspaper the following is true: John now believes it to have been stolen.

However, one presumably doesn't want to say that 'yesterday' takes scope over 'now' in the above sentence, even when the sentence is given its de re reading. Hence the proponent of the scope proposal must explain why in this case, paraphrase isn't a guide to indexical-scope.

12. Conclusion

The MOP is a serious challenge to the Kaplanian theory of indexicals: the original account assigns the wrong contents to a large number of sentence-utterances. Further, the puzzling sentence-utterances are utterances of ordinary sentences, utterances that we encounter in everyday conversations. Thus the MOP is not a marginal phenomenon that can be ignored, but a serious problem that must be solved.

The CST offers a promising solution to the MOP. Since it assigns contents to sentences in contexts as well as to sentence-utterances, the CST is much more attractive to the Kaplanian as a solution than the Token-Reflexive Theory. Further, the CST assigns a particular content to an utterance of a given sentence, only if the CST also implies that the sentence has the content in question at some contexts. This is the CST's major advantage over the MCA. In section ten, I further showed that the CST can provide an adequate definition of sentence-validity, so that a logic of indexicals could be developed within the CST.

The last section showed however that the CST has a major flaw: it struggles to define the notion of synonymy. Within a broadly Kaplanian framework, the best definition of this notion is in terms of sameness of character. However, if the CST holds that all words are context-shifter, this definition is not available to the CST as intuitively synonymous sentences like (16) and (17) then have different characters according to the CST. And even if indexicals are the only context-shifters, one can still find pairs of intuitively synonymous sentences to which the CST assigns different characters. Thus whether or not all words are context-shifter, the CST cannot provide a satisfying definition of synonymy.

In conclusion, while the CST is flawed as a theory of indexicals, its central notion, the notion of a context-shifter, may well turn out to be crucial to solving the MOP.

References

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

Cappelen, H. (1999). 'Intentions in Words', *Nous* **33**, 92 – 102.

Cappelen, H. and Lepore, E. (2005). 'Insensitive Semantics. A Defense of Semantic Minimalism and Speech Act Pluralism', Blackwell.

Davies, M. & Humberstone, L. (1980). 'Two Notions of Necessity', *Philosophical Studies* **38**, 1 – 30.

DeRose, K. (2009). 'The Case for Contextualism. Knowledge, Skepticism, and Context, Volume 1', Oxford University Press.

Dorr, C. (forthcoming). 'Transparency and the Context-Sensitivity of Attitude Reports', in M. García-Carpintero and G. Martí (eds.), 'Empty Representations: Reference and Non-existence', Oxford University Press 2014.

Egan, A. (2009). 'Billboards, bombs and shotgun weddings', *Synthese* **166**, 251 – 279.

Field, H. (2008). 'Saving Truth from Paradox', Oxford University Press.

García-Carpintero, M. (1998). 'Indexicals as Token-Reflexive', *Mind* **107**, 529 – 563.

Giegerich, Heinz (2004). 'Compound or phrase? English noun-plus-noun constructions and the stress criterion', *English Language and Linguistics* **8**, 1 – 24.

Goodman, N. & Quine, W.V. (1947), 'Steps Toward a Constructive Nominalism', *Journal of Symbolic Logic* **12**, 105 – 122.

Kaplan, D. (1989a). 'Demonstratives', in J. Almong, J. Perry, and H. Wettstein (eds.), 'Themes from Kaplan', Oxford University Press 1989, 481 – 563.

Kaplan, D. (1989b). 'Afterthoughts', in J. Almong, J. Perry, and H. Wettstein (eds.), 'Themes from Kaplan', Oxford University Press 1989, 565 – 614.

Kaplan, D. (1990). 'Words', *Proceedings of the Aristotelian Society* **64** (suppl.), 93 – 119.

King, J (2006). 'Singular Terms, Reference and Methodology in Semantics', *Philosophical Issues* **16**, 141 – 161.

Nunberg, G. (1993). 'Indexicality and Deixis', *Linguistics and Philosophy* **16**, 1 – 43.

Perry, J. (1997). 'Indexicals and demonstratives', in B. Hale & C. Wright (eds.), 'A Companion to the Philosophy of Language', Blackwell 1997, 586 – 612.

Perry, J. (2003). 'Predelli's threatening note: contexts, utterances, and tokens in the philosophy of language', *Journal of Pragmatics* **35**, 373 – 387.

Predelli, S. (1998). 'I am not here now', *Analysis* **58**, 107 – 115.

Predelli, S. (2004). 'Think Before You Speak: Utterances and the Logic of Indexicals', *Argumentation* **18**, 445 – 463.

Predelli, S. (2006). 'The Problem with Token-Reflexivity', *Synthese* **148**, 5 – 29.

Richard, M. (1993). 'Attitudes in Context', *Linguistics and Philosophy* **16**, 123 – 148.

Sidelle, A. (1991). 'The Answering Machine Paradox', *Canadian Journal of Philosophy* **21**, 525 – 539.

Stalnaker, R. (1978). ‘Assertion’, *Syntax and Semantics* **9**, pp. 315 – 332, reprinted in his ‘Context and Content’, Oxford University Press 1999, 79 – 95.

Subtle Indexicals: Kaplanian Semantics and the Answering Machine Paradox

Kaplan's theory of indexicals predicts that utterances of 'I', 'now', 'here', and 'today' always refer, respectively, to the speaker, the time of the speaking, the location of the speaking, and the day of the speaking. However, there is some well-known linguistic evidence suggesting that this prediction is false. For instance, suppose that a history professor holds a lecture in 2014 and says: 'It is June 22, 1941. Now Hitler begins his invasion of the URSS'. Then the professor's utterance of 'now' intuitively refers to June 22, 1941, rather than to the time of the speaking, contrary to Kaplan's prediction. Similar puzzle cases can be constructed for 'I', 'here', and 'today'. Thus, there is linguistic evidence suggesting that Kaplan's theory makes false predictions about some indexical-utterances. The resulting problem is commonly known as the Answering Machine Paradox.

In this paper, I shall attempt three things. Firstly, I shall summarise the relevant parts of Kaplan's theory and present the linguistic data for which the theory cannot account. (Sections one and two.) Secondly, I shall discuss five potential solutions to the problem at hand. Section three surveys four solutions, explaining their strengths and weaknesses. Sections four to eight discuss a fifth solution in more detail, and argue that it fails. Finally, in section nine, I shall present some new data which none of the discussed solutions can account for.

1. The Kaplanian theory

According to Kaplan, indexicals like 'I', 'now', 'here', and 'today' do not have contents *simpliciter*, but only relative to entities that he calls *contexts*. For present

purposes, we shall assume that contexts are ordered triples, containing an agent, a time interval, and a region of space, respectively called the speaker, time, and location of the context. Further, each indexical is associated with a function from contexts to contents, called its *character*. The content of indexical *i* relative to context *c* is the value that *i*'s character yields for *c*. More particularly, the character of 'I' maps each context *c* to the speaker of *c*. The characters of 'now' and 'here' map each context *c* to the time and location of *c*.

Moving on from indexicals, Kaplan holds that *every* word can be evaluated for content at any given context. If a word is context-*in*sensitive, it has the same content at all contexts. If there are context-sensitive words other than indexicals, then these words vary in content from context to context.

Sentences, too, are assigned contents relative to contexts, in accordance with the following principle[1]:

> SENTENCE-CONTENT. For all sentences σ and contexts *c*, the content of σ at *c* is determined by σ's compositional structure and the contents that σ's constituents have at *c*.

Thus to evaluate what content σ has at *c*, one must first evaluate what contents the constituents of σ have at *c*. These contents, together with σ's compositional structure, then yield the content of σ at *c*. To illustrate, consider the sentence 'I like Oxford'. To evaluate this sentence at a given context *c*, assume that 'like' and 'Oxford' denote the *liking*-relation and Oxford at all contexts, and that John is the speaker of *c*. Then the content of 'I' at *c* is John, and the sentence expresses at *c* that John likes Oxford.

[1] Kaplan holds that demonstrative-containing sentences don't have contents relative to contexts: they only have contents relative to *pairs* of contexts and demonstrations – where a demonstration is an extra-linguistic act, like a pointing gesture. (See Kaplan 1989a, b for details.) If this is correct, SENTENCE-CONTENT does not apply to sentences that have demonstratives as constituents. Since I won't consider demonstratives in any detail, this complication is irrelevant for present purposes, however.

So far the theory has only assigned contents to expression *types* (relative to contexts). Yet the theory can also assign contents to expression *tokens*, i.e. to utterances[2]. For present purposes, it is enough to explain how the theory assigns contents to utterances of sentences.

To start, utterances are linked to contexts, by saying that sentence-utterances *occur* in contexts. In particular, if *u* is a sentence-utterance, then *u* occurs in <S, *t*, *l*> iff: S uttered *u* during time *t*[3], and *l* is the smallest location that S occupies during *t*[4]. Note that this implies that each sentence-utterance occurs in exactly one context. The context in which an utterance occurs is commonly called its *context of origin*.

With the notion of occurrence analysed, the theory can assign contents to sentence-utterances thus: if *u* is an utterance of sentence σ, and if *u* occurs in context *c*, then the content of *u* is the content that σ has relative to *c*. Thus if John utters 'I like Oxford', his utterance expresses that John likes Oxford. For, his utterance occurs in a context that has John as speaker, and at such a context, the content of the sentence 'I like Oxford' is the proposition that John likes Oxford.

Officially, the theory only assigns contents to sentence-utterances, and not to word-utterances. However, for simplicity, I shall occasionally say that this-and-this word-utterance has that-and-that content. E.g. if Jane says 'John is home now', and if her sentence-utterance expresses that John is home at *t*, I shall say that Jane's utterance of 'now' refers to *t*.

---

[2] Kaplan's original theory doesn't assign contents to utterances, but it is now common to extend his theory so that it can assign contents to utterances also. E.g. Predelli (1998), Sidelle (1991), and Perry (2003) propose extensions of this kind.

[3] By the time *during* which S utters *u* I mean the time interval that starts when S starts uttering *u* and ends when S has finished uttering *u*.

[4] As the first paper in this dissertation explains, the above analysis of occurrence does not sit particularly well with the assumption that the characters of 'now' and 'here' map each context *c* to the time and location of *c*. Yet for present purposes, this does not matter: the Kaplanian could refine her notion of occurrence, or her assumption about the characters of 'now' and 'here'. (See *Indexicality and Co-Reference*, section 2.3, for a refinement of the latter kind.) However, no (plausible) refinement of this kind will allow her to avoid the problems that I discuss in this paper. For simplicity, I shall therefore assume that the Kaplanian analyses occurrence as described, and that she assigns the said characters to 'now' and 'here'.

2. The Puzzle

In this section, I shall provide utterances of indexical-containing sentences that appear to contradict the Kaplanian account. In the literature, the problem that these utterances give rise to is commonly called the *Answering Machine Paradox*, due to the fact that one of the first known instances of the problem involved an utterance recorded on an answering machine[5]. In this paper, I won't consider the case involving the answering machine, but shall nevertheless refer to the problem that the Kaplanian faces as the Answering Machine Paradox (AMP).

2.1 The linguistic data

Only the first of the following puzzle cases will be spelled out in full detail. The remaining cases follow exactly the same pattern, however.

HISTORY LECTURE. The history professor holds a lecture in 2014 and utters: 'In this lecture, I'd like to talk about World War Two.' She continues:

(1) It is June 22, 1941. Now Hitler begins his invasion of the URSS.

Let *c* be the context in which the professor's last sentence-utterance occurs, and let *t* be the time of *c*. Then *t* is also the content of 'now' at *c*, and the second sentence in (1) expresses at *c* that Hitler began his invasion at *t*. Hence according to the Kaplanian account, the professor's last sentence-utterance expresses that Hitler began his invasion at *t*. However, the utterance intuitively *doesn't* express the said content: while *t* is a time in 2014, the professor's sentence-utterance intuitively expresses that Hitler began his invasion *on June 22, 1941*. Thus, the Kaplanian assigns the wrong content to the professor's sentence-utterance.

[5] See Sidelle 1991.

CRIMINOLOGY LECTURE. The criminology professor holds a lecture in 2014 and utters: 'In this lecture, I'd like to talk about famous assassinations.' She continues:

(2) It is March 15, 44 BC. Today Brutus assassinates Caesar.

The professor's last sentence-utterance intuitively expresses that Brutus assassinated Caesar on March 15, 44 BC. Yet as the professor is speaking in 2014, it follows that March 15, 44 BC isn't the time of the context in which the professor's sentence-utterance occurs. Thus the Kaplanian account implies that the utterance does not express the said proposition.

PRISON BREAK. A tour guide at a retired prison facility describes a spectacular prison break that happened 50 years ago. She and the tourists are standing on a watch tower while she is reporting: 'After digging for weeks, Pete the prisoner can get out of his cell and into the sewer system.' Still on the watch tower, she says:

(3) Pete next accessed the infirmary. Here he was spotted by a guard and arrested.

The tour guide's last sentence-utterance intuitively expresses that Pete was spotted and arrested in the infirmary, but the Kaplanian account does not assign this proposition as content to her utterance.

STICKY NOTE[6]. John's office is next to Frank's. John knows that Frank is on holiday for a week, but worries that some of Frank's colleagues do not know this. So John puts a note on Frank's office door saying:

(4) I am not in this week.

Although John produced the utterance of (4) at hand, it intuitively expresses that Frank is not in his office during the week in question. As before, the Kaplanian account does not yield this result.

[6] This case is due to Corazza *et al*. (2002).

Having presented the puzzle cases, I'd like to compare the first three cases to the last case. Note that HISTORY LECTURE, CRIMINOLOGY LECTURE, and PRISON BREAK involve spoken sentence-utterances. Yet it is easy to construct parallel cases that involve written utterances instead. For instance, imagine that the history professor writes a paper on World War Two; then she could felicitously include (1) in her paper, and the second sentence in (1) would still express that Hitler began his invasion on June 22, 1941. Similarly, the criminology professor could include (2) in a paper on famous assassinations. As for PRISON BREAK, imagine that there is a small museum near the prison (but outside the infirmary). If (3) appears on an information board that describes Pete's famous escape, the second sentence in (3) still expresses that Pete was spotted and arrested in the infirmary. In sum, there are now two instances of the AMP for each of 'now', 'today', and 'here': one features a spoken sentence-utterance, and one features a written sentence-utterance.

Turning to STICKY NOTE and 'I', matters are different: STICKY NOTE features a written utterance, and it is not clear that one can construct a parallel case that involves a spoken utterance[7]. More generally, it is not easy to construct a convincing instance of the AMP for 'I' that features a spoken utterance. To see this, suppose the history professor begins his lecture with: 'This lecture discusses Claus von Stauffenberg'. He continues:

(5a) I was born in Swabia on November15, 1907.

In my opinion, it is almost impossible to read the professor's utterance of (5a) as expressing a statement about Stauffenberg: prima facie, the statement is plainly false. The situation is somewhat improved if the professor begins his lecture as follows: he

[7] E.g. suppose that John hears one of Frank's colleagues knocking on Frank's office door. Peeking out his own office, John says 'I am not in this week'. Although Frank's colleague may understand what John is trying to communicate, it is not very plausible to suppose that John's sentence-utterance literally expresses that Frank is not in his office during the week in question.

announces as before that he will speak on Stauffenberg and then utters: 'When I think about Stauffenberg, I often try to put myself in his shoes'. He continues:

(5b) I am an officer of the Wehrmacht, and plan to assassinate Hitler.

Since the professor announced that he will put himself in Stauffenberg's shoes, it is fairly natural to read his utterance of (5b) as expressing a statement about Stauffenberg[8]. Still, one might think that the professor is speaking non-literally, or engaging in some kind of pretence. Yet for present purposes, we needn't decide the matter: HISTORY LECTURE, CRIMINOLOGY LECTURE, PRISON BREAK, and STICKY NOTE are clearly instances of the AMP – as are the cases for 'now', 'today', and 'here' which feature written sentence-utterances. Hence even if it turns out that the professor's utterance of (5b) is not an instance of the AMP, there is enough data to challenge the Kaplanian account.

2.2 Related data

Before moving on, it will be useful to compare the data in the previous sub-section to the following case, which is often discussed in the literature[9].

MAP. While in London, Peter points at Munich on a map and utters (6).

(6) The Oktoberfest takes place here.

Intuitively, Peter's utterance of (6) expresses that the Oktoberfest takes place in Munich. However, as Munich isn't the location of the context in which Peter's utterance occurs, the Kaplanian account does not assign this content to Peter's utterance.

Most authors who discuss cases like MAP agree that they are *not* instances of the AMP. Rather, it is thought that Peter uses 'here' in a special way, and that Kaplan's

[8] Uttering (5a) after saying 'When I think about Stauffenberg, I often try to put myself in his shoes' still seems odd.
[9] See e.g. Sidelle 1991, pp. 528 – 529.

theory of indexicals simply does not apply to these special uses. In particular, Peter uses 'here' *demonstratively*. Let me elaborate.

The theory I outlined above is Kaplan's theory of indexicals. He develops a separate theory for demonstratives – i.e. a theory for expressions like 'this', 'that', and 'there'. (A major difference between the theories will be explained below.) Importantly, Kaplan's theory of indexicals does not apply to demonstratives, and doesn't assign contents to utterances of demonstrative-containing sentences. Now the consensus in the literature is that (6) contains a demonstrative rather than an indexical: the idea is that there are two words spelled 'here' in English; one is an indexical, the other is a demonstrative[10]. Further, (6) contains the demonstrative 'here', rather than the indexical, whence Peter's utterance of (6) is an utterance of a demonstrative-containing sentence. Thus Kaplan's theory of indexicals doesn't apply to (6), and doesn't assign a content to Peter's sentence-utterance. *A fortiori*, it doesn't assign the *wrong* content to his sentence-utterance, whence MAP is clearly not an instance of the AMP.

I shan't discuss whether the proposed reply to MAP is correct. In the literature, it is normally assumed that the proposal is correct, and I shall accept this assumption: I shall accept that MAP is not an instance of the AMP.

The reader may wonder whether one could generalise the current proposal, by saying that *all* the above puzzle cases involve demonstratives rather than indexicals. This idea will be discussed in the next section.

3. Four potential solutions to the AMP

In this section I shall present four potential solutions to the AMP, and point out some of the problems they face.

[10] According to an alternative proposal, there is just one word 'here' in English, viz. the indexical. However, there are two ways to pronounce the demonstrative 'there': normally, it is pronounced 'there', but in certain circumstances – such as Peter's – one can pronounce it 'here' instead.

3.1 The Character Solution

The first solution, which I shall call the *Character Solution*, may seem obviously correct at first, but quickly turns out to be unsatisfying. The general idea is to ascribe more complex characters to the indexicals. I shall focus on 'now' in explaining the proposal, but everything I say applies mutatis mutandis to the remaining indexicals.

To start, recall that the original Kaplanian theory holds that the character of 'now' functions as described in NOW.

NOW. For all contexts *c*, the content of 'now' at *c* is the time of *c*.

According to the Character Solution, NOW must be rejected and replaced with a principle of the following kind.

NOW*. The character of now is such that:

(a) It maps most contexts *c* to the time of *c*.

(b) It maps the context in which the professor's utterance of (1) occurs to June 22, 1941.

As should be clear, NOW* is incomplete. E.g. suppose that Bill utters (7) = 'It is December 3, 1804. Napoleon is now Emperor of the French.' Then the proponent of the Character Solution will add the following clause to NOW*.

(c) It maps the context in which Bill's utterance of (7) occurs to December 3, 1804.

More generally, the proponent of the Character Solution will add a clause like (b) and (c) to NOW* for every instance of the AMP that involves 'now'. She will further formulate statements like NOW* for 'I', 'here', and 'today'. Consequently, the Character Solution assigns the correct content to all utterances of indexical-containing sentences.

Despite the seeming success of the proposal, it does not yield a satisfying solution to the AMP. For, while NOW* may be an adequate *description* of the character of 'now', it does not yield a *theory* for 'now'. In particular, it does not make any *predictions* about

the contents of new sentence-utterances. E.g. suppose that I utter 'Now Smith courts Mary', and that my utterance of 'now' does not refer to the time of the utterance. *Once it's been discovered* what time my utterance refers to instead, a new clause will be added to NOW*. Yet until this discovery has been made, the Character Solution does not make any predictions about the content of my utterance. In short, the Character Solution is an uninformative list, rather than a genuine semantic theory of indexicals. Therefore, it does not provide a solution to the AMP.

3.2 The Pragmatic Solution

The second solution I'd like to consider will be called the *Pragmatic Solution*. According to this proposal, our intuitions about the puzzling utterances of (1) to (4) are simply misleading. For instance, when the history professor utters (1), her second sentence-utterance does *not* express that Hitler begins his invasion on June 22, 1941, contrary to our intuitions. Her utterance *pragmatically communicates* this proposition, but semantically expresses whatever proposition the Kaplanian account assigns to the utterance. The Pragmatic Solution makes similar claims about the remaining instances of the AMP.

In the literature, the Pragmatic Solution is normally only mentioned in passing and immediately set aside. Hence the consensus in the literature seems to be that the AMP requires a semantic solution: a theory of indexicals is required that implies – among other things – that the history professor's utterance semantically expresses that Hitler begins his invasion on June 22, 1941. I agree with this, but would like to mention two more specific arguments against the Pragmatic Solution, due to Michaelson and Cohen[11].

[11] See Michaelson and Cohen 2013.

Firstly, as Michaelson and Cohen understand the proposal, the Pragmatic Solution claims that the instances of the AMP involve *conversational* rather than *conventional* implicature. However, conversational implicature can normally be cancelled. For instance, if John tells Mary that he is out of petrol, she might say 'There is a petrol station around the corner'. Her utterance conversationally implies that the petrol station is open, but she can cancel this implicature: 'There is a petrol station around the corner, *but it is closed*'. By contrast, the history professor cannot felicitously utter (1a).

(1a) It is June 22, 1941. Now, *but not on June 22, 1941*, Hitler begins his invasion of the URSS.

Yet if the professor cannot cancel that Hitler began his invasion on June 22, 1941, the said proposition is not merely a conversational implicature of his utterance, but rather its semantic content.

The proponent of the Pragmatic Solution could reply that the professor *can* cancel that Hitler began his invasion on June 22, 1941. Consider the following speech.

(8) Smith is in London now – and by 'now' I don't mean this morning: I mean this month.

Prima facie, (8) is felicitous: the clause that follows the hyphen simply makes the semantic content of 'now' explicit. It therefore stands to reason that the professor can use a similar clause to make explicit the content of her utterance of 'now', thereby cancelling that Hitler began his invasion on June 22, 1941. Consider:

(1b) It is June 22, 1941. Now Hitler begins his invasion of the URSS – and by 'now' I don't mean June 22, 1941: I mean the present time.

While (1b) is marginally better than (1a), the statement is nevertheless rather difficult to understand. If the professor doesn't intend 'Now Hitler begins his invasion of the

URSS' to be a statement about June 22, 1941, why does she say 'It is June 22, 1941'? Hence although (1b) may not be contradictory, it is certainly very misleading. Therefore, it is doubtful that (1b) successfully cancels that Hitler began his invasion on June 22, 1941.

Michaelson and Cohen raise a second objection against the Pragmatic Solution: they object that it is "potentially unstable"[12]. For ease of reference, let (1.1) be the first sentence in (1), and let (1.2) be the second sentence. Now, the professor's utterance of (1.2) intuitively expresses that Hitler began his invasion on June 22, 1941. Yet the professor's utterance is not special in this regard: virtually *all* utterances of (1.2) that follow an utterance of (1.1) intuitively express this content. Hence the Pragmatic Solution will claim that all the utterances of (1.2) in question pragmatically communicate that Hitler began his invasion on June 22, 1941. However, Michaelson and Cohen argue, this is implausible: if a content is "*regularly* and *systematically*"[13] associated with (utterances of) a certain expression, one would expect that that content eventually becomes the expression's semantic content. Thus, they conclude, whatever the semantics of (1.2) may have been in the past, its semantics are now such that the professor's utterance semantically expresses that Hitler began his invasion on June 22, 1941. Similar remarks apply, of course, to sentence (2) to (4).

In response, the proponent of the Pragmatic Solution could appeal to the notion of *generalised* conversational implicature[14]: while some conversational implicatures are context-specific, others appear to be context-independent or generalised. To illustrate, suppose that Tim and Tom are at a party, and compare the following two conversations.

CON.1 Tim: What time is it?

Tom: Some guests have already left.

[12] Michaelson and Cohen 2013, p. 581.
[13] Michaelson and Cohen 2013, p. 581, emphasis in the original.
[14] For a detailed discussion of this notion, see e.g. Levinson 2000.

CON.2 Tim: Where is John?

Tom: Some guests have already left.

Tom's statement in CON.1 plausibly implies that Tom suspects that it is already quite late. By contrast, Tom's statement doesn't imply this in CON.2, but rather implies that Tom suspects that John has already left. Hence both implicatures are context-specific. However, there is also a context-independent implicature in place: Tom's remarks imply in both CON.1 and CON.2 that some guests have not yet left. In fact, most – if not all – utterances of 'Some guests have already left' carry this implicature. Thus, that some guests have not yet left is a generalised conversational implicature of 'Some guests have already left'. That is to say, the said proposition is regularly and systematically associated with utterances of this sentence, although none of the sentence's utterances express this content semantically.

Returning to (1), the proponent of the Pragmatic Solution could try to argue that (1.2) also carries a generalised conversational implicature, viz. that Hitler began his invasion on June 22, 1941. This would account for why utterances of (1.2) are regularly and systematically associated with the said proposition, without making the proposition the sentence's semantic content (at certain contexts). The problem with this proposal is that the said implicature does not appear to be general enough: As indicated, *most* utterances of 'Some guests have already left' imply that some guests have not yet left. By contrast, even if the Pragmatic Solution is correct, only a relatively small number of utterances of (1.2) imply that Hitler began his invasion on June 22, 1941, viz. only those utterances that follow an utterance of (1.1). Thus, that (1.2) carries the said content as generalised conversational implicature is doubtful.

To summarise the discussion of the Pragmatic Solution, the proposal firstly struggles to account for the infelicity of (1a) and (1b), and is secondly potentially unstable.

Thirdly and more generally, the consensus in the literature is that the best solution to the AMP will be in semantic, rather than pragmatic terms. Thus, the Pragmatic Solution is not very attractive.

3.3 The Demonstrative Solution

The next solution I'd like to consider will be called the *Demonstrative Solution*. It attempts to generalise the claims that were made about MAP and Peter's utterance of 'here'.

Recall, the consensus in the literature is that Peter utters a demonstrative when he says 'here', so that Kaplan's theory of indexicals does not apply to his utterance of (6). Now, if there is a demonstrative spelled 'here' in English, then perhaps there are also demonstratives spelled 'now', 'today', and 'I' in English. In turn, if all these demonstratives exist, one could hold that they – rather than the homonymous indexicals – occur in sentences (1) to (4). Hence when the speakers in the above puzzle cases produce their sentence-utterances, they do not utter any indexicals. Thus, the Demonstrative Solution concludes, Kaplan's theory of indexicals does not apply to sentences (1) to (4), and doesn't assign a content to the sentence-utterances in question. *A fortiori*, the utterances don't pose a problem for Kaplan's theory of indexicals, and the AMP is solved.

To clarify, the Demonstrative Solution holds that the demonstratives 'here', 'now', 'I', and 'today' exist *alongside* the indexicals 'here', 'now', 'I', and 'today'. Further, Kaplan's theory of indexicals gives the correct semantics for the latter four expressions, according to the Demonstrative Solution. In other words, the Demonstrative Solution *defends* Kaplan's theory of indexicals against the AMP. Hence the Demonstrative Solution must not be confused with proposals that *reject* Kaplan's theory of indexicals:

according to such a proposal, there is just one set of expressions spelled 'here', 'now', 'I', and 'today' in English, and these expressions are demonstratives[15]. On this alternative proposal, Kaplan's theory of indexicals does not give the correct semantics for 'here', 'now', 'I', and 'today'. In fact, it doesn't give the correct semantics for anything. While the said proposal is interesting in its own right, I won't discuss it in the present paper. For, I am interested in whether one can solve the AMP *within* Kaplan's theory of indexicals, i.e. without abandoning the claim that English contains indexicals spelled 'here', 'now', 'I', and 'today'. The Demonstrative Solution attempts this; yet as I shall now argue, the attempt fails.

To reply, then, to the Demonstrative Solution, let me explain one of the major differences that Kaplan sees between indexicals and demonstratives. Kaplan holds that indexicals are *automatic*: What content a given indexical-utterance has is determined only by the context in which the utterance occurs. In particular, the utterance's content is determined independently of any intentions that the speaker may have: If the speaker has no reference-fixing intentions when producing her utterance, the utterance can still acquire a content. On the other hand, the utterance will acquire whatever content the context supplies, even if the speaker intends the utterance to refer to another individual, time, or location.

By contrast, demonstratives are not automatic, according to Kaplan: if one utters a demonstrative-containing sentence, one's utterance won't acquire a content unless the utterance is accompanied by an extra-linguistic act, a 'demonstration'[16].

Now, prima facie the puzzling utterances of (1) to (4) are not accompanied by demonstrations: e.g. while Peter points at the map when uttering (6), John presumably

[15] E.g. Mount (2008) defends such a proposal.

[16] Pointing gestures are the paradigm example of demonstrations, but most philosophers think that not all demonstrations are pointing gestures. E.g. Kaplan (1989b, pp. 582 – 584) holds in his later work that directing intentions are a kind of demonstration. But whatever demonstrations are, utterances of demonstratives must be accompanied by *them* in order to acquire a referent, according to Kaplan.

does not point at Frank when writing (4). Nevertheless, the utterances of (1) to (4) appear to have contents. Hence if Kaplan is right in saying that demonstratives aren't automatic, it follows that the utterances of (1) to (4) aren't utterances of demonstrative-containing sentences, but rather utterances of indexical-containing sentences – which, in turn, means that Kaplan's theory of indexicals applies to the utterances. Thus, the Demonstrative Proposal fails as a solution to the AMP.

Perhaps someone will reply that the utterances of (1) to (4) *are* accompanied by demonstrations: granted, John does not point at Frank when uttering (4); but he does *intend* his utterance of 'I' to refer to Frank. The speakers of (1) to (3) have similar intentions. These intentions, the reply continues, are demonstration.

To respond to this, let's continue to focus on John's utterance of (4). Further, recall that the Demonstrative Solution wants to maintain the distinction between indexicals and demonstratives: the proposal agrees that English contains an indexical spelled 'I' (as well as a demonstrative spelled 'I'), and also that this indexical is automatic (while the demonstrative is not). However, this will be difficult to maintain if John utters the demonstrative 'I' in uttering (4), rather than the indexical: John's utterance of (4) is just too similar to typical indexical-utterances.

To illustrate this, consider a sentence that contains the indexical 'I' and an utterance thereof. (Again, the Demonstrative Solution agrees that such sentences and utterances exist.) E.g. suppose Jane utters 'I am not in this week'. If John intends his utterance of 'I' to refer to Frank, Jane presumably intends to refer to herself with her utterance of 'I'. If John's intention is a demonstration, it would be implausible to suppose that Jane's intention isn't a demonstration. And if John must supply a demonstration for his sentence-utterance to acquire a content, it will be difficult to maintain that Jane needn't do this.

More generally, seeing how similar John's and Jane's sentence-utterances are, the proponent of the Demonstrative Solution cannot plausibly hold that Jane's sentence-utterance is automatic while John's is not: if she holds that John utters the demonstrative 'I', she must also hold that Jane utters the demonstrative 'I'. Yet if Jane utters the demonstrative 'I', it is not clear that there are *any* utterances of the indexical 'I': Jane's utterance is arguably the most typical example of an indexical-utterance. Hence if it turns out that she uttered the demonstrative 'I', it is unclear what utterances of the indexical 'I' would look like.

Thus in sum, if the proponent of the Demonstrative Proposal holds that John utters the demonstrative 'I' in uttering (4), she will be hard-pressed to find utterances of the indexical 'I'. Since similar considerations apply to 'now', 'here', and 'today', she will more generally have difficulties finding utterances of indexicals. In turn, if all utterances of 'I', 'now', 'here', and 'today' turn out to be utterances of demonstratives, the Demonstrative Proposal must abandon the claim that English contains indexicals spelled 'I', 'now', 'here', and 'today'. Thus, maintaining that (1) to (4) contain demonstratives is not a viable option and the Demonstrative Solution fails.

3.4 The Intentionalist Solution

The last solution I consider in this section is due to Stefano Predelli[17], and is usually called *Intentionalism*. In a nutshell, Intentionalism holds that speaker intentions play a crucial part in fixing the referents of indexical-utterances: e.g. the history professor's utterance is a statement about June 22, 1941 because the professor *intends* her utterance

[17] See Predelli 1998.

to refer to this date. Since Intentionalism is discussed extensively in the literature[18], I shall only summarise the theory and point out its major flaw.

3.4.1 Intentionalism and the AMP

According to Intentionalism, every sentence-utterance is associated with *two* contexts, viz. its context of origin and its *context of intention*. Contexts of origin are defined as before[19]. As for contexts of intention, suppose that John utters 'I am here now', and suppose that he intends his utterance of 'I' to refer to S, and that he intends his utterances of 'here' and 'now' to refer to *l* and *t*. Then <S, *t*, *l*> is his utterance's *context of intention*. Importantly, the time / location of his utterance's context of intention need not include the time / location of the speaking. On the other hand, if John intends his utterance of 'I' to refer to himself, and intends his utterances of 'here' and 'now' to refer to the time and location of the speaking, his utterance's context of intention and his utterance's context of origin will be identical.

Next suppose that John utters 'I am married now'. Since John does not use 'here', his intentions do not determine a location for his utterance's context of intention. Therefore, the location of his utterance's context of origin will also be the location of his utterance's context of intention. An utterance's context of intention is constructed similarly if the speaker's intentions do not single out an individual or a time.

Turning to the semantic evaluation of sentence-utterances, the Intentionalist makes the same claim as the Kaplanian, except that she says 'context of intention' where the Kaplanian said 'context of origin': For all utterances *u*, if *u* is an utterance of sentence σ, and if *c* is *u*'s context of intention, then the content of *u* is the content that σ has

[18] See e.g. Corazza *et al.* 2002, Predelli 2011, and Michaelson and Cohen 2013.
[19] Recall, if S utters *u* during *t* and if *l* is the smallest location that S occupies during *t*, then *u* occurs in <S, *t*, *l*>. The time *during* which S utters *u* is the time interval that starts when S starts uttering *u* and ends when S has finished uttering *u*.

relative to *c*. Further, the characters of 'I', 'now', and 'here' remain unchanged. Hence Intentionalism and the original Kaplanian theory assign the same content to a sentence-utterance *u* iff *u*'s contexts of intention and origin are identical.

It should be obvious that Intentionalsim solves the AMP: when the history professor in HISTORY LECTURE utters 'Now Hitler begins his invasion of the URSS', she intends her utterance of 'now' to refer to June 22, 1941. Thus her utterance's context of intention has June 22, 1941 as time, whence her utterance expresses that Hitler begins his invasion on June 22, 1941. The utterances of (2) to (4) are treated similarly, whence the AMP is solved.

3.4.2 The Humpty-Dumpty problem

Although Intentionalism looks promising at first, it faces a serious problem, sometimes called the Humpty-Dumpty problem[20]: the theory implies that indexical-utterances can refer to virtually any person, time, or location, as long as the speaker has the appropriate intentions. However, this is plainly false. To illustrate, suppose that Frank and Mary are in London. He asks her where she went to school, and she replies: 'I went to school here'[21]. Suppose further that Mary went to school in Boston, and that she intended to refer to Boston with her utterance of 'here'. Then if Intentionalism is true, Mary's sentence-utterance expresses a true proposition. However, Mary's utterance is clearly false: it expresses that she went to school in London. A reading on which it expresses that she went to school in Boston is simply not available. More generally, speakers do not have as much semantic control over their indexical-

[20] E.g. Corazza *et al*. (2002) use this term.
[21] I am imagining a case in which Mary doesn't point at a map or the like. She just looks at Frank and utters 'I went to school here'.

utterances as the Intentionalist position ascribes to them. For that reason, most authors think that Intentionalism fails as a solution to the AMP[22].

Let me briefly mention two potential solutions to the Humpty-Dumpty problem. The first is due to Kent Bach[23] and insists that reference-fixing intentions are of a fairly specific form and subject to Gricean constraints. In particular, for Mary's utterance of 'here' to refer to Boston, it's not enough that she intends: 'I want my utterance of 'here' to refer to Boston'. Rather, she must intend her sentence-utterance to communicate a content about Boston to Frank. Roughly, she must intend that upon hearing her sentence-utterance, Frank comes to believe that she went to school in Boston. However, given the communicative setting, Mary knows that Frank won't form this belief, as he has no way of figuring out that Mary is attempting to talk about Boston. Thus, the proposal concludes, Mary can't rationally form the required intention, whence her utterance refers to London.

In reply, suppose that Mary believes that she *can* communicate a content about Boston with her utterance of 'I went to school here'. (Perhaps she is confused about the communicative setting, or generally holds strange beliefs about what utterances can and can't communicate.) Then it seems plausible that Mary has the relevant kind of intention. However, her utterance of 'here' intuitively still refers to London, and not to Boston. Thus, the proposed adjustment to Intentionalism will find it difficult to account for cases that involve speakers who hold unusual beliefs about communicative possibilities.

The second reply to the Humpty-Dumpty problem is due to Predelli[24]. He insists that contrary to our intuitions, Mary's utterance of 'here' *does* refer to Boston: Her

[22] E.g. Corazza *et al*. (2002), Romdenh-Romluc (2006), and Michaelson and Cohen (2013) draw this conclusion.
[23] See Bach 1992.
[24] See Predelli 2002.

utterance may pragmatically communicate that she went to school in London; semantically, however, it expresses that she went to school in Boston.

One problem with this proposal is that it calls into question the reliability of our linguistic intuitions: if we can't rely on our intuitions when trying to identify the semantic content of Mary's utterance of 'here', why think that we can rely on our intuitions when assessing the original AMP-cases? Indeed, if one is willing to say that the semantic and pragmatic contents of indexical-utterances sometimes come apart, why not say that they come apart in the original cases and endorse the Pragmatic Solution? At least in the original cases, it is apparent that something out of the ordinary is happening.

There is a second, related problem. In connection with the Pragmatic Solution, I pointed out that the consensus in the literature is that the AMP requires a *semantic* solution. However, Predelli's suggestion strays from this: while Intentionalism approaches the original AMP-cases in semantic terms, as desired, Predelli's proposed explanation of Mary's case is of a pragmatic nature. Thus if the AMP really requires a *purely* semantic solution, Predelli's proposal is ruled out[25].

As always, one can try to refine Bach's and Predelli's proposals, but I won't consider any such refinements here. Instead, I shall turn to the Anaphoric Theory of Indexicals.

4. Corazza's Anaphoric Theory of Indexicals

In this section, I shall present Corraza's Anaphoric Theory of Indexicals. In describing his theory, Corazza focusses on 'now' and 'here', but the theory can easily

[25] His proposal may also be vulnerable to the objections raised against the Pragmatic Solution.

be generalised so that it applies to all indexicals[26]. I shall focus on 'now' and 'here' also. Note that all page references in the text refer to Corazza (2004).

4.1 The major tenets

To outline the Anaphoric Theory of Indexicals, return to HISTORY LECTURE and recall that in 2014, the professor utters: 'It is June 22, 1941. Now Hitler begins his invasion of the URSS.' According to Corazza, the professor's utterance of 'now' is *anaphoric* on 'June 22, 1941': it inherits its referent from this noun phrase and therefore refers to June 22, 1941, rather than to a time in 2014. Corazza makes a similar claim about PRISON BREAK: while on the watch tower, the tour guide utters: 'Pete next accessed the infirmary. Here he was spotted by a guard and arrested'. According to Corazza, the guide's utterance of 'here' inherits its referent from the noun phrase 'the infirmary' and hence refers to the prison's infirmary rather than to the watch tower. The next step will be to formulate the Anaphoric Theory in more general terms, by presenting Corazza's two major tenets.

4.1.1 The Anaphoric Interpretation Thesis

As his first major tenet, Corazza proposes what he calls the *Anaphoric Interpretation Thesis* (p. 296):

> AI. When a use of 'now' / 'here' does not select the time / location of the utterance as referent, 'now' / 'here' works like an anaphoric term, inheriting its reference from another noun phrase.

[26] In fact, the theory *must* apply to indexicals other than 'now' and 'here'; otherwise it is not a solution to the AMP in general.

Note that one can understand AI in two ways: firstly, one could read it as saying that the indexicals 'now' and 'here' themselves sometimes function like anaphoric terms, while functioning in accordance with Kaplan's theory at other times. Secondly, one could understand AI as stating that there are two anaphoric expressions in English, 'now' and 'here', which are homonymous with, but distinct from the indexicals 'now' and 'here'. Corazza is not entirely clear on which of the two interpretations he has in mind: he claims at one point that his theory "avoid[s] the need to posit ambiguity" (p. 295), while stating later that "one can maintain that 'now' and 'here' are systematically ambiguous" (p. 300). I shall take AI to mean that the indexicals themselves have anaphoric uses, but nothing depends on this.

Note further that Corazza intends AI as a modification to Kaplan's original theory, and not as a replacement: according to Corazza, Kaplan's theory of indexicals is *correct* if it is restricted to the non-anaphoric uses of 'now' and 'here'. In other words, as Corazza sees it, the AMP does not show that Kaplan's theory is false, but only that there are uses of indexicals to which it does not apply; the Anaphoric Theory gives the correct semantics for these latter uses.

Corazza further explains that according to his theory, 'now' and 'here' are *unbound* anaphors, like 'he' and 'it' in 'John owns a donkey. He beats it.' The reader may wonder whether Corazza has a particular theory of unbound anaphora in mind, but this is not the case. He states: "For the sake of my argument I do not need to commit myself to a specific theory of anaphora. A satisfactory theory of cross-sentential or unbound anaphora will do the job required" (p. 307, ftn. 27). Since the present paper concerns indexicals rather than anaphora, I won't investigate this claim, but shall assume that Corazza can supply a suitable theory of unbound anaphora. The next sub-section presents Corazza's second major tenet.

4.1.2 The Tacit Initiator Thesis

According to AI, the professor's utterance of 'now' in HISTORY LECTURE is anaphoric on 'June 22, 1941', i.e. it is anaphoric on an expression that the professor *uttered*. Yet now consider the following case.

CAESAR. Jane holds a lecture in 2014. She starts with: 'In this lecture, I'd like to talk about Caesar's assassination. It took the conspirators several months to prepare their plan'. She continues:

(9) Now they are ready to carry their scheme into execution.

Intuitively, Jane's utterance of 'now' does not refer to a time in 2014, but rather to March 15, 44 BC (i.e. to the day on which Caesar was killed). Hence by AI, Jane's utterance of 'now' "works like an anaphoric term" (p. 296). However, Jane does not utter any expression that denotes March 15, 44 BC, which means that there is no expression that her utterance of 'now' could be anaphoric on: apparently, no antecedent is present from which the utterance could inherit March 15, 44 BC as referent.

In reply, Corazza holds that an antecedent *is* present in CAESAR, except that the antecedent is *tacit*. More explicitly, Corazza endorses the *Tacit Initiator Thesis* (p. 297):

> TI. If an utterance of 'now' / 'here' refers to time *t* / to location *l*, rather than to the time / location of the utterance, and if no expression has been uttered that denotes *t* / *l*, then the utterance of 'now' / 'here' in question "inherits its referent from a *tacit initiator*", also called a "*tacit antecedent*".

To a first approximation, the idea is this: Jane only says 'Now they are ready to carry their scheme into execution'. Yet semantically, her utterance is to be treated *as if* she *had* uttered 'Now, *on March 15, 44 BC*, they are ready to carry their scheme into execution.' The next sub-section tries to spell this idea out more clearly.

4.2 Tacit initiators

Unfortunately, Corazza does not explain the notion of a tacit initiator very clearly. On the one hand, several of Corazza's remarks indicate that he takes tacit initiators to be *unarticulated expressions*. For instance, AI states that all anaphoric uses of 'now' and 'here' inherit their referents from a *noun phrase*. Hence if some anaphoric uses of 'now' and 'here' inherit their referents from a tacit initiator, it follows that tacit initiators are noun phrases. Elsewhere, Corazza further says that "the tacit initiator is a NP, referring to [a certain] object" (p. 308). On the other hand, Corazza states that tacit initiators are part of what "is presupposed by the discourse situation" (p. 297), and that the "initiators may be understood as what is taken for granted during a speech act." (*ibid*.) This suggests that tacit initiators are *presupposed propositions*. Thus, there are two positions that one can plausibly ascribe to Corazza, PROPOSITION and EXPRESSION.

> EXPRESSION. When an utterance of 'now' / 'here' is used anaphorically, it inherits its referent from an expression, but the reference-endowing expression need not be phonetically realised.
>
> PROPOSITION. When an utterance of 'now' / 'here' is used anaphorically, it either inherits its referent from a (phonetically realised) expression, or – if no such expression is present – from a proposition that is presupposed by the speaker at the time at which the utterance is produced.

To illustrate the difference between EXPRESSION and PROPOSITION, return to CAESAR. According to the former principle, Jane utters (9a) rather than (9), but 'on March 15, 44 BC' remains phonetically unrealised.

(9a) Now, on March 15, 44 BC, they are ready to carry their scheme into execution.

According to the latter principle, Jane utters (9) – which does not contain 'on March 15, 44 BC' as constituent –, but presupposes a proposition from which her utterance of 'now' can inherit March 15, 44 BC as referent[27]. Further, if EXPRESSION is true, Jane's utterance of 'now' is anaphoric on the phonetically unrealised 'on March 15, 44 BC'. If PROPOSITION is true, Jane's utterance of 'now' is anaphoric on her presupposition. On both accounts, Jane's sentence-utterance expresses a statement about March 15, 44 BC.

As indicated, it is not clear from Corazza's paper whether he endorses EXPRESSION or PROPOSITION. I shall assume that he endorses PROPOSITION. For, as the next section shows, there is a reason for rejecting EXPRESSION.

5. EXPRESSION and the Humpty-Dumpty Problem

According to EXPRESSION, Jane utters (9a) rather than (9). However, since 'on March 15, 44 BC' remains unarticulated, her utterance of (9a) sounds exactly the same as ordinary utterances of (9). More generally, if EXPRESSION is true, it is quite common for speakers to utter an indexical-containing sentence in such a way that the utterance looks or sounds like an utterance of a slightly different sentence. E.g. some utterances of (10a) will look or sound like utterances of (10b).

(10a) I went to school here, in Boston.

(10b) I went to school here.

To put the matter slightly differently, if EXPRESSION is true, at least some indexical-utterances give rise to *sentential ambiguity*: when a speaker *seems* to utter 'Now this-and-this happens', her utterance is ambiguous between being an utterance of 'Now this-

[27] Below I will discuss what proposition this might be. For now, the reader may assume that Jane presupposes that Caesar was killed on March 15, 44 BC.

and-this happens', and 'Now, at time *t*, this-and-this happens'. Similarly for utterances of sentences that contain 'here'.

Next, recall the following case which was first presented in connection with Intentionalism: Frank and Mary are in London. He asks her where she went to school and she replies: 'I went to school here'. Her utterance intuitively expresses that she went to school in London. It cannot be read as expressing that she went to school in Boston. However, EXPRESSION implies that her utterance can be read as expressing this latter proposition, whence EXPRESSION faces the Humpty-Dumpty Problem.

In particular, if EXPRESSION is true, Mary's sentence-utterance is ambiguous between being an utterance of (10a) and (10b)[28]. In turn, if Mary's utterance is ambiguous between (10a) and (10b), one would expect that her utterance can be read in two ways, corresponding to the two possible disambiguations: it can not only be read as expressing that she went to school in London, but also as expressing that she went to school in Boston. Thus, EXPRESSION falsely predicts that one can read Mary's utterance as expressing this latter proposition.

The proponent of EXPRESSION could try to adjust her account, by modifying AI, TI, or EXPRESSION itself. However, I won't consider any adjustments of this kind: as will be seen further below, PROPOSITION can avoid the Humpty-Dumpty Problem, which gives it an advantage over EXPRESSION. I shall therefore focus on the former tenet in discussing the Anaphoric Theory.

6. PROPOSITION: Some details

According to PROPOSITION, utterances of 'now' and 'here' can be anaphoric on presupposed propositions. Yet the reader may find this puzzling: normally, (utterances

[28] More precisely, it is ambiguous between *infinitely many* statements.

of) expressions are anaphoric on (utterances of) other *expressions*, and not on propositions. In reply, while I agree that the proponent of PROPOSITION needs to address this point[29], I shall grant, for argument's sake, that expressions can be anaphoric on propositions. This and the next section consider how this idea might be implemented.

To start, note that Corazza endorses a notion of *speaker* presupposition, rather than a notion of *sentence* presupposition. More particularly, he endorses that speaker S presupposes *p* just in case S believes *p*, S believes that her audience believes *p*, and S believes that her audience believes that S believes *p*[30]. Hence when Jane utters (9), her utterance of 'now' is anaphoric on a proposition that *Jane* presupposes, rather than on a proposition that is presupposed by the *sentence* she utters. Moreover, by PROPOSITION, the utterance is anaphoric on a proposition that Jane presupposes *at the time at which she produces the utterance*.

Next, for an utterance of 'now' or 'here' to inherit a certain time or location as referent from a proposition, the proposition must presumably stand in some special relation to the time or location in question. E.g. it would be odd to think that an utterance of 'now' can inherit March 15, 44 BC from the proposition that two plus two equals four. At first glance, one might think that this special relation is constituency, and hold that an utterance of 'now' / 'here' inherits time *t* / location *l* as referent from proposition *p*, only if *t* / *l* is a *constituent* of *p*. However, we shall be slightly more liberal, for reasons that will become apparent in the next section. In particular, we shall allow that an utterance of 'now' / 'here' can inherit *t* / *l* not only from *singular* propositions that contain *t* / *l* as constituent, but also from *particularised general*

[29] Corazza doesn't address this point.
[30] This is Stalnaker's (1974) definition. Corazza cites, and explicitly endorses, this definition (p. 298), but note that nothing will depend on whether Corazza endorses the said definition of speaker presupposition, or an alternative definition.

propositions that are about *t* / *l*[31]. For instance, if *t* is the time at which Caesar died, then *t* can be inherited not only from P_{sing} but also from P_{gen}.

P_{sing}: that Brutus was in Rome at *t*.

P_{gen}: that Brutus was in Rome at the time at which Caesar died.

For ease of reference, let's say that a proposition *supplies* time *t* / location *l* to mean that *t* / *l* can be inherited from the proposition in question. A proposition that supplies *t* / *l* can contain *t* / *l* as constituent, but it needn't.

One shouldn't find it too surprising that the Anaphoric Theory of Indexicals allows 'now' and 'here' to be anaphoric on general propositions. By way of comparison, suppose that I say: 'The farmer owns a donkey. He beats it'. In uttering the first sentence, I do not utter an expression that directly refers to the farmer: I only utter a definite description that denotes the farmer. Hence no expression in the first sentence has the farmer himself as constituent. Still, it is normally thought that 'he' is used anaphorically, and that the utterance of 'he' directly refers to the farmer. This being so, it is fairly natural to think that 'now' and 'here' can be anaphoric on general propositions, at least if one is willing in general to grant that 'now' and 'here' can be anaphoric on propositions.

In sum, for an utterance of 'now' / 'here' to inherit *t* / *l* from a proposition, the proposition must supply *t* / *l*, and must be presupposed by the speaker at the time at which she utters 'now' / 'here'. The next section compares different kinds of propositions that fit this description.

[31] For the difference between singular and general propositions, see e.g. Armstrong and Stanley 2011.

7. PROPOSITION: finding suitable antecedents

Return to CAESAR: Jane utters (9) = 'Now they are ready to carry their scheme into execution', and her utterance of 'now' intuitively refers to March 15, 44 BC. Since Jane does not utter an expression that refers to this date, her utterance of 'now' is anaphoric on a proposition. The proposition in question must supply March 15, 44 BC, and must be presupposed by Jane at the time at which she utters 'now'. As I shall now show, several different propositions fit this description. Yet as I shall also show, not all propositions that fit the description are suitable antecedents for Jane's utterance. After identifying a suitable antecedent for her utterance of 'now', I shall show that PROPOSITION can avoid the Humpty-Dumpty Problem.

7.1 Ordinary and uninformative propositions

At first glance, one might think that Jane's utterance of 'now' is simply anaphoric on an 'ordinary' proposition like P_1 or P_2.

P_1 that Caesar was killed on March 15, 44 BC.

P_2 that on March 15, 44 BC, the conspirators were ready to assassinate Caesar.

Both propositions supply March 15, 44 BC, as desired. Yet whether Jane presupposes P_1 and P_2 at the time at which she utters 'now' greatly depends on how CAESAR is spelled out. Let's focus on P_1; then if Jane is speaking to an audience of classical historians, it is plausible that she presupposes P_1 at the time in question: given her audience's field of expertise, it is plausible that she believes that her audience believes P_1, and also plausible that she believes that her audience believes that she believes P_1. Yet suppose that Jane knows that her audience does not know when Caesar was assassinated. Then her utterance of (9) intuitively still expresses a proposition about 15 March, 44 BC: that the audience does not know when Caesar was assassinated does

not seem to influence what content her sentence-utterance has. However, if the case is spelt out as described, Jane cannot rationally presuppose P_1, as she cannot rationally believe that her audience believes P_1. Thus if the case is spelt out as described, P_1 can't be the proposition that Jane's utterance of 'now' is anaphoric on.

As should be obvious, the foregoing remarks apply to P_2 also. More generally, it seems that CAESAR could be spelled out so that Jane cannot rationally presuppose any 'ordinary' proposition that supplies March 15, 44 BC. This being so, holding that Jane's utterance of 'now' is anaphoric on an 'ordinary' proposition does not seem like a viable option[32].

Next consider P_3 and P_4, which differ from P_1 and P_2 in that they are much less informative. (In that sense, they aren't 'ordinary' propositions.)

P_3 that Caesar was killed on the day he was killed.

P_4 that it once was March 15, 44 BC.

Both propositions supply March 15, 44 BC. More importantly, Jane can rationally presuppose P_3 and P_4 *regardless* of how CAESAR is spelled out. Nevertheless, P_3 and P_4 aren't suitable antecedents for Jane's utterance. The problem is that P_3 and P_4 aren't the only 'uninformative' propositions that Jane presupposes when uttering 'now'. For instance, she plausibly also presupposes P_5 and P_6.

P_5 that Caesar was born on the day he was born.

P_6 that it once was June 3, 44 BC.

Hence if the proponent of PROPOSITION were to maintain that Jane's utterance of 'now' is anaphoric on an 'uninformative' proposition, she would have to explain why it is anaphoric on P_3 or P_4, rather than on P_5 or P_6. Yet this will be difficult to explain: P_3

[32] 'Ordinary' propositions are the only kind of proposition that Corazza considers. He does not address the problems I raise in this section, and I do not know how he would respond to them.

and P_4 don't bear any special relation to Jane's utterance of 'now' that P_5 and P_6 don't also bear to it[33]. Thus, holding that Jane's utterance is anaphoric on a proposition like P_3 or P_4 doesn't seem to be the best option for the proponent of PROPOSITION.

7.2 Utterance-referential propositions

Next consider proposition P_7.

P_7 that at the time Jane utters 'now', she is talking about the time at which the conspirators were ready to carry their scheme into execution.

Since March 15, 44 BC is the time at which the conspirators were ready to carry their scheme into execution, P_7 supplies March 15, 44 BC, as desired. Further, we can be certain that Jane presupposes P_7 when uttering 'now'[34]. Thirdly and most importantly, P_7 bears a special relation to Jane's utterance of 'now', as the utterance is explicitly mentioned in P_7. Hence, the proponent of PROPOSITION could claim that Jane's utterance is anaphoric on P_7.

Let us speak of *utterance-referential* propositions to mean propositions like P_7, and note that utterance-referential propositions always supply *two* times. In particular, P_7 supplies the time at which Jane utters 'now', as well as March 15, 44 BC. Hence if the proponent of PROPOSITION claims that Jane's utterance of 'now' is anaphoric on P_7, she must explain why the utterance inherits March 15, 44 BC, rather than the time at which Jane utters 'now'.

[33] Even if the proponent of PROPOSITION could narrow the set of suitable antecedents down to P_3 and P_4, she would still need to decide between the two, as Jane's utterance of 'now' is presumably anaphoric on *exactly one* proposition.

[34] By uttering (9), Jane informs her audience that she is talking about the time at which the conspirators were ready to carry their scheme into execution. Hence at least if the audience is attentive, they will believe P_7, and believe that Jane believes P_7. Therefore, P_7 avoids the problem that 'ordinary' propositions face: spelling out Caesar in such a way that Jane cannot rationally presuppose P_7, is much harder than spelling out the case in such a way that Jane cannot rationally presuppose P_1 or P_2.

Perhaps the proponent of PROPOSITION can explain this. Suppose I utter: 'The glass fell off the table. It broke.' Then my utterance of 'it' is ambiguous between denoting the glass and the table. However, the reading on which 'it' denotes the glass is more natural than the reading on which 'it' denotes the table. Along the same lines, the proponent of PROPOSITION could claim that Jane's utterance of 'now' is ambiguous between referring to March 15, 44 BC, and the time at which Jane utters 'now'. Further, the reading on which 'now' refers to March 15, 44 BC, is more natural than the reading on which 'now' refers to the time at which Jane utters 'now'.

While the proposed explanation may not be entirely convincing, it is not completely implausible either: utterances of anaphoric terms are sometimes ambiguous; hence if one takes serious the claim that 'now' can function like an anaphoric term, one shouldn't be surprised that utterances of 'now' are sometimes ambiguous. In other words, *within the framework of the Anaphoric Theory*, the proposed explanation is plausible enough.

In sum, the proponent of PROPOSITION now holds that Jane's utterance of 'now' is anaphoric on P_7. Before moving on, let's generalise this tenet. The best version of PROPOSITION holds that all anaphoric utterances of 'now' and 'here' (that aren't anaphoric on an expression) are anaphoric on *the same kind* of proposition. As CAESAR shows, it is not plausible that all of the utterances in question are anaphoric on 'ordinary' or 'uninformative' propositions. Therefore, the proponent of PROPOSITION will hold that only utterance-referential propositions are suitable antecedents: if an utterance of 'now' or 'here' is anaphoric on a proposition, it is anaphoric on an utterance-referential proposition. As the next sub-section shows, this allows the proponent of PROPOSITION to avoid the Humpty-Dumpty Problem.

7.3 PROPOSITION and the Humpty-Dumpty Problem

Recall that Mary and Frank are in London and that Mary utters 'I went to school here'. As the case was spelled out, Mary does not utter an expression that her utterance of 'here' could be anaphoric on. Hence if her utterance is anaphoric at all, it is anaphoric on a proposition. Given the results from the previous sub-section, this proposition would furthermore have to be an utterance-referential proposition. Hence according to the current version of PROPOSITION, Mary's sentence-utterance expresses that Mary went to school in Boston, only if Mary presupposes an utterance-referential proposition that supplies Boston. P_M is a proposition of this kind.

P_M that at the time Mary utters 'here', she is talking about Boston.

However, Mary cannot rationally presuppose P_M at the time at which she utters 'here'. At the time in question, Mary hasn't given Frank any indication that she is talking about Boston rather than London. Hence she cannot plausibly believe that Frank believes that she is talking about Boston. Since a speaker presupposes *p* only if she believes that the audience believes *p*, this means that Mary cannot rationally presuppose P_M.

In sum, by insisting that only utterance-referential propositions are suitable antecedents, the proponent of PROPOSITION can avoid the Humpty-Dumpty problem. So far, then, PROPOSITION offers a promising solution to the AMP. However, as the next section shows, PROPOSITION and the Anaphoric Theory fail in the end.

8. PROPOSITION: lack of suitable antecedents

To repeat, the proponent of PROPOSITION now claims that only utterance-referential propositions are suitable antecedents. However, there are instances of the AMP in

which the speaker cannot rationally presuppose an utterance-referential proposition. The following case is an example.

GERMANIC LIMES. Tom wishes to discuss a historically important attack on the Germanic Limes, call it the First Germanic Invasion. He is aware, however, that his audience does not know that this attack will be the topic of the lecture. He begins by uttering 'In this lecture, I'd like to describe the Germanic Limes, between 100 and 110 AD. On most days, everything is quiet on both sides of the border.' He continues:

(11) But now the Roman soldiers see enemy troops advance. Now the First Germanic Invasion begins.

Intuitively, both of Tom's utterances of 'now' refer to the time at which the First Germanic Invasion begins. Hence by PROPOSITION, at least his first utterance of 'now' is anaphoric on a proposition that supplies this time[35]. More particularly, since only utterance-referential propositions are suitable antecedents, the utterance must be anaphoric on a proposition like P_8 or P_9.

P_8 that at the time Tom utters 'now', he is talking about the time at which the First Germanic Invasion begins.

P_9 that at the time Tom utters 'now', he is talking about the time at which the Roman soldiers see enemy troops advance.

However, Tom cannot rationally presuppose P_8 at the time at which he produces his first utterance of 'now'. For at the time in question, he has not yet informed his audience that he will be talking about the First Germanic Invasion. Therefore, he cannot rationally believe that his audience believes P_8.

[35] If the Anaphoric Theory is correct, his second utterance of 'now' is plausibly anaphoric on his first utterance of 'now'.

Turning to P_9, GERMANIC LIMES can be spelled out in such a way that Tom cannot presuppose this proposition either. Suppose that the time at which the First Germanic Invasion begins is just one of *many* times at which the Roman soldiers see enemy troops advance. Then Tom cannot presuppose that he is talking about *the* time at which the Roman soldiers see enemy troops advance, simply because there is no unique time that fits the description. Thus if the case is spelled out as described, Tom cannot rationally presuppose P_9.

There is a different kind of utterance-referential proposition that is worth considering, however.

P_{10} that when Tom utters 'now', he is talking about the time he is talking about when uttering 'now'.

While *the* time at which the Roman soldiers see enemy troops advance does not exist, *the* time that Tom is talking about when uttering 'now' apparently does exists: when Tom utters 'now', he is talking about the time at which the First Germanic Invasion began, and about no other time. Further, Tom's audience plausibly believes P_{10}, and believes that Tom believes P_{10}. Thus, P_{10} is a plausible candidate for being the proposition that Tom's utterance of 'now' is anaphoric on.

Let us say that a proposition like P_{10} is *empty* (in addition to being utterance-referential). Then the proponent of PROPOSITION could refine her tenet and hold that empty, utterance-referential propositions are suitable antecedents. Let's assess this proposal.

One problem with the proposal at hand is that it leads back to the Humpty-Dumpty Problem. If empty propositions are suitable antecedents, then P_M* is a suitable antecedent for Mary's utterance of 'here'.

P_M* that when Mary utters ‘here’, she is talking about the location she is talking about when uttering ‘here’.

In turn, if P_M* is a suitable antecedent for Mary’s utterance ‘here’, the proponent of PROPOSITION can no longer explain why Mary’s utterance of ‘I went to school here’ can’t be read as expressing that Mary went to school in Boston: Mary can rationally presuppose P_M* when uttering ‘here’. Further, Mary is talking (or, trying to talk) about Boston when uttering ‘here’, whence P_M* supplies Boston. Thus, P_M* is a suitable antecedent which supplies Boston and is presupposed by Mary at the time at which she utters ‘here’. By PROPOSITION, that means that her utterance of ‘here’ refers to Boston, and that her sentence-utterance expresses that she went to school in Boston.

The proposal at hand faces a second problem. To see it, consider the following case.

GUMMY BEARS. Emma starts her lecture with: ‘In this lecture, I’d like to talk about the Sweet Tooth Problem. The French politicians noticed this problem as early as March 1856, but couldn’t agree on a solution.’ She continues:

(12) Now the conservatives propose the law that I already mentioned in yesterday’s lecture: the gummy bear law.

According to the current proposal, Emma’s utterance of ‘now’ is anaphoric on an empty, utterance-referential proposition. The desired proposition for Emma’s utterance is P_{11}.

P_{11} that when Emma utters ‘now’, she is talking about the time she is talking about when uttering ‘now’.

However, it is not clear that Emma can rationally presuppose P_{11}. The problem is similar to the problem encountered in connection with (11) and P_8: as P_{11} concerns *the* time that Emma is talking about when uttering ‘now’, Emma cannot presuppose P_{11}

unless she is talking about *exactly one* time when uttering 'now'. Yet in a fairly natural sense of *talking about*, she is talking about *two* times when uttering 'now': firstly, she is talking about the time at which the conservatives proposed the gummy bear law; secondly, she is talking about 'yesterday', i.e. about the day before the day on which the lecture takes place. Thus unless the proponent of PROPOSITION can specify a sense of *talking about* in which Emma is talking about just one time, she cannot hold that Emma presupposes P_{11}.

Every utterance of 'now' refers to exactly one time. So, there is a sense of *talking about* in which speakers talk about just one time when they utter 'now': they talk about the time their utterance refers to. Yet the proponent of PROPOSITION cannot make use of this: she wants to define the content of Emma's utterance of 'now' in terms of the time that Emma (presupposes to be) talking about. Therefore, she cannot also define the time that Emma (presupposes to be) talking about in terms of the content of Emma's utterance of 'now': such a definition would be circular. Put slightly differently, as far as the proponent of PROPOSITION is concerned, Emma's utterance of 'now' does not have a content until a proposition has been specified that her utterance can be anaphoric on. Therefore, one cannot specify this proposition by considering the content of Emma's utterance.

In sum, the proponent of PROPOSITION doesn't seem able to specify a sense of *talking about* in which Emma is talking about just one time when she says 'now'[36]. Therefore, she cannot hold that Emma presupposes P_{11} when uttering 'now'. This, in turn, means that Emma's utterance of 'now' is not anaphoric on P_{11}. More generally, the claim that

[36] Perhaps it will be suggested that Emma presupposes that she is talking about the time to which she *intends* her utterance of 'now' to refer. Yet this proposal will also fail: suppose Emma *falsely* believes that the conservatives proposed the law in 1756. Then she intends to refer to 1756, but she also intends to refer to the time at which the conservatives in fact proposed the law. Thus, *the* time to which Emma intends her utterance of 'now' to refer does not exist, which means that she cannot presuppose that she is talking about *the* time to which she intends her utterance of 'now' to refer.

anaphoric utterances of ‘now’ and ‘here’ are anaphoric on empty, utterance-referential propositions fails.

Let me conclude the discussion of PROPOSITION. The tenet can account for cases like CAESAR and it potentially avoids the Humpty-Dumpty Problem. Yet it fails to account for cases like GERMANIC LIMES and GUMMY BEARS: while the said cases are apparently instances of the AMP, no proposition emerges that could serve as antecedent. I therefore conclude that PROPOSITION and the Anaphoric Theory fail.

9. Strengthening the paradox

In this final section, I shall draw attention to a previously unnoticed feature of the AMP: roughly, whether an utterance of ‘now’ / ‘here’ can be read as referring to a time in the past / to a far-away location is a very subtle matter, and none of the theories presented in this paper are fine-grained enough to account for this subtlety.

9.1 The data

Suppose that the criminology lecturer describes a famous bank robbery. At 3 pm, she says: ‘As guards are patrolling the bank’s premises day and night, Rob the robber has to time his movements very carefully.’ She continues:

(13a) It is 1 am. Rob can now enter the bank’s premises unseen.

The utterance at hand is a typical instance of the AMP: intuitively, it expresses that Rob can enter the bank’s premises at 1 am (on the day of the robbery). By contrast, consider the following sentence.

(13b) At 1 am, Rob can now enter the bank’s premises unseen.

If the lecturer utters (13b) instead of (13a), her utterance isn't easily read as expressing that Rob can enter the bank's premises at 1 am: at least according to my intuitions, the relevant utterance of 'now' refers to 3 pm, i.e. to the time at which the lecturer speaks. If this is correct, there is an interesting and surprising difference between the lecturer's utterances of (13a) and (13b), a difference that theories of indexicals must explain. Yet before I consider possible explanations, I'd like to comment on my intuitions regarding (13b).

Upon reading (13b), most speakers I consulted agreed that the relevant utterance of 'now' cannot be read as referring to 1 am: they shared my intuition that it refers to the time at which the lecturer speaks[37]. Still, to strengthen this intuition, I'd like to present a few sentences that are similar to (13b).

In his *Demonstratives*, and more particular in the chapter on Monsters, Kaplan considers (14a) and points out that utterances of this sentence cannot be read as expressing proposition M_1 (with 'M' for 'Monster')[38].

(14a) In some contexts, it is true that I am not tired now.
M_1: that for some contexts *c*, the speaker of *c* is tired at the time of *c*.

If the sentence can be interpreted at all, it expresses (in Kaplan's mouth) that Kaplan isn't tired at the time at which he utters the sentence; 'in some contexts, it is true that' is redundant. Kaplan concludes that one cannot operate on character, but this conclusion needn't concern us here. What matters is that virtually everybody[39] shares Kaplan's intuition that utterances of (14a) can't express M_1. With this in mind, consider (14b) and M_2.

(14b) In the context that Smith was in at 3 pm, it is true that I am not tired now.
M_2: that Smith isn't tired at the time of the context that Smith was in at 3 pm.

[37] The same speakers agreed that if the lecturer utters (13a), her utterance of 'now' refers to 1 am.
[38] See Kaplan 1989, p. 510.
[39] Including Schlenker (2003), who argues *pace* Kaplan that one *can* operate on character.

Since (14a) cannot be read as expressing M_1, I'd be very surprised if speakers felt they can read (14b) as expressing M_2. Of course, 'in some contexts, it is true that' attempts to *quantify* over contexts while 'the context that Smith was in at 3 pm' describes a *particular* context. Still, as far as I can tell, (14b) cannot express M_2 – no more than (14a) can express M_1.

Note that 'the context that Smith was in at 3 pm' fully describes a context; but in describing it, the context's speaker and time are explicitly mentioned. Nevertheless, 'I' and 'now' don't acquire this speaker and time as referents. Now consider the following sentence.

(14c) Last Thursday at 1 am, it was true that Rob is not tired now.

Unlike (14b), (14c) doesn't fully describe a context; like (14b), it explicitly mentions a time. Yet since 'now' in (14b) doesn't acquire 3 pm as referent, it would be rather surprising if (one had the intuition that) 'now' in (14c) acquired 1 am as referent: reducing the full description of a context to a description of a point in time shouldn't impact the referent of 'now'[40]. Hence (14c) doesn't express M_3.

M_3: that Rob wasn't tired last Thursday at 1 am.

Next consider (14d) and M_4:

(14d) Last Thursday at 1 am, it was true that Rob can now enter the bank's premises unseen.

M_4: that Rob was able to enter the bank's premises unseen last Thursday at 1 am.

The only difference between (14c) and (14d) is that they embed different, but structurally very similar sentences under 'last Thursday at 1 am, it was true that'. One would expect, therefore, that the intuitions we have about (14c) carry over to (14d) Hence (14d) doesn't express M_4.

[40] If anything, it should make it *less* likely for 'now' to receive its referent from a context other than the context of utterance.

Now return to (13b) = 'At 1 am, Rob can now enter the bank's premises unseen'. Once one has acquired the intuition that (14c) and (14d) can't express M_3 and M_4, the intuition that 'now' in (13b) doesn't refer to 1 am should come very naturally. After all, the only difference between (13b) and (14d) is that 'last Thursday' and 'it was true that' have been deleted. In turn, it is difficult to imagine that the deletion of these elements impacts our intuitions or the referent of 'now'. In sum, the chain from (14a) to (14d) bolsters the claim that in the lecturer's mouth, (13b) doesn't express that Rob can enter the bank's premises at 1 am. Rather, the relevant utterance of 'now' refers to 3 pm, i.e. to the time at which the lecturer speaks.

As mentioned above, the difference between (13a) and (13b) needs to be explained. Yet as (13c) shows, such an explanation won't be easy to give.

(13c) At 1 am, the guards change, and Rob can now enter the bank's premises unseen.

For, if the lecturer utters (13c), her utterance of 'now' intuitively refers to 1 am again; the sentence-utterance intuitively expresses that the guards change at 1 am, and that Rob can enter the bank's premises *at 1 am*[41]. Hence although both (13b) and (13c) open with 'at 1 am', 'now' acquires 1 am as referent only in one of the two sentences. Theories of indexicals must explain why this is so.

Before considering possible explanations, note briefly that one can construct a similar case with 'here': The criminology lecturer says: 'After cutting through the wire fence, Rob the robber can enter the bank's premises.' Now consider the following three sentences.

(15a) Rob first needs to cross the yard. Here he has to move especially quietly.

[41] To tease out this intuition, suppose the lecturer says:

(13c*) At 1 am, the guards change. Rob can now enter the bank's premises unseen.

If one has the intuition that 'now' in (13a) refers to 1 am, I think one will also find it intuitive that 'now' in (13c*) refers to 1 am. Further, I believe that this intuition survives if we conjoin the two sentence in (13c*) and form the conjunction of (13c).

(15b) In the yard, Rob has to move especially quietly here.

(15c) In the yard, a guard is on patrol, so that Rob has to move especially quietly here.

If the lecturer utters (15a), her utterance intuitively expresses that that Rob has to move especially quietly in the yard. By contrast, an utterance of (15b) would difficult to interpret: If it can be interpreted at all, it expresses that in the yard, it is the case that Rob has to move quietly in the lecture theatre. Yet as before, a different reading is available for (15c): if the lecturer utters this sentence, her utterance of 'here' intuitively refer to the yard; the sentence-utterance intuitively expresses that Rob has to move especially quietly *in the yard*, because a guard is on patrol there.

With the data presented, let's ask whether it can be explained by any of the previously presented theories.

9.2 Explaining the data

The challenge is twofold. Each theory must explain why in the lecturer's mouth, (13b) cannot be read as expressing P_{12}, but only as expressing P_{13}. At the same time, each theory must explain why the lecturer's utterance of (13c) is most naturally read as expressing P_{14}.

P_{12} that Rob can enter the bank's premises at 1 am.

P_{13} that at 1 am, it is the case that Rob can enter the bank's premises at 3 pm.

P_{14} that the guards change at 1 am, and that Rob can enter the bank's premises at 1 am.

At least prima facie, none of the five theories that were presented earlier can meet the challenge. To show this, consider first what the theories have to say about (13c).

The Anaphoric Theory: The lecturer's utterance of 'now' is anaphoric on 'at 1 am' and inherits this time as referent. Her sentence-utterance therefore expresses P_{14}.

Intentionalism: The lecturer intends her utterance of 'now' to refer to 1 am. Her sentence-utterance's context of intention therefore has 1 am as time. As (13c) expresses P_{14} at such a context, her sentence-utterance expresses P_{14}.

The Pragmatic Solution: The lecturer's sentence-utterance semantically does not express P_{14}, but merely communicates this proposition on a pragmatic level. Semantically, it expresses that the guards change at 1 am, and that Rob can enter the bank's premises at 3 pm.

The Demonstrative Solution: In uttering (13c), the lecturer utters the demonstrative 'now', rather than the indexical 'now'. Further, she supplies a demonstration that picks out 1 am. Therefore, her utterance expresses P_{14}.

The Character Solution: The character of 'now' maps the context in which the lecturer's utterance of (13c) occurs to 1 am. As (13c) expresses P_{14} at such a context, her sentence-utterance expresses P_{14}.

With these claims in place, let's consider each theory in turn and ask whether it can account for the fact that the lecturer's utterance of (13b) can only be read as expressing P_{13}. Starting with the Anaphoric Theory of Indexicals, its proponent could try to explain the data in syntactic terms. When Kaplan discusses sentences like (13b), he treats expressions like 'at 1 am' as operators[42]. For argument's sake, let's accept that this is correct. Then the syntactic form of (13b) is: Oφ. By contrast, it is not so clear that Oφ is also the form of (13c): arguably, 'at 1 am' operates only on 'the guards change'. If this is correct, then 'now' occurs within the scope of 'at 1 am' in (13b), while occurring outside its scope in (13c).

Now, according to the Anaphoric Theory, the lecturer's utterance of 'now' is anaphoric on 'at 1 am' in the case of (13c). Therefore, the theory's proponent cannot

[42] See Kaplan 1989a, pp. 510 - 512.

maintain that utterances of ‘now’ cannot be anaphoric on operators. However, she could hold that utterances of ‘now’ cannot be anaphoric on operators *under which they are embedded*. This claim implies that the lecturer’s utterance of ‘now’ isn’t anaphoric on ‘at 1 am’ in the case of (13b), without implying that the lecturer’s utterance of ‘now’ isn’t anaphoric on ‘at 1 am’ in the case of (13c).

The proposal is not very plausible: according to the Anaphoric Theory, anaphoric uses of ‘here’ and ‘now’ work just like anaphoric uses of ‘he’ or ‘it’. However, utterances of ‘he’ and ‘it’ *can* be anaphoric on operators under which they are embedded. E.g. ‘he’ can be read as being anaphoric on ‘Bill’ in: ‘In Bill’s opinion, he is the best athlete’.

As an alternative, the proponent of the Anaphoric Theory could hold that neither utterance of ‘now’ is anaphoric on ‘at 1 am’, and try to explain the data in terms of presuppositions: in the case of (13c), the lecturer presupposes a proposition that supplies 1 am; in the case of (13b), no such proposition is presupposed. Therefore, the former utterance of ‘now’ is anaphoric and refers to 1 am, while the latter utterance of ‘now’ is not anaphoric and refers to 3pm.

In reply, there is no reason to think that the lecturer can presuppose a suitable proposition in the case of (13c), but not in the case of (13b). For instance, it seems that she can presuppose in both cases that she is talking about 1 am when uttering ‘now’. In sum, the Anaphoric Theory cannot explain the data.

Next consider Intentionalism. According to this theory, the lecturer’s utterance of (13c) expresses P_{14} because the lecturer intends her utterance of ‘now’ to refer to 1 am. However, if the lecturer can intend this in the case of (13c), there is no reason to think that she cannot have the same intention in the case of (13b). In turn, if she intends to refer to 1 am when uttering (13b), her sentence-utterance’s context of intention has 1 am

as time. At such a context, (13b) expresses P_{12} rather than P_{13}, and the lecturer's sentence-utterance consequently expresses P_{12}.

The Intentionalist could try to adjust her account: a sentence-utterance must be evaluated relative to its context of intention, *unless* the uttered sentence is of the form of Oφ, in which case the sentence-utterance must be evaluated at its context of origin. Since (13b) is of the form of Oφ, while (13c) is not, this would ensure that the correct contents are assigned to both of the lecturer's utterances. However, the proposal is not very convincing. In fact, it is false by the Intentionalist's own lights. Suppose that in 2014, the history lecturer talks about the Storming of the Bastille:

(16) It is July 15, 1789. Today, the Republicans celebrate yesterday's success.

The second sentence in (16) is of the form of Oφ. Hence according to the current proposal, the lecturer's utterance of this sentence must be evaluated at its context of origin. However, if the utterance is so evaluated, it isn't assigned the content that it intuitively expresses, viz. that on July 15, 1789, the Republicans celebrate the success of July 14, 1789. To assign this content to the lecturer's utterance, the Intentionalist must hold that the utterance's context of intention has July 15, 1789 as time, and must evaluate the utterance at this context. Thus, the above proposal fails by the Intentionalist's own lights.

Turning to the Pragmatic Solution, the dialectic situation is slightly different: the proponent of this theory agrees that the lecturer's utterance of (13b) expresses P_{13}, but denies that the utterance of (13c) expresses P_{14}: she holds that it expresses P_{15}, while merely communicating P_{14}.

P_{15} that the guards change at 1 am, and that Rob can enter the bank's premises at 3 pm.

The question is whether the proponent of the Pragmatic Solution can explain why utterances of (13c) can communicate P_{14}, although utterances of (13b) cannot communicate P_{12}. Yet no such explanation seems available. In fact, one of the Gricean Maxims suggests that utterances of (13b) *can* communicate P_{12}: the Maxim of Quality requires speakers to make only true statements. Yet P_{13} is patently false. Hence when the lecturer utters (13b), it can't be her aim to communicate P_{13}. By contrast, it could be her aim to communicate P_{12} as this proposition is both true and relevant. Thus, the Pragmatic Solution seems to predict, falsely, that utterances of (13b) can communicate P_{12}.

The proponent of the Pragmatic Solution could try to invoke the Maxim of Manner, which requires speakers to be as brief as possible: (13b) is slightly longer than 'At 1 am, Rob can enter the bank's premises unseen'. Hence, since the lecturer utters (13b), rather than the slightly shorter sentence, she must be trying to communicate something that one cannot communicate with the shorter sentence. However, one can communicate P_{12} with the shorter sentence. Thus, the lecturer is trying to communicate something other than P_{12}.

In reply, even if the Maxim of Manner warrants the above derivation, it does not yield a satisfying explanation of the data. For, one also needn't utter (13c) to communicate P_{14}: an utterance of the slightly shorter 'At 1 am, the guards change, and Rob can enter the bank's premises unseen' apparently also communicates this content. Hence if the proponent of the Pragmatic Solution endorses the above derivation, she must also conclude that the lecturer's utterance of (13c) doesn't communicate P_{14}, and that conclusion is false.

Turning briefly to the Demonstrative Solution, its proponent would have to maintain that the lecturer utters the indexical 'now' in uttering (13b), while uttering the

demonstrative ‘now’ in uttering (13c). Yet as far as I can see, there is no reason to maintain this. For instance, it is not true that demonstratives cannot be embedded under operators: ‘At 4 pm, John realised that that wasn’t the time of his doctor’s appointment’. Similarly, it is not plausible that the lecturer cannot demonstrate 1 am when uttering (13b), although she can demonstrate 1 am when uttering (13c). In sum, the Demonstrative Solution cannot account for the data.

Lastly, the Character Solution holds that the character of ‘now’ is such that:

(i) It maps the context in which the lecturer’s utterance of (13b) occurs to 3 pm.

(ii) It maps the context in which the lecturer’s utterance of (13c) occurs to 1 am.

While these statements may be true, they don’t *explain* why the utterances of (13b) and (13c) have the contents they have. As stated in section three, the Character Solution doesn’t amount to a genuine semantic theory of indexicals, but is merely a record of utterance-contents. Therefore, (i) and (ii) at best *describe* the semantic difference between (13b) and (13c), but do not account for it.

In sum, at least at first glance the five theories at hand are unable to explain the semantic difference between (13b) and (13c). And while a more detailed study would be needed to confirm this result, it seems correct to say that sentence-pairs like (13b) and (13c) considerably strengthen the Answering Machine Paradox.

9.3 More data

As (13b) and (13c) show, it is a rather subtle matter whether an utterance of ‘now’ can be read as referring to a time in the past. (14b) and (14c) establish a similar result for ‘here’. To conclude this section, let me present a few more sample sentences, which may help to bring out just how subtle the matter is. The context is the same as before:

the criminology lecturer says at 3 pm: ‘As guards are patrolling the bank’s premises day and night, Rob the robber has to time his movements very carefully’. Now consider:

(13d) At 1 am, knowing that the guards are changing, Rob now enters the bank’s premises.

(13e) At 1 am, knowing that the guards are now changing, Rob enters the bank’s premises.

(13f) At 1 am, Rob realises that he can now enter the bank’s premises.

(13g) At 1 am, Rob realises that the guards are changing and that he can now enter the bank’s premises.

(13h) At 1 am, every guard on duty is now asleep.

(13i) At 1 am, every guard who is now on duty is asleep.

(13j) At 1 am, everywhere John looks, a guard is now on patrol.

(13k) At 1 am, Rob can get to the bank unseen if he moves now.

(13l) At 1 am, Rob can now get to the bank unseen if he moves quickly.

If the lecturer utters (13d), her utterance of ‘now’ cannot be read as referring to 1 am. It can be so read if the lecturer utters (13e). Hence how the sentence-utterance can be interpreted isn’t just a matter of how far removed ‘now’ is from ‘at 1 am’.

If the lecturer utters (13f), it is not quite clear whether her utterance of ‘now’ can be read as referring to 1 am: the reading is not entirely natural, but more readily available than in the case of (13b). This suggests that embedding ‘now’ under a propositional-attitude verb makes a difference. If the lecturer utters (13g), her utterance of ‘now’ can be read as referring to 1 am.

Involving a quantifier doesn’t appear to make a difference: if the lecturer utters (13h), (13i), or (13j), her utterance of ‘now’ cannot be read as referring to 1 am, at least according to my intuitions. I do not have any clear intuitions about (13k) and (13l).

To draw any definitive conclusion about the semantics of 'now', one would need to analyse the presented data in much more detail. Yet even without such an analysis, the data makes clear that the semantics of indexicals is even more complex than the original instances of the AMP suggest.

References

Armstrong, J. and Stanley, J. (2011). 'Singular Thoughts and Singular Propositions', *Philosophical Studies* **154**, 205 – 222.

Bach, K. (1992). 'Giorgione was So-Called Because of His Name' *Philosophical Perspectives* **16**, 73–103.

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

Corazza, E. (2004). 'On the Alleged Ambiguity of 'Now' and 'Here'', *Synthese* 138, 289 – 313.

Kaplan, D. (1989a). 'Demonstratives', in J. Almong, J. Perry, and H. Wettstein (eds.), 'Themes from Kaplan', Oxford University Press 1989, 481 – 563.

Kaplan, D. (1989b). 'Afterthoughts', in J. Almong, J. Perry, and H. Wettstein (eds.), 'Themes from Kaplan', Oxford University Press 1989, 565 – 614.

Levinson, S. (2000). 'Presumptive Meanings: The Theory of Generalized Conversational Implicature', Cambridge, MA: MIT Press.

Michaelson, E. and Cohen, J (2013). 'Indexicality and The Answering Machine Paradox', *Philosophy Compass* **8**, 580 – 592.

Mount, A. (2008). 'The impurity of "pure" indexicals', *Philosophical Studies* 138, 193 – 209.

Perry, J. (2003). 'Predelli's threatening note: contexts, utterances, and tokens in the philosophy of language', *Journal of Pragmatics* **35**, 373 – 387.

Predelli, S. (1998). 'I am not here now', *Analysis* **58**, 107 – 115.

Predelli, S. (2002). 'Intentions, Indexicals, and Communication', *Analysis* **62**, 310 – 316.

Predelli, S. (2011). 'I am still not here now', *Erkenntnis* **74**, 289 – 303.

Romdenh-Romluc, K. (2006). 'I', *Philosophical Studies* **128**, 257 – 283.

Schlenker, P. (2003). 'A plea for Monsters', *Linguistics and Philosophy* **26**, 29 – 120.

Sidelle, A. (1991). ‘The Answering Machine Paradox’, *Canadian Journal of Philosophy* **21**, 525 – 539.

Stalnaker, R. (1974). ‘Pragmatic Presuppositions’, in his ‘Context and Content’, Oxford University Press, 47 – 62.

Compounds in Context

There is evidence that at least some NN-compounds (noun-noun compounds) are context-sensitive. For instance, one can use 'toy shop' to denote a shop that sells toys, but can also use it to denote the toy model of a shop. Similarly, one can use 'carbon filter' to denote a device for filtering carbon, but can also use it to denote a device that is made of carbon and used for filtering e.g. smoke. 'Library book' can denote a book that belongs to a library, and can – like 'philosophy book' – also denote a book that is about libraries.

Now consider the Principle of Compositionality, according to which the content of any complex expression *e*, relative to a context *c*, is determined by *e*'s structure and the contents that *e*'s constituents have at *c*. Assuming that NN-compounds are complex expressions, context-sensitive NN-compounds are arguably a counterexample to Compositionality: suppose that 'carbon filter' is indeed context-sensitive, so that there are contexts at which 'carbon filter' denotes a device for filtering carbon, and contexts at which it denotes a filter that is made of carbon. Plausibly, 'carbon' denotes *carbon* in both kinds of contexts and 'filter' denotes a *filter* in both kinds of contexts. Yet if this is correct, the contents of 'carbon' and 'filter' do not co-vary with the content of 'carbon filter'. Hence what content 'carbon filter' has at a given context is not determined by the structure of 'carbon filter' and the contents that 'carbon' and 'filter' have at the context in question, contrary to the Principle of Compositionality. Thus, context-sensitive NN-compounds like 'carbon filter' are a counterexample to Compositionality, at least prima facie.

Call the argument in the previous paragraph the Compound Argument against Compositionality. In the present paper, I shall defend Compositionality against the

Compound Argument. Section one spells the argument out in more detail. The remaining sections discuss its premises and ask whether any of them can be rejected.

For ease of reference, I shall speak of the *Attacker* to mean the proponent of the Compound Argument; I shall speak of the *Defender* to mean the proponent of Compositionality. Further, since NN-compounds are the only compounds I consider, I shall use the term 'compound' to mean NN-compounds.

1. The Compound-Argument against Compositionality

In presenting the Compound Argument, I shall focus on 'carbon filter' but one could develop the argument using other compounds, such as 'toy shop' or 'library book'. To start, the Compound Argument relies on the following three premises[1]:

(C1) The compound 'carbon filter' is a complex rather than simple expression, and 'carbon' and 'filter' are its only constituents.

(C2) The compound 'carbon filter' is context-sensitive, denoting in some contexts a filter that is made of carbon, while denoting in other contexts a filter for purifying carbon.

(C3) The contents of 'carbon' and 'filter' do not co-vary with the content of 'carbon filter'.

Let me begin by spelling these premises out in some more detail.

1.1 The premises of the Compound Argument

Starting with (C1), note that the premise is twofold: its first part is the claim that 'carbon filter' is syntactically complex, that it is a phrase rather than a word. The second

[1] Premise (C3) will be slightly adjusted below.

part is the claim that ‘carbon filter’ has exactly two constituents – viz. the words ‘carbon’ and ‘filter’. One can combine the two claims by saying that the syntactic tree for ‘carbon filter’ has exactly two terminal nodes, viz. ‘carbon’ and ‘filter’.

Note that (C1) is not trivial. Being spelled with a space between ‘carbon’ and ‘filter’, ‘carbon filter’ certainly *looks* like a phrase. However, it is not clear that spelling is a reliable guide to syntactic structure, for at least two reasons. Firstly, idioms like ‘kick the bucket’ or ‘spill the beans’ look like phrases. Yet many theorists hold that they are unstructured lexical items[2]. If this is correct, it isn’t true in general that an expression is a phrase if it looks like one. Secondly, many compounds – such as ‘wheelchair’ and ‘policeman’ – are spelled without a space and therefore look like words. At the same time, one might think that all compounds are syntactically alike: either all compounds are words, or all are phrases[3]. If this is correct, then spelling *can’t* be a guide to the syntactic structure of compounds: if all compounds are phrases, the spelling of e.g. ‘wheelchair’ is misleading; and if all compounds are words, the spelling of ‘carbon filter’ is misleading. In sum, that ‘carbon filter’ is spelled with a space doesn’t automatically mean that ‘carbon filter’ is a phrase: one cannot infer (C1) just from looking at ‘carbon filter’.

There is another reason why (C1) is not trivial. The premise is true only if ‘carbon filter’ is a phrase with *exactly two* constituents. Yet – if the compound is a phrase at all – it is conceivable that ‘carbon filter’ contains *three* constituents: ‘carbon’, ‘filter’, and a third constituent, which is not realised at the compound’s surface form. The idea would be that the third constituent is a hidden, context-sensitive variable; it has the *made-of*-relation as value in contexts in which ‘carbon filter’ denotes a filter that is made of

[2] E.g. Katz (1973), Chomsky (1980), and Davies (1983) hold this. To give another example, pub names like ‘The Red Lion’, ‘The Three Horseshoes’, and ‘The Two Brewers’ also look like phrases although they are arguably words.
[3] Bauer (1998) defends this position.

carbon, while having the *used-for*-relation as value in contexts in which 'carbon filter' denotes a device for filtering carbon. If this is correct, (C1) fails.

The idea that 'carbon filter' contains a hidden variable will be considered more closely below; I shall also reconsider the connection between spelling and syntactic structure. For now, it suffices to note that (C1) is not trivial.

Turning to (C2), the Attacker can support this premise with what is often called a *Context-Shifting Argument*, or CSA[4]. The argument can be developed thus: John wants to introduce his friend Henry to pipe smoking, but Henry is worried that smoking is too unhealthy. John knows that the pipe company Vauen produces small devices that are made of carbon and that can be inserted into pipes to reduce the harmful substances in the smoke. So to ease Henry's worries, John asserts (1a). Intuitively, his utterance is true.

(1a) Vauen produces carbon filters.

Now suppose that two chemists, Jane and Mary, have a carbon sample that they want to purify. Jane knows that there are devices that can be used to achieve this, and wonders whether she and Mary can buy such a device from Vauen. Yet Mary tells her:

(1b) Vauen doesn't produce carbon filters.

Since Vauen indeed doesn't produce the kind of device that Jane and Mary are after, Mary's utterance of (1b) is intuitively true.

Now, if (1a) and (1b) are context-*in*sensitive, then John's utterance expresses a true proposition iff Mary's utterance expresses a false proposition. In other words, if (1a) and (1b) are context-insensitive, John's and Mary's utterances cannot *both* be true. Yet both utterances *are* true, whence (1a) and (1b) are context-sensitive. Further, since 'Vauen' and 'produces' plausibly express the same contents in both John's and Mary's

[4] E.g. Cappelen and Lepore (2005) use this term.

contexts, it follows that 'carbon filter' expresses different contents in the two contexts, i.e. it follows that 'carbon filter' is context-sensitive. Finally, it is plausible that in the first context, 'carbon filter' denotes a filter that is made of carbon, while denoting a filter for purifying carbon in the second context. Thus, the CSA establishes (C2): 'carbon filter' is context-sensitive and denotes in some contexts a filter that is made of carbon, while denoting in other contexts a filter for purifying carbon[5].

One might wonder whether there are contexts at which 'carbon filter' denotes, say, a filter that removes carbon from exhaust gasses. The above CSA does not support this, but the proponent of (C2) has no reason to deny that 'carbon filter' denotes such a device at some contexts. For simplicity, however, I shall suppose that (the proponent of (C2) holds that) 'carbon filter' always either denotes a filter that is made of carbon or a filter for purifying carbon.

Turning to (C3), the claim is that the contents of 'carbon' and 'filter' *do not co-vary with* the content of 'carbon filter'. This is true as long as there is a pair of contexts c_1 and c_2, such that the content of 'carbon filter' varies across c_1 and c_2 while the contents of 'carbon' and 'filter' do not vary between the two contexts. Hence one can make (C3) more precise thus:

(C3#) There is a pair of contexts c_1 and c_2, such that 'carbon' and 'filter' have the same contents at c_1 and c_2, while 'carbon filter' does not have the same content at c_1 and c_2.

[5] I should say that I don't actually want to commit to the claim that the CSA establishes that 'carbon filter' is context-sensitive: *the Attacker* holds that the CSA establishes this, and many theorists would agree with her. (E.g. Austin (1979), Szabó, (2006), and DeRose (2009) accept CSAs as valid.) However, other theorists hold that CSAs are invalid. (E.g. Cappelen and Lepore (2005) reject CSAs, arguing that they over-generate.) *Qua* proponent of Compositionality, the Defender can remain neutral about whether CSAs are valid. However, when (C2) is discussed below, she will grant in the end that the CSA at hand is successful. For simplicity I shall occasionally say that this-and-this CSA establishes that-and-that, to mean that *CSA lovers will claim that* this-and-this CSA establishes that-and-that.

The CSA I presented above can be used to support (C3#): Let us refer to John's and Mary's contexts as c_{pipe} and c_{chem} respectively. In c_{pipe}, 'carbon' intuitively denotes the substance that Vauen's filters are made of; in c_{chem}, 'carbon' intuitively denotes the substance that Jane and Mary want to purify. Hence 'carbon' denotes *carbon* in both contexts. Similarly, 'filter' intuitively denotes filters in both contexts. Thus, both 'carbon' and 'filter' have the same contents at c_{pipe} and c_{chem}, while 'carbon filter' does not have the same content at the two contexts. From this, (C3#) follows by existential generalisation.

Although the Attacker certainly endorses (C3#), there is a caveat regarding this statement: note that it presupposes that 'carbon filter' is context-sensitive, i.e. (C3#) presuppose (part of) the Attacker's second premise. In the end, this won't matter since the Defender will grant (C2). Yet for the purpose of formulating the Compound Argument, it will be better to use a premise that does not presuppose that 'carbon filter' is context-sensitive. The following statement is a premise of this kind.

(C3*) If the compound 'carbon filter' is context-sensitive, then there is a pair of contexts c_1 and c_2, such that 'carbon' and 'filter' have the same contents at c_1 and c_2, while 'carbon filter' does not have the same content at c_1 and c_2.

Since (C3#) implies (C3*), and since the Attacker endorses the former claim, she endorses the latter claim also. The next step is to formulate the Compound Argument.

1.2 The Compound Argument

The aim is to show that (C1), (C2), and (C3*) together imply that Compositionality fails. To show this, note firstly that Compositionality implies the following claim:

VARY. For all complex expressions *e* and for all contexts c_1 and c_2, if *e* does not have the same content at c_1 and c_2, then at least one of *e*'s constituents does not have the same content at c_1 and c_2.

Assume VARY for reductio. (C2) and (C3*) together imply that there is a pair of contexts c_1 and c_2 such that 'carbon filter' does not have the same content at c_1 and c_2, although 'carbon' and 'filter' have the same contents at c_1 and c_2. Together with VARY, this implies that 'carbon filter' has a constituent which does not have the same content at c_1 and c_2. However, 'carbon' and 'filter' are the only two constituents of 'carbon filter', by (C1). This completes the reductio. Thus since Compositionality implies VARY, premises (C1), (C2), and (C3*) together imply that Compositionality fails.

The Compound-Argument is valid: anybody who endorses its premises must reject Compositionality. In what follows, I shall therefore examine each premise in turn and ask whether the Defender can plausibly reject it. Sections two and three consider (C1); although it is not entirely implausible to reject this premise, the Defender will grant it in the end. In section four, I shall present some evidence against (C2), but the evidence turns out to be inconclusive. Sections five and six show that the Defender has good reasons to reject the Attacker's third premise. The conclusion will be that the Defender *can* resist the Compound Argument.

2. The syntactic status of 'carbon filter': word or phrase?

As stated above, (C1) is twofold. Its first part is the claim that 'carbon filter' is a phrase rather than a word. The second part is the claim that 'carbon filter' has exactly two constituents. The current section focusses on the former claim and asks whether the Defender can plausibly reject it; the next section turns to the second part of (C1).

2.1 Testing ‘carbon filter’

To start, there is no consensus in the linguistic literature about the syntactic status of compounds. Some linguists hold that *all* compounds are words[6]. According to an alternative view, at least some compounds are phrases[7]. Yet even if this latter view is correct, it doesn’t follow that ‘carbon filter’ is a phrase. After all, ‘carbon filter’ could be one of the compounds that are words. Thus the onus is on the Attacker to convince the Defender that ‘carbon filter’ is a phrase.

Since the view that all compounds are words is fairly idiosyncratic, the Defender will grant that some, but not all compounds are phrases. Now, proponents of this view hold that there are certain tests with which one can determine whether a particular compound is a word or a phrase. Hence the question is whether the Attacker can find a test that supports her claim that ‘carbon filter’ is a phrase. In what follows, I shall discuss four tests.

2.1.1 The spelling test

The first test to be discussed is the orthographic or spelling test. First, some terminology: if a compound is spelled as one string of symbols, like ‘wheelchair’, it is said to be *closed*; a compound like ‘carbon filter’ which is spelled with a space is said to be *open*; a compound like ‘front-runner’ is said to be *hyphenated*. Given this terminology, the spelling test can be formulated thus:

> SPELL. To determine whether a particular compound is a word or a phrase, consider how the compound is spelled: if the compound is closed, it is a word; if the compound is open or hyphenated, it is a phrase.

[6] E.g. Levi (1978) subscribes to this view.
[7] This view is defended e.g. by Lees (1963), Payne and Huddleston (2002), and Giegerich (2004). To my knowledge, nobody holds that all compounds are phrases.

Since ‘carbon filter’ is spelled with a space between ‘carbon’ and ‘filter’, the test yields that ‘carbon filter’ is a phrase. Hence provided that SPELL is a reliable test, it provides strong support for (C1).

At first glance, SPELL may seem *obviously* correct, as one usually can rely on spelling when trying to tell words apart from phrases. E.g. ‘this house’ or ‘the giraffe’ are clearly phrases; despite its etymology, ‘today’ is syntactically simple, and not a phrase that has ‘to’ and ‘day’ as constituents. Despite this, the Attacker cannot use SPELL to support (C1) as there are several reasons for rejecting the test as unreliable.

First, many compounds receive multiple spellings. For instance, the OED lists each of the following[8]:

breakpoint	break-point	break point
copperplate	copper-plate	
daycare		day care
sandpit	sand-pit	
snakebite	snake-bite	snake bite

Now, it is not very appealing to hold that dictionaries are always mistaken when assigning multiple spellings to a compound. Rather, if dictionaries list both ‘day care’ and ‘daycare’, we should accept that ‘day care’ and ‘daycare’ are equally correct spellings of the compound in question. Yet as the spelling test was formulated, this means that the compound is both a phrase and a word, and that is clearly false[9]. One could reformulate the test so that it counts as inconclusive if a compound receives multiple spellings, but this would considerably limit the test’s usefulness. Thus, that

[8] OED Online, Oxford University Press, September 2014.

[9] One might conclude instead that ‘day care’ and ‘daycare’ are two distinct compounds, and that the former is a word while the latter is a phrase. Yet this conclusion does not seem very plausible. Presumably, the two compounds are synonymous. However, if two expressions are synonyms, homophones, and near-homographs, the most natural conclusion seems to be that the two expressions are in fact one, and that this expression can be spelled in two ways. We draw this conclusion e.g. in cases like ‘et cetera’ and ‘etcetera’.

some compounds receive multiple spellings gives one a reason for thinking that the spelling test is not very reliable[10, 11].

A second problem with the spelling test concerns hyphenated compounds. According to SPELL, these compounds are phrases; yet this is somewhat *ad hoc*. One can see the basic intuition underlying the spelling test: words are 'spelled together' while phrases are spelled with a space. However, it is not clear that the intuition extends to hyphenated compounds. For, hyphenation is a compromise between the open and the closed spelling. On the one hand, in hyphenated compounds, the head and dependent noun[12] are more clearly separated than they are in closed compounds. On the other hand, when a compound is spelled with a hyphen, it *is* spelled as one string of symbols, except that one of the symbols isn't a letter. Further, expressions like 'pre-Christian', 'co-authored', and 'de-escalate' are arguably words, so that it is not true that words are never spelled with a hyphen. Thus it is unclear whether hyphenated compounds should be counted as words or as phrases. Choosing to count them as phrases, as SPELL does, seems arbitrary.

There is a third problem with the spelling test. We can find pairs of very similar compounds, of which one is closed while the other is open. For example, consider 'worktop' and 'work surface'[13]. According to the spelling test, the former compound is a word while the latter is a phrase. Yet it would be at least somewhat odd if this was correct: the two compounds denote very similar objects: indeed, they appear to be synonyms; both appear to be equally common in use, and both have a similar etymology. Therefore, it would be more plausible to hold either that both compounds

[10] If we consider the corpus, an even larger number of compounds receives multiple spellings. E.g. while the OED lists 'lifestyle' with only one spelling (viz. the closed spelling), all three spellings occur in the quotations that the OED provides for the compound: of the seven quotations, two use the closed spelling, four use hyphenation, and one uses the open spelling.

[11] Occasionally, different dictionaries assign different spellings to a given compound. E.g. the OED lists 'stock-room', while both Collins (HarperCollins, 11th edition, 2011) and Chambers (Chambers Harrap, 12th edition, 2011) list 'stockroom'. This gives rise to a similar problem.

[12] In a compound 'N_1N_2', 'N_2' is the head noun while 'N_1' is the determiner noun.

[13] According to all dictionaries that I have consulted, each compound receives only one spelling: 'worktop' is always closed and 'work surface' is always open. Another pair may be 'beefsteak' (always closed) and 'pork steak' (always open).

are words, or that both are phrases. Since the spelling test doesn't yield this, pairs like 'worktop' and 'work surface' provide another reason for thinking that the spelling test is unreliable.

Finally, SPELL gives rise to implausible cross-linguistic conclusions. In English, open and closed compounds appear to be equally common. By contrast, virtually all compounds are closed in German, while all compounds are open in e.g. Vietnamese. Given SPELL, this means that the distribution of verbal and phrasal compounds[14] is very different in the three languages in question. Yet this is not particularly plausible: if only about half the compounds are lexicalised in English, why think that German lexicalises almost all its compounds? Of course, the proponent of the spelling test could claim that SPELL applies only to English. Yet this claim is not plausible either: there is no reason to think that spelling is a reliable guide to syntactic form in English, but not in German or Vietnamese.

In sum, although SPELL yields that 'carbon filter' is a phrase, the Attacker cannot use the spelling test to support (C1) as the test is unreliable.

2.1.2 Modification, Replacement, and Insertion

The remaining three tests can be discussed together as they are all based on the same idea[15]: if a compound is a phrase, it should behave like a phrase. In particular, it should be possible to operate on the compound, and on both its constituent nouns, in the same way in which one can operate on other noun phrases and their constituents. In what follows, I shall consider three tests that follow this line of thought, each based on a different operation. As before, the question will be whether the Attacker can use the tests to support her claim that 'carbon filter' is a phrase.

[14] A verbal compound is a compound that is a word; a phrasal compound is a compound that is a phrase.
[15] The tests are discussed in some more detail e.g. in Giegerich 2004, pp. 7 – 8.

Given a noun phrase like ‘large desk’, one can modify ‘large’ separately, as in ‘very large desk’, and one can insert an expression into the phrase, as in ‘large wooden desk’. Finally, if two noun phrases share their head, as do ‘large desk’ and ‘small desk’, one can conjoin them and replace the head noun’s second occurrence with ‘one’: ‘a large desk and a small one’. Accordingly, the question is whether one can modify ‘carbon’ separately in ‘carbon filter’, whether one can insert a phrase into ‘carbon filter’, and whether one can replace ‘filter’ with ‘one’.

Starting with the modification test, consider the compound ‘steel bridge’ and note that one can modify ‘steel’ separately: it is perfectly possible to read ‘stainless steel bridge’ as denoting a bridge that is made of stainless steel; one needn’t interpret the compound as denoting a steel bridge that is stainless. Hence in ‘stainless steel bridge’, ‘stainless’ can modify ‘steel’ separately. By contrast, it is much more difficult to modify ‘watch’ in ‘watch maker’: e.g. ‘Swiss watch maker’ is most naturally read as denoting a watch maker who is Swiss. It is difficult to read the compound as denoting a maker of Swiss watches. Hence when it comes to modification, ‘steel bridge’ behaves like a phrase, while ‘watch maker’ does not. According to the modification test, this means that the former compound is a phrase while the latter is a word.

To apply the modification test to ‘carbon filter’, consider ‘black carbon filter’. According to my intuitions, this cannot easily be read as denoting a device for filtering black carbon, or as denoting a filter that is made of black carbon. It can only be read as denoting a carbon filter that is black. If this is correct, ‘black’ modifies the whole compound rather than modifying ‘carbon’ separately. Similarly, ‘oxidised carbon filter’ isn’t easily read as denoting a carbon filter that is oxidised. Thus provided that my intuitions regarding ‘black carbon filter’ and ‘oxidised carbon filter’ are veridical,

'carbon filter' fails the modification test[16]. In other words, the test result contradicts the Attacker's claim that 'carbon filter' is a phrase.

The next test is the replacement test. To introduce the test, note that 'steel bridge' and 'iron bridge' share a head noun, as do 'watch maker' and 'shoe maker'. Yet while 'an iron bridge and a steel one' is felicitous, 'a watch maker and a shoe one' is not. According to the replacement test, this means that 'steel bridge' is a phrase while 'shoe maker' is a word.

Whether 'carbon filter' passes the replacement test is not entirely clear: e.g. 'a water filter and a carbon one' is certainly better than 'a watch maker and a shoe one', but seems less felicitous than 'an iron bridge and a steel one'. One might therefore be inclined to say that the replacement test is inconclusive for 'carbon filter'. On the other hand, if the reader feels strongly that 'a water filter and a carbon one' is felicitous, she will conclude that 'carbon filter' passes the test and this conclusion seems feasible to me.

The last test to be discussed is the insertion test. Consider again 'steel bridge' and note that 'steel suspension bridge' is perfectly felicitous, denoting a suspension bridge that is made of steel. Contrast this with e.g. 'milk bottle' and note that constructions like 'milk glass bottle', 'milk two-litre bottle' and 'milk Codd bottle' are infelicitous. According to the insertion test, this means that the former compound is a phrase while the latter is a word.

[16] Note that 'activated carbon filter' is attested, denoting a filter that is made of activated carbon. However, I don't think that this shows that 'carbon filter' passes the modification test: if 'carbon filter' really is a phrase, one would expect that it *always* behaves like a phrase, so that 'carbon' is modified separately in 'black carbon filter' and 'oxidised carbon filter'. (By comparison, any word that can modify 'large' can modify it separately in 'large desk': very large desk, rather large desk, incredibly large desk, surprisingly large desk …) Hence as long as one agrees that 'carbon' isn't modified separately in 'black carbon filter' and 'oxidised carbon filter', the most one can conclude from the existence of 'activated carbon filter' is that the modification test is inconclusive for 'carbon filter'. Since the Defender will grant in the end that 'carbon filter' is a phrase, this conclusion won't affect the overall debate.

As far as I can see, ‘carbon filter’ passes the insertion test, which supports the Attacker’s claim that the compound is a phrase: While ‘carbon airlift filter’[17] may not be as easy to understand as ‘steel suspension bridge’, it is clearly more felicitous than e.g. ‘milk two-litre bottle’. It can be read as denoting an airlift filter that is made of carbon, and arguably also as denoting a device that uses an airlift to filter carbon.

To summarise, the insertion test, and arguably also the replacement test, support the Attacker’s claim that ‘carbon filter’ is a phrase, which puts at least some pressure on the Defender to grant the first part of (C1). Of course, the Defender could insist that the modification test – which contradicts the Attacker’s claim – is more pertinent than the other two tests; or she could side after all with those linguists who hold that all compounds are words. I won’t pursue these lines of thought, however, but shall grant, on the Defender’s behalf, that ‘carbon filter’ is a phrase.

2.2 Creating phrases

Before moving on, I’d like to examine a more direct argument for the claim that ‘carbon filter’ is a phrase. While the argument is initially quite convincing, I think that it fails in the end, or anyway that the Attacker would be ill-advised to rely on it.

To start, assume again that some, but not all compounds are phrases. If some compounds are phrases, then some English phrases are the result of conjoining two nouns. Hence that some compounds are phrases implies that it is possible in English to create a grammatically correct phrase by conjoining two nouns.

Next, suppose that e.g. ‘glass bottle’ is a *word*. Clearly, ‘glass’ and ‘bottle’ are nouns. Therefore, since it is possible in English to create a phrase by conjoining two nouns, the following is a phrase: ‘glass bottle’. Thus there are two compounds spelled

[17] An airlift filter is a filter that is powered by compressed air.

'glass bottle' in English, the original verbal compound and the newly introduced phrasal compound. As the example was chosen arbitrarily, the view that some compounds are phrases hence arguably implies the following principle.

CREATE. For every English verbal compound k_1, there is a phrasal compound k_2 in English that is homonymous with k_1.

Note that CREATE guarantees that there is a phrasal compound spelled 'carbon filter' in English[18]. Yet note also that CREATE does *not* imply that the phrasal compound is the *only* compound spelled 'carbon filter' in English: the principle leaves open the possibility that the verbal compound 'carbon filter' exists alongside the phrasal compound. In other words, with CREATE in place 'carbon filter' is potentially *ambiguous* between the phrasal and the verbal compound. With this in mind, let's ask whether CREATE supports the Attacker's position.

To start, CREATE does *not* imply the first part of (C1). For, the first part of (C1) states that *the* compound 'carbon filter' is a complex expression, and hence presupposes that there is just one compound spelled 'carbon filter' in English. Yet as just noted, CREATE leaves open that there are two compounds spelled 'carbon filter' in English, and hence does not imply the first part of (C1). Still, one might think that the Attacker can just set the verbal compound aside and focus on the phrasal compound when running the Compound Argument. In particular, she could just reformulate (C1) – and (C2) and (C3) – to get rid of the now problematic presupposition. Consider[19]:

(C1a) The *phrasal* compound 'carbon filter' has exactly two constituents, 'carbon' and 'filter'.

[18] An expression spelled 'carbon filter' clearly exists in English; it is either a phrase or a word. If it's a phrase, English contains a phrasal compound spelled 'carbon filter'. If it's a word, CREATE implies that a second, phrasal compound spelled 'carbon filter' exists alongside the verbal compound. Hence English contains a phrasal compound spelled 'carbon filter' in both cases.

[19] I focus on the original (C3), but (C3#) and (C3*) could obviously be reformulated as well.

(C2a) The *phrasal* compound 'carbon filter' is context-sensitive, denoting in some contexts a filter that is made of carbon, while denoting in other contexts a filter for purifying carbon.

(C3a) The contents of 'carbon' and 'filter' do not co-vary with the content of the *phrasal* compound 'carbon filter'.

The new premises no longer make the problematic presupposition. They do presuppose that English contains a phrasal compound spelled 'carbon filter'; yet with CREATE in place, this presupposition is unproblematic. Of course, CREATE does nothing to establish the reformulated premises. Let's focus on (C2a) and ask how the Attacker might support it.

The Attacker supported the original (C2) with a CSA. One might therefore think that she can support (C2a) with a CSA also. In outline, this CSA would run as follows[20]:

> In c_{pipe}, John truthfully utters 'Vauen produces carbon filters'. In c_{chem}, Mary truthfully utters 'Vauen does not produce carbon filters'. If the uttered sentences were context-insensitive, the utterances couldn't both be true. Since both utterances are true, the sentences are context-sensitive. Since 'Vauen' and 'produces' have the same contents in c_{pipe} and c_{chem}, John's utterance of 'carbon filter' does not have the same content as Mary's utterance of 'carbon filter'.

The conclusion is supposed to be (C2a), the claim that the phrasal compound 'carbon filter' is context-sensitive. However as 'carbon filter' is potentially ambiguous in the current framework, the data does not support this conclusion: given CREATE, it stands to reason that only one of the two speakers uttered the phrasal compound 'carbon filter' while the other speaker uttered the verbal compound. If this is the case, perhaps John's utterance of 'carbon filter' has a different content than Mary's utterance of 'carbon

[20] Recall, c_{pipe} is a context in which John and Henry discuss pipe smoking; c_{chem} is a context in which Jane and Mary consider how they might purify their carbon sample.

filter', simply because John and Mary uttered different expressions, and not because they uttered a context-sensitive expression. In short, in the current framework, one needn't postulate any context-sensitivity to explain the data. Thus with CREATE in place, the Attacker's CSA becomes invalid.

To summarise, CREATE does the Attacker more harm than good: it undermines her second premise, unless she can establish that both John and Mary uttered the phrasal compound. Since this will be rather difficult to establish in the current framework, the Attacker seems better off rejecting CREATE and assuming that there is just one compound 'carbon filter' in English. Yet as we saw in the previous sub-section, the Attacker can argue that this compound is a phrase. The current result is hence not all too devastating.

3. The syntactic status of 'carbon filter': constituents

The second part of premise (C1) is the claim that 'carbon filter' has exactly two constituents, viz. the words 'carbon' and 'filter'. In this section, I shall discuss whether the Defender can plausibly reject this claim.

3.1 Hidden variables

To start, consider sentence (2).

(2) It rains.

This sentence can be used to communicate that it rains in Oxford, that it rains in Bonn, that it rains in Saigon, and so on. It therefore stands to reason that (2) is context-sensitive. One way to account for the context-sensitivity of (2) is by postulating a hidden location-variable: Assume that a variable is present in the syntactic structure of

(2), that this variable is not realised in the sentence's surface form, and that the variable ranges over locations. Then the syntactic structure of (2) is (2a)[21, 22], and the context-sensitivity of (2) can be explained by saying that *l* has different contents at different contexts.

(2a) It rains in *l*.

Now consider the following principle.

HIDE. A variable *r* is present in the syntactic structure of 'carbon filter'; *r* is not realised in the compound's surface form and ranges over relations.

If HIDE is true, the context-sensitivity of 'carbon filter' can be explained: at some contexts, the *made-from*-relation is the content of *r*; at other contexts, the *used-for*-relation is the content of *r*. At all and only the former contexts, 'carbon filter' denotes a filter that is made of carbon; it denotes a device for filtering carbon at all and only the latter contexts.

Note that HIDE allows reconciling the context-sensitivity of 'carbon filter' with VARY: by VARY, at least one constituent of 'carbon filter' has different contents in c_1 and c_2 if 'carbon filter' has different contents in c_1 and c_2. Given the current assumptions, this is correct: if 'carbon filter' denotes a filter that is made of carbon in c_1, while denoting a device for filtering carbon in c_2, the hidden variable in 'carbon filter' has the *made-of*-relation as value in c_1 while having the *used-for*-relation as value in c_2. Hence HIDE allows the Defender to explain the context-sensitivity of 'carbon filter' while maintaining VARY and Compositionality.

[21] For simplicity, I'm ignoring the possibility that (2) contains a second hidden variable ranging over times.

[22] The claim that (2) contains a hidden variable isn't tantamount to the claim that (2a) is the syntactic structure of (2). Written in bracket notation, the structure of (2a) is: [$_{NP}$ It [$_{VP}$ rains [$_{PP}$ in *l*]]]. Thus the variable appears within a PP and is a terminal node of (2a). However, there is an alternative view, on which the variable in (2) merely *co-habits* a terminal node with 'rains'. On this view, the structure of (2) is: [$_{NP}$ It [$_{VP}$ <rains, *l*>]]. See Stanley 2002 for further details. For present purposes, I'd like to set the latter view aside, however: if (2) contains a variable at all, (2a) is its syntactic structure.

The dialectic situation is now as follows. HIDE implies that 'carbon filter' has more than two constituents. Hence if the Defender endorses HIDE, she thereby rejects (the second half of) premise (C1). In that sense, HIDE constitutes a reply to the Compound Argument. Whether HIDE constitutes a *convincing* reply depends, of course, on how plausible HIDE really is. To assess this, the next sub-section considers whether there is any evidence to support HIDE, while the following sub-section discusses an argument against HIDE.

3.2 Binding

HIDE is an assumption about the syntactic structure of 'carbon filter', and one should not make assumptions about syntactic structure in the absence of syntactic evidence. In other words, if the Defender cannot support HIDE with syntactic evidence, it is not very plausible for her to endorse HIDE. In what follows, I shall consider one particular kind of evidence – evidence from binding – which is often used to support claims about hidden variables. It turns out that this type of evidence does *not* support HIDE.

To start, consider the following sentence, in which (2) is embedded under a quantifier.

(3) Everywhere John goes, it rains.

Those who claim that a hidden location-variable is present in (2) think that (3) supports their claim, for the following reason[23]. On its most natural reading, (3) expresses that every location *l* that John goes to is such that it rains **in *l***. This means that the quantifier phrase 'everywhere John goes' is not a *vacuous* quantifier in (3); rather, it *binds a variable*. More particularly, it binds a variable in the sentence that is embedded under it. Since (2) is this sentence, it follows that 'everywhere John goes' binds a

[23] See e.g. Stanley 2002.

variable in (2). And as a quantifier cannot bind a variable where none is present, it follows *a fortiori* that a variable is present in (2). Call this the *Binding Argument*.

To illustrate the Binding Argument a little further, note that it has also been used to argue that nominal and relational expressions are associated with a hidden variable[24]. To outline the Binding Argument as applied to these expressions, consider the (a)-sentences in (4) and (5); they are embedded under a quantifier in the (b)-sentences, and the quantifiers appear to bind a variable in the embedded sentences[25]. The *that*-clauses in (c) indicate the (b)-sentences' most natural readings.

(4a) John failed every student.

(4b) In most of his classes, John failed every student.

(4c) that in most of his classes *x*, John failed every student **in *x***.

(5a) Goods that are imported from China don't sell well.

(5b) In some countries, goods that are imported from China don't sell well.

(5c) that in some countries *x*, the goods that are imported **to *x*** from China don't sell well.

For present purposes, we can remain neutral on whether the Binding Argument is persuasive[26]. The point is only that many would consider the Binding Argument to be the best kind of evidence for a hidden variable in (2), (4a), and (5a). Hence, it is worth

[24] See Stanley 2000.

[25] Stanley's view is slightly more complex than the view I am describing: according to him, (4a) and (5a) contain both a function variable and an individual variable. For present purposes, however, it will suffice to consider the simpler view.

[26] Cappelen and Lepore (2002, 2005) object that the Binding Argument over-generates: if (3) is evidence that (2) contains a hidden variable, then – they argue – a sentence like 'Everywhere John goes, 2 + 2 = 4' is evidence that '2 + 2 = 4' contains a hidden variable; and they take it to be obviously false that '2 +2 = 4' contains a hidden variable. See e.g. Cohen and Rickless 2009 for a critical discussion of Cappelen and Lepore's objection. Another objection challenges the Binding Argument's assumption that binding is a syntactic phenomenon. See e.g. Farkas 1997. Finally, the Binding Argument assumes that if sentence σ contains a variable when embedded under a quantifier, then σ already contains a variable in its *unembedded* form. Recanati (2002, p. 329) objects to this.

asking whether one can use the Binding Argument to show that a hidden variable is present in 'carbon filter'.

As far as I can see, one cannot run the Binding Argument for 'carbon filter'. The reason is simply that there is no expression in English that could bind the variable that is assumed to be present in 'carbon filter'. For, that variable is supposed to range over *relations*, while ordinary English quantifier expressions like 'every', 'some', 'most', 'whenever', etc. bind variables that range over *individuals*. Thus although the quantifier in e.g. (6) apparently binds something, it doesn't bind the right kind of variable.

(6) Whenever John is in London, Vauen produces carbon filters.

One can embed 'Vauen produces carbon filters' under a quantifier that explicitly ranges over relations, as in (7). Yet it is far from obvious that such a quantifier binds anything.

(7) ?There is a relation *r*, such that Vauen produces carbon filters.

Stanley claims that some English sentences contain a hidden variable that ranges over functions[27]. To support his claim with the Binding Argument, he argues that (8b) is the best interpretation of the question in (8a); he concludes from this that 'what' in (8a) binds a variable that ranges over functions.

(8a) Q: What does every author like?

A: Her first book.

(8b) What function *f* is such that for every author *x*, *x* likes *f*(*x*)?

Supposing that Stanley is correct, and seeing that functions are a special sort of relation, one might think that a question like the one in (8a) can also be used to show

[27] See Stanley 2002.

that some sentences involve quantification over relations. The question in (9a) may be a sentence of this sort[28].

(9a) Q: What does arsenic do to rats?

A: It kills them.

(9b) What relation R is such that arsenic stands in R to rats?

Now, if the question in (9a) indeed involves quantification over a relation variable, perhaps one can also find a question that involves quantification over the relation variable that is assumed to be present in 'carbon filter'. However, this doesn't seem possible: what is needed is a sentence that can be interpreted according to (10). While (11a) can perhaps be so interpreted, it does not contain 'carbon filter'. (11b) contains the compound, but cannot be interpreted according to (10). Finally, (11c) is ungrammatical.

(10) What relation R is such that (some) filters stand in R to carbon?

(11a) What do filters have to do with carbon?

(11b) What do carbon filters have to do with carbon?

(11c) #What do carbon filters?

In sum, it seems impossible to bind the variable that is assumed to be present in 'carbon filter'. Consequently, one cannot run the Binding Argument for 'carbon filter'. Of course, this does not show that there is *no* syntactic evidence for HIDE: the Defender may be able to find syntactic evidence of another sort. However, I don't know of any such evidence. Thus for all that's been said, the Defender has no syntactic evidence to support HIDE. It would therefore be rather contentious if she claimed that a hidden

[28] I don't mean to commit to the claim that (9a) involves quantification over a relation variable. *If* one thinks that Stanley is correct in saying that 'what' in (8a) binds a function variable, and *if* one thinks that 'what' can also bind relation variables, *then* the question in (9a) is a good candidate for being a sentence that involves quantification over a relation variable.

relation variable is present in 'carbon filter'. The next sub-section considers an argument for the same conclusion.

3.3 Syntactic siblings

While the previous sub-section argued that the Defender doesn't have any evidence *for* HIDE, the current section presents an objection *against* HIDE. In outline, the objection is as follows:

> If 'It rains in *l*' is the structure of (2), then (2) and (2*) = 'It rains in Bonn' have very similar syntactic structures. An important difference is that the PP in (2*) is overt while the PP in (2) is hidden or covert. If (1a) = 'Vauen produces carbon filters' contains a hidden variable, we'd expect to find a sentence that is to (1a) as (2*) is to (2), i.e. a sentence which is structurally similar to (1a), but has an overt constituent where (1a) has the covert variable. Yet no such sentence exists. Thus, (1a) does not contain a hidden variable.

After spelling out the objection in more detail, I shall argue that it is not successful in the end: it is not clear that there is no sentence that is to (1a) as (2*) is to (2).

Assume that (2a) is the syntactic structure of (2). Under that assumption, (2) and (2*) have very similar syntactic structures: The sentences' NPs consist of just one noun; their VPs have exactly two constituents, viz. a verb and a PP; and their PPs consist of just a preposition and a noun. In short, (2) and (2*) are identical, except that (2*) contains 'in Bonn' where (2) contains 'in *x*'. Putting it in terms of syntactic trees, (2) and (2*) differ *only with respect to their terminal nodes*. For ease of reference, let us say that two

sentences are *syntactic siblings* to mean that they differ only with regard to (some of) their terminal nodes. Then all the sentences in (12) are syntactic siblings of (2)[29].

(12a) It rains near London.

(12b) It snows.

(12c) It snows in Munich.

(12d) It hails.

(12e) It hails outside Glasgow.

Like (2), (12b) and (12d) contain a covert element. By contrast, (12a), (12c), and (12e) are like (2*) in that all of their constituents are overt. Let us therefore say that the latter four sentences are *fully overt* syntactic siblings of (2).

Now, it is not at all surprising that (2) has syntactic siblings: after all, finding a sibling is just a matter of replacing like-for-like, a matter of substituting syntactically equivalent terminal nodes. Hence since 'rains' isn't the only verb that appears in 'it'-constructions, the result of replacing 'rains' with a verb like 'snows' is a sibling of (2) – just as one would expect. Likewise, it is not surprising that (2) has fully overt syntactic siblings: as (2) contains the PP 'in *x*' as constituent, all its syntactic siblings will also contain a PP. Yet there is no reason to expect that all siblings of (2) will have their PP hidden like (2) does. On the contrary: since PPs are normally overt, one would expect to find a syntactic sibling of (2) the PP of which is overt. As (2*), (12a), (12c), and (12e) show, this is exactly what one does find.

To illustrate the notion of a fully overt syntactic sibling further, return to (4a) = 'John failed every student' and (5a) = 'Goods that are imported from China don't sell well',

[29] If (2a) is the syntactic structure of (2), then it is eminently plausible that 'It snows in *l*' and 'It hails in *l*' are the syntactic structures of 'It snows' and 'It hails'.

and assume that these sentences contain hidden variables[30]. For the above-stated reasons, one would expect to find fully overt syntactic siblings of (4a) and (5a). This expectation is met in (4a*) and (5a*).

(4a*) John failed every student in his logic class.

(5a*) Goods that are imported to Germany from China don't sell well.

In sum, we saw that (2), (4a), and (5a) have fully overt syntactic siblings. Importantly, we also saw that it is not at all surprising that they have these siblings. With this in mind, return to (1a) = 'Vauen produces carbon filters'.

Assume with HIDE that (1a) contains a hidden variable, just like (2), (4a), and (5a). Then we'd expect that one can find a fully overt syntactic sibling of (1a), just like one could find fully overt syntactic siblings for (2), (4a), and (5a). For, whatever the syntactic structure of (1a) is exactly, one would expect that there are other sentences with the same structure – modulo the terminal nodes; and one has no reason to expect that all these sentences keep one of their constituents covert. To put it the other way around, if (1a) doesn't have a fully overt syntactic sibling, then this is a reason to reject HIDE.

Now, HIDE implies that the variable in (1a) ranges over relations. Therefore, every sibling of (1a) contains an expression that denotes a relation. Since verbs often express relations, as in 'These filters *purify* carbon', it stands to reason that one can create a sibling of (1a) by inserting a verb into (1a). However, there is no grammatically correct way to do this:

(i) #Vauen produces carbon purify filters.

(ii) #Vauen produces purify carbon filters.

[30] More particularly, assume that 'John failed every student in *x*' and 'Goods that are imported to *x* from China don't sell well' are the syntactic structures of (4a) and (5a).

(iii) #Vauen produces carbon filters purify.

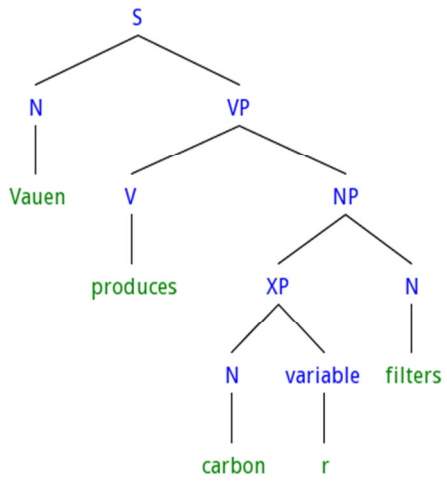
The following sentence is grammatical, but it is not syntactic siblings of (1a).

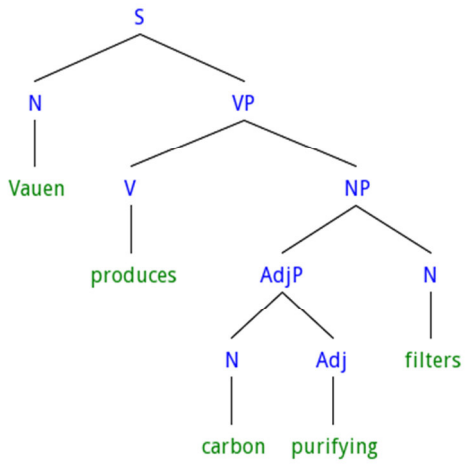
(iv) Vauen produces filters that purify carbon.

So far, then, it looks like (1a) doesn't have a fully overt syntactic sibling, which would indicate that (1a) does not contain a hidden variable. However, there is one more sentence to consider:

(v) Vauen produces carbon-purifying filters.

Prima facie, (v) may not look like a syntactic sibling of (1a) as it does not contain 'carbon filters'. However, if HIDE is true, (v) and (1a) resemble each other more closely than one might think. Let's assume that 'carbon-purifying filter' has three constituents, so that its tree has three terminal nodes. Then if HIDE is true, 'carbon filter' and 'carbon-purifying filter' are syntactically very similar: for, if HIDE is true, the tree for 'carbon filter' *also* has three terminal nodes, viz. 'carbon', 'filter', and the variable. This being so, it does look like (v) and (1a) differ only with respect to their terminal nodes. The following diagrams may help to bring this out more clearly – where '*r*' is the variable that is (assumed to be) present in 'carbon filter'.





The trees suggest that (1a) and (v) are syntactic siblings. However, one might think that the tree to the right isn't the tree for (v): above I *assumed* that 'carbon-purifying filter' has three constituents, in which case the tree for (v) will indeed have five terminal nodes, as shown. However, it might turn out that 'carbon-purifying' is not itself a phrase but a word, in which case 'carbon-purifying filter' has only two constituents. The tree for (v) will then have only four terminal nodes. Since the tree for (1b) has five terminal nodes if HIDE is true, this would mean that (1a) and (v) aren't siblings after all[31].

Even if the trees are correct, one might worry that the NP 'carbon filter' is too dissimilar from the NP 'carbon-purifying filter' for (1a) and (v) to be siblings. As the notion of syntactic sibling was defined, that two sentences have the same number of terminal nodes doesn't make them siblings. The sentences must furthermore agree in all their non-terminal nodes: e.g. if a sentence contains a PP, all its siblings must also contain a PP, in the same place. However, it is not clear that (1a) and (v) agree in this regard: the tree for (v) contains an AdjP, but it is far from clear that the phrase labelled 'XP' is also an AdjP: the phrase does not contain an adjective but only a noun and a variable, and the variable appears to be the head. Thus, it is far from clear that (1a) and (v) are syntactic siblings.

There is one other point to take into consideration, however. Although the notion of syntactic sibling is an interesting notion, it isn't the only relevant notion of syntactic resemblance. E.g. 'John kissed Hanna' and 'The boy kissed the girl' aren't siblings, but still resemble each other syntactically. So perhaps the question of whether (1a) has a fully overt syntactic sibling isn't as relevant as the argument made it out to be: even if (1a) and (v) aren't syntactic siblings, they still resemble each other quite closely in their syntactic structures. And perhaps that's enough to show that HIDE is plausible.

[31] Bell (2005) argues that 'carbon filter' isn't an NP but an AdjP, which would mean that the tree to the left is also incorrect.

In sum, the proposed argument against HIDE is inconclusive: On the one hand, it isn't clear that (v) and (1a) are syntactic siblings. On the other hand, even if they are not, this may prove less than the argument's proponent claimed.

To conclude the discussion of HIDE, the Defender does not have any evidence to support it, but the claim has also not been disproved. Therefore, the Defender could decide to endorse HIDE as a reply to the Compound Argument. However, as I will argue in section six, the Defender can formulate a much more convincing reply by rejecting premise (C3#). Therefore, I shall assume that the Defender eschews HIDE and grants (the second half of) premise (C1).

4. The context-sensitivity of 'carbon filter'

In this section, I shall examine premise (C2), the claim that 'carbon filter' is context-sensitive. For this purpose, I shall present two tests, SPEECH and MENTION, devised by Cappelen and Lepore[32]. While the tests indicate that 'carbon filter' is not context-sensitive, I shall argue that the Defender should not rely on the tests to reply to the Compound Argument.

As many linguistic tests do, SPEECH and MENTION rely on the having of certain linguistic intuitions. In discussing the tests, I shall report Cappelen and Lepore's intuitions as well as my own. Of course, it is possible that the reader – and the Attacker – have different intuitions. Yet as I will advise the Defender in the end not to rely on the test anyway, a difference in linguistic intuitions won't affect the overall dialectic situation too much.

[32] See Cappelen and Lepore 2005, cap. 7. They do not use the names 'SPEECH' and 'MENTION'.

4.1 Speech reports

The first test centres on speech reports, and can be formulated thus:

SPEECH. To test whether expression *e* is context-sensitive, pick two contexts c_S and c_R. Further, consider a sentence σ that contains *e*, and suppose that a speaker S utters σ in c_S. If *e* is context-sensitive, 'S said that σ' will be false in c_R.

To explain SPEECH, suppose one wants to test expression *e*. The first step is to choose c_S and c_R – 'S' for 'speech' and 'R' for 'report'. Starting with the latter context, Cappelen and Lepore insist that theorists use their *own* context as c_R, rather than some made-up context in which a fictional speaker makes a speech report. This has the advantage that the theorist has "to confront intuitions about her own language *in use*, and not just about other people's use of language (or her own use [of the sentence] in other contexts"[33]. According to Cappelen and Lepore, this guarantees a more reliable verdict on 'S said that σ'.

Turning to the choice of c_S, Cappelen and Lepore require that this context must be *relevantly different* from c_R: when testing an expression *e* with SPEECH, one considers what contextual feature the meaning of *e* is likely to be sensitive to; then c_S is chosen so that it differs from c_R as much as possible with regard to *those* features. For instance, suppose one wanted to test 'today' with SPEECH. Since the meaning of 'today' is sensitive to time, c_S must be chosen so that c_S and c_R take place on different days. Similarly, when testing 'know' with SPEECH, c_S must be chosen so that two different epistemic standards[34] are operative in c_S and c_R.

To illustrate SPEECH, let's apply the test to 'today', which everybody agrees to be context-sensitive. Since I am writing this on 1 March, c_R will be a context that takes

[33] Cappelen and Lepore 2005, p. 104, emphasis in the original.
[34] Very roughly, an epistemic standard is a criterion by which to judge claims of the forms of 'S knows that *p*'. See e.g. DeRose 2009 for further details.

place on 1 March. Accordingly, c_S will be a context that takes place on a different day – 1 February, say – to ensure that c_S and c_R are relevantly different. To run the test, suppose that John utters 'Mike kissed Jane today' in c_S, and that I report his speech in c_R with (SP_1). In other words, John speaks on 1 February and I report his speech on 1 March.

(SP_1) John said that Mike kissed Jane today.

Clearly, my utterance of (SP_1) is false: it expresses that John said that Mike kissed Jane on 1 March; yet John said that Mike kissed Jane on 1 February. Thus SPEECH correctly yields that 'today' is context-sensitive.

Next consider 'ready', which some authors believe to be context-sensitive. If 'ready' is context-sensitive, its content at a context *c* may depend on the plans and expectations that are salient in *c*: if we plan to go out, 'ready' is likely to express the property of being ready to go out; if we expect to sit an exam, 'ready' is likely to express the property of being ready to take the exam. As I am not expecting to go out right now, but planning to submit this paper for examination, my current context, and hence c_R, is a context in which the plan to submit a paper is salient. Accordingly, c_S will be a context in which different plans and expectations are salient.

To run the test, suppose that John and Anna want to go out with Mike; they are waiting for him as he is getting dressed. In that context, John says 'Mike is ready'. The question is whether I can truthfully report John's speech in my current context with (SP_2).

(SP_2) John said that Mike is ready.

At least according to Cappelen and Lepore's intuitions[35], I *can* truthfully report John's speech by uttering (SP_2) in c_R. They conclude that 'ready' fails the SPEECH test. My own intuitions coincide with Cappelen and Lepore's.

Suppose Cappelen and Lepore are correct in saying that my speech report is true. Perhaps someone will try to resist the conclusion that 'ready' is context-insensitive as follows: since I am not planning to go out right now, it *seems* that the plan to go out is not salient in my current context. However, the objection continues, this plan *becomes* salient when I report John's speech: by reporting John's speech, I move into a context in which the plan to go out is salient, and hence into a context in which 'ready' means what it meant in John's context. In short, the proposal is that my context is *automatically adjusted* when I produce the speech report[36]. Thus that my speech report is true doesn't mean that 'ready' is context-insensitive. It only means that contexts of speech report are always adjusted to the context of the original speech.

The proposed reply fails. Suppose that Jim and Tara speak about an upcoming exam, wondering whether Mike needs to study more before he can sit the exam. In other words, in Jim and Tara's context, the plan to sit an exam is salient. In that context, Tara says: 'Mike is ready'. Now return to my current context and suppose I utter 'Both John and Tara said that Mike is ready.' According to Cappelen and Lepore, this joint speech report is fully accurate and true. Yet the suggested process of automatic context-adjustment cannot explain this: if 'ready' is context-sensitive, then the content it has in Jim and Tara's context is *different* from the content it has in John and Anna's context. Therefore, 'ready' in my mouth cannot have the content it has in John's mouth *and* the content it has in Tara's mouth. Thus, context-adjustment cannot explain Cappelen and Lepore's data.

[35] See Cappelen and Lepore 2005, p. 90 sq. Of course, if Cappelen and Lepore were to test 'ready', they would use as c_R whatever context they are in when running the test, instead of my current context.
[36] See DeRose 2009, cap 4, for a discussion of context adjustment.

Return now to 'carbon filter'. It is not easy to say what features of context the compound's meaning is sensitive to (if it is context-sensitive at all). Therefore, it is not easy to select c_S. Of course, the onus is on the Attacker to specify the relevant features: if she can't specify any such features, her claim that 'carbon filter' is context-sensitive becomes much less plausible. For argument's sake, let us assume that the Attacker holds that the meaning of 'carbon filter' is sensitive to the speaker's interests: if the speaker is interested in carbon processing, 'carbon filter' denotes in her context a device for purifying carbon; it denotes a filter that is made of carbon if the speaker is interested in how carbon can be used. This assumption probably oversimplifies the matter, but it will help to select c_S and c_R.

To run the test, note that I am no engineer, and that I have no interest in processing carbon. I do own a pipe, however, and know that some pipe filters are made of carbon. Therefore, I am at least a little bit interested in how carbon can be used, and this interest is reflected in c_R. To construct c_S, suppose that Jake and Selma are engineers and interested in purifying a carbon sample. Suppose further that Jake utters (1a) = 'Vauen produces carbon filters' in c_S; I report his speech in c_R with an utterance of (SP_3). As before, the question is whether my speech report is accurate and true.

(SP_3) Jake said that Vauen produces carbon filters.

According to my own intuitions, my speech report *is* true. And although Cappelen and Lepore do not consider 'carbon filter', I believe that they would share my intuition[37]. I thus conclude that 'carbon filter' fails the SPEECH test. Before discussing this conclusion further, let me present Cappelen and Lepore's second test.

[37] Cappelen and Lepore hold that only indexicals and demonstratives, and possibly expressions like 'local' and 'foreign', pass the SPEECH test. See Cappelen and Lepore 2005, p. 108.

4.2 Use and Mention

The second test relies on the distinction between use and mention and can be formulated as follows.

MENTION. To test whether expression *e* is context-sensitive, pick a sentence σ which contains *e*. If *e* is context-sensitive, there is a true utterance of the following: Although not-σ, there are true utterances of 'σ'.

To illustrate MENTION, consider again 'today' and note that I am writing this paragraph on a Tuesday. Now I'd like to make the following statement.

(SP_4) Although it isn't Monday today, there are true utterances of 'It is Monday today'.

Clearly, my statement is true: since I am writing on a Tuesday, it is true that it isn't Monday today; at the same time, it is clear that there are true utterances of 'It is Monday today'. Hence, there are true utterances of (SP_4), and 'today' passes MENTION. Thus the test correctly yields that 'today' is context-sensitive.

By contrast, consider 'ready'. Let *t* be the time at which I am writing this paragraph, and ask whether there are any true utterances of the following sentence.

(SP_5) Although Kate isn't ready at *t*, there are true utterances of 'Kate is ready at *t*'.

Cappelen and Lepore deny that there are any true utterances of (SP_5): for, if there are true utterances of 'Kate is ready at *t*', Kate had better be ready at *t*[38]. I am inclined to agree with Cappelen and Lepore, and thus conclude that 'ready' fails the MENTION test.

Finally, let's apply MENTION to 'carbon filter', by considering whether there are any true utterances of the following sentence.

(SP_6) Although carbon filters aren't made of metal, there are true utterances of 'Carbon filters are made of metal'.

[38] See Cappelen and Lepore 2005, p. 109.

Even supposing that filters for purifying carbon are made of metal, I am inclined to deny that there are any true utterances of (SP_6). As before, it is safe to assume that Cappelen and Lepore would agree. Thus, 'carbon filter' fails the MENTION test.

4.3 Assessing the tests

Cappelen and Lepore claim that it is "utterly obvious"[39] that my utterance of (SP_3) is true, and equally obvious that there are no true utterances of (SP_6). If that is correct, then 'carbon filter' is context-insensitive. However, one might not find Cappelen and Lepore's claims as obvious: perhaps (SP_3) seems true only because I have not described the context in which I uttered (SP_3) in sufficient detail; perhaps with enough ingenuity, one can find a context in which (SP_6) can be truthfully uttered[40].

There is a second, bigger problem. The Defender is currently replying to (C2), the claim that 'carbon filter' is context-sensitive. Hence for the Defender's purposes, it is enough if SPEECH and MENTION make it plausible that *one particular expression* – 'carbon filter' – is context-insensitive. However, Cappelen and Lepore designed the tests to establish a much more general claim, viz. Semantic Minimalism – the claim that the set of context-sensitive English expression is limited, roughly, to indexicals, demonstratives, and pronouns.

Now, although a more detailed discussion may be necessary, the following conjecture seems plausible: either the tests work the way Cappelen and Lepore intend them to work and can be used to establish their general claim; or the tests do not work at all and cannot be used to establish anything.

If the said conjecture is correct, the dialectic situation is as follows. If SPEECH and MENTION can be used to establish that 'carbon filter' is context-insensitive, then the

[39] Cappelen and Lepore 2005, p. 91.
[40] For a more detailed criticism of Cappelen and Lepore's tests, see e.g. Hawthorne 2006 and Szabó 2006.

tests also establish Semantic Minimalism. Thus if the Defender endorses that the tests can be used to reply to (C2), she must also endorse Semantic Minimalism. To be sure, Semantic Minimalism may be true; yet the Defender may not want to be committed to this position. If she does not want to be thus committed, and if the above conjecture is true, she cannot use the tests to reply to (C2).

In sum, the results of both tests are inconclusive. On the one hand, both tests indicate that (C2) is false and hence challenge the Attacker's argument. On the other, it is not clear that the Defender *wants* to pose this challenge: for her purposes, it seems best to remain neutral about Semantic Minimalism. Thus, the Defender grants (C2).

5. The meanings of 'carbon' and 'filter': preliminaries

This and the next section discuss the Attacker's third premise. Recall, there are two versions of this premise: while (C3#) presupposes that 'carbon filter' is context-sensitive, the conditional (C3*) does not carry this presupposition. When formulating the Compound Argument, it was best to work with (C3*). Yet as the Defender now grants that 'carbon filter' is context-sensitive, we can focus on (C3#). To repeat it:

(C3#) There is a pair of contexts c_1 and c_2, such that 'carbon' and 'filter' have the same contents at c_1 and c_2, while 'carbon filter' does not have the same content at c_1 and c_2.

When combined with the claim that 'carbon' and 'filter' are context-*in*sensitive, (C2) implies (C3#)[41]. Hence if she can, the Attacker will maintain that 'carbon' and 'filter' *are* context-insensitive. Yet in this section, I shall show that she *cannot* plausibly maintain this: given her other commitments, the Attacker must grant that 'carbon' and

[41] If 'carbon filter' is context-sensitive, as (C2) has it, there is a pair of contexts c_1 and c_2, such that 'carbon filter' has different contents at c_1 and c_2. Thus if 'carbon' and 'filter' are context-insensitive, they have the same contents at c_1 and c_2, while 'carbon filter' does not have the same content at c_1 and c_2.

'filter' are context-sensitive. Hence in this section, the dialectic situation is reversed: the aim isn't to establish that the *Defender* can plausibly *reject* that 'carbon' and 'filter' are context-insensitive, but rather that the *Attacker* cannot plausibly *endorse* this claim. Once this has been established, I will turn to the Defender's reply to (C3#).

5.1 Two CSAs for 'carbon' and 'filter'

Consider the following two CSAs, which respectively attempt to establish that 'carbon' and 'filter' are context-sensitive[42].

ARGUMENT 1. Suppose that two engineers want to conduct a highly sophisticated experiment, for which they require a perfectly pure sample of the sixth element in the periodic table. One of them wonders whether they could use a pencil lead as their sample. However, after discovering that the material from which the lead is made isn't pure enough for their purposes, the other engineer states: 'That isn't carbon'. Intuitively, the utterance is true. By contrast, suppose that a pencil maker asks his apprentice what the lead is made of, and that the apprentice says: 'That is carbon'. Despite the impurity in the material, the utterance is intuitively true[43].

ARGUMENT 2. Petrovic has invented a machine that purifies petroleum. The heart of this invention is a device that the scientist calls the *Petrovic membrane*. Yet this membrane has a strange property: if petroleum is poured straight into it, the sample is not purified in any way. For the membrane to work, one must first add a small amount of a certain substance, XYZ, to the petroleum. As it so happens, the easiest way to get

[42] I focus on presenting the relevant data. This data can be used to show that 'carbon' and 'filter' are context-sensitive in the same way in which the Attacker used the utterances of (1a) and (1b) to show that 'carbon filter' is context-sensitive.

[43] An alternative CSA compares a context in which diamond doesn't count as 'carbon' to a context in which diamond does count as 'carbon'. If 'That isn't carbon' is said of a diamond in the first context, the utterance is intuitively true. Similarly, if 'That is carbon' is said of the same diamond in the second context, the utterance is intuitively also true.

the right amount of XYZ into the petroleum is by pouring it through the kind of cartridge that is commonly found in certain water jugs[44]. So the machine purifies petroleum in two stages: first XYZ is added in the cartridge, and then the petroleum is processed through the Petrovic membrane.

Now Petrovic's machine is presented at a science exhibition. One of the visitors, John, suddenly becomes desperate to purify a small amount of water. As he sees Petrovic's machine, John points at the cartridge and exclaims 'That is a filter', planning to steal the cartridge later on; his utterance is intuitively true. At the same time, another scientist points at the cartridge and asks Petrovic whether the petroleum is purified in the cartridge. Petrovic replies that this is not the case and utters 'That isn't a filter'. Intuitively, his utterance is true: as he explains his machine, only the Petrovic membrane deserves to be called 'filter'[45].

As ARGUMENT 1 and ARGUMENT 2 are CSAs, it is difficult to see how the Attacker could reject the arguments, and how she could resist concluding that 'carbon' and 'filter' are context-sensitive. For, the Attacker used a CSA to support (C2), whereby she is committed to holding that CSAs are valid arguments in principle. Further, the data in ARGUMENT 1 and ARGUMENT 2 seems no less compelling than the data in the Attacker's CSA. E.g. the intuition that Petrovic's utterance is true is no weaker than the intuition that Mary's utterance of 'Vauen doesn't produce carbon filters' is true. Likewise, that 'that' has the same content in both the engineers' and the apprentice's context is no less plausible than the claim that 'Vauen' and 'produces' have the same contents in c_{pipe} and c_{chem}. In sum, since the Attacker uses a CSA to support (C2), she must endorse ARGUMENT 1 and ARGUMENT 2 also. In other words, the Attacker cannot plausibly reject that 'carbon' and 'filter' are context-sensitive.

[44] These cartridges, or the jugs themselves, are what one would normally call a water *filter*, but in describing the case, I want to avoid using 'filter'.
[45] One can suppose that the cartridges do not have any purifying effect on the petroleum.

5.2 A Pragmatic response to CSAs

To reinforce the conclusion just reached, I shall consider what is arguably the most common response to CSAs[46]. I shall call it the Pragmatic Error Response (PER). As I shall argue, the PER doesn't allow the Attacker to reject ARGUMENT 1 or ARGUMENT 2.

In outline, the PER makes three claims about ARGUMENT 1 and ARGUMENT 2: (i) the engineer's utterance of 'That isn't carbon', and Petrovic's utterance of 'That isn't a filter', *semantically express* a false proposition; (ii) they *pragmatically communicate* a true proposition, however; (iii) ordinary speakers mistake the communicated proposition for the semantically expressed proposition and therefore intuit the two sentence-utterance to be true[47].

To elaborate the PER for ARGUMENT 1, its proponent holds that 'carbon' is context-insensitive. More particularly, she holds that the content of 'carbon' is a property that is instantiated by typical carbon samples, so that 'That is carbon' can truthfully be said of all typical carbon samples. For present purposes, let's assume that the property in question is the property of being a carbon sample that is at least as pure as the average carbon sample[48]. Let's refer to this property as AVR.

Next, since 'carbon' expresses AVR at all contexts, the engineer's utterance of 'That isn't carbon' expresses that the demonstrated carbon sample doesn't instantiate AVR. Consequently, the sentence-utterance semantically expresses a false proposition, at least assuming that the demonstrated carbon sample is an average sample.

Turning to the utterance's pragmatic content, the PER lover holds that the utterance pragmatically communicates the *strongest true claim* about carbon that the engineer can

[46] E.g. Unger (2002) and Stanley (2005) consider this type of reply.
[47] This is only one version of the PER. Another version reverses (i) and (ii): the utterances of 'That isn't carbon' and 'That isn't a filter' are *true*, while the utterances of 'That is carbon' and 'That is a filter' are *false*; the latter two utterances communicate a true content. I shall only consider the version that is mentioned in the main text, but everything I say applies mutatis mutandis to the alternative version.
[48] According to the PER's alternative version, 'carbon' always expresses the property of being 100 percent pure carbon.

make. To spell this out, note that there are several AVR-like properties: one of these is only instantiated by carbon samples that are at least 98 percent pure; another is only instantiated by carbon samples that are at least 99 percent pure; and a third is only instantiated by carbon samples that are 100 percent pure. Further, suppose that the engineers' carbon sample instantiates the first of these properties, but not the second or the third. Then the PER lover will hold that the engineer's utterance communicates that the carbon sample is less than 99 percent pure. This proposition being true, the engineer's utterance communicates a true proposition, and the PER for ARGUMENT 1 is complete.

Turning to ARGUMENT 2, the PER lover will make almost exactly the same claims about 'filter' and 'That isn't a filter' that she made about 'carbon' and 'That isn't carbon'. In particular, she holds that 'filter' is context-insensitive, and that it expresses a property that typical filters instantiate. Let's say it expresses the property of being a device for purifying some substance or other. Since the cartridge in Petrovic's machine instantiates this property, Petrovic's utterance of 'That isn't a filter' semantically expresses a false content.

It pragmatically communicates a true content, however, according to the PER lover: like there are several AVR-like properties, there are also properties that resemble the property of being a device for purifying some substance or other. E.g. there is the property of being a device for purifying water, and the property of being a device for purifying petroleum. The cartridge in Petrovic's machine instantiates the former property, but not the latter. Therefore, and because Petrovic is talking about purifying petroleum, his utterance of 'That isn't a filter' pragmatically communicates that the cartridge in his machine doesn't instantiate the property of being a device for purifying petroleum. This proposition being true, Petrovic's utterance communicates a true proposition, whence the PER for ARGUMENT 2 is complete.

For present purposes, we needn't decide whether the PER yields convincing semantic and pragmatic theories for 'carbon' and 'filter'. All that matters is that the PER can also be applied to the Attacker's CSA, as follows. 'Carbon filter' is context-insensitive; at all contexts, it expresses a property that is instantiated both by filters that are made of carbon and by filters for purifying carbon. Let's say it expresses the property of being a filter that stands to carbon in some relation or other. Then Mary's utterance of (1b) = 'Vauen does not produce carbon filters' semantically expresses a false content. It communicates a true content, however, according to the PER lover: besides the property of being a filter that stands to carbon in some relation or other, there is the property of being a filter that stands in the *made-of*-relation to carbon, and the property of being a filter that stands in the *used-for*-relation to carbon. The filters that Vauen produces instantiate the former property, but not the latter. Therefore, and because Mary is talking about devices for purifying carbon, her utterance of (1b) pragmatically communicates that Vauen doesn't produce filters that stand in the *used-for*-relation to carbon. This proposition being true, Mary's utterance communicates a true proposition, and the PER for the Attacker's CSA is complete.

As far as I can see, the PER that replies to the Attacker's CSA is just as convincing as the PERs for ARGUMENT 1 and ARGUMENT 2. In particular, that 'carbon filter' always expresses the property of being a filter that stands to carbon in some relation or other, is prima facie no less plausible than the claim that 'filter' always expresses the property of being a device for purifying some substance or other. Likewise, that the engineer's utterance merely communicates that his carbon sample is less than 99 percent pure, is no more plausible than the claim that Mary's utterance merely communicates that Vauen's filters aren't filters for purifying carbon.

Now, if all three PERs are equally plausible, one cannot plausibly reject one PER without rejecting the other two PERs also. As should be obvious, this means that the

Attacker must reject all three PERs: she can't accept that the third PER is a valid reply to her own CSA, whence she must admit that the other two PERs aren't valid replies to ARGUMENT 1 and ARGUMENT 2. Granted, this doesn't *quite* prove that she cannot accept her own CSA while rejecting ARGUMENT 1 and ARGUMENT 2: it is *possible* that there is an objection to ARGUMENT 1 and ARGUMENT 2 which isn't also an objection to her CSA. However, I do not know of any objection of this kind[49]. Therefore, I conclude that the Attacker must accept ARGUMENT 1 and ARGUMENT 2 as valid. In other words, the Attacker must accept that 'carbon' and 'filter' are context-sensitive.

5.3 Switching examples?

The foregoing considerations establish a result about 'carbon', 'filter', and 'carbon filter'; they don't establish the following, universal claim: For all phrasal compounds 'N_1N_2', if 'N_1N_2' is context-sensitive, 'N_1' and 'N_2' are also context-sensitive. Hence, if the Attacker really wanted to run the Compound Argument with a compound that has context-insensitive constituents, perhaps she can just switch to a different example: as pointed out at the very outset, e.g. 'toy shop' and 'library book' are intuitively also context-sensitive. So maybe the Attacker can maintain that 'toy' and 'shop', and 'library' and 'book' are context-insensitive, and run the Compound Argument with 'toy shop' or 'library book' instead of 'carbon filter'. To conclude this section, let me briefly discuss this option.

[49] It has been pointed out that CSAs that compare an utterance of 'It *is* the case that *p*' to an utterance of 'It is *not* the case that *p*' are stronger than CSAs that compare two utterances of 'It is the case that *p*'. (See DeRose 2009, cap. 2.3) The Attacker's CSA meets this requirement, as do ARGUMENT 1 and ARGUMENT 2. Further, in connection with Contextualism about 'know', it has been pointed out that CSAs that involve third-person knowledge-ascriptions ('*Keith* knows that *p*') are more reliable than CSAs that involve first-person knowledge-ascriptions ('*I* know that *p*'). (See DeRose 2009, cap. 2.6) Yet even if this restriction applies also to CSAs for 'carbon' and 'filter', ARGUMENT 1 and ARGUMENT 2 are not affected as they do not involve any first-person statements.

Switching to a different example isn't as easy as one might think: if the Attacker wanted to run the Compound Argument with 'toy shop' or 'library book', she would need to establish that the two compounds in question are phrases rather than words. Yet as can be seen from the discussion in section two, it will not be trivial to establish these claims[50]. And even if the Attacker can show that 'toy shop' and 'library book' are phrases, there is little hope that she can resist the CSAs for 'toy', 'shop', 'library', and 'book' while accepting the CSAs for 'toy shop' and 'library book'. To illustrate this, I shall briefly sketch a CSA for each of the six expressions in question.

TOY SHOP. Bill's store sells toys. As a hobby, he collects antique toys: so far he has got a tin soldier, a small wooden carriage, and a doll house. When asked what Bill does for a living, his wife says: 'Bill has a toy shop'. When discussing what toys are missing from Bill's collection, one of his friends remarks 'Bill doesn't have a toy shop'.

TOY. The company Bark & Howl sells rubber bones for dogs to chew on. When Kevin suggests getting the dog something to play with, his wife remarks 'Bark & Howl sells toys'. When Martha wants to get the children a gift, her husband recommends against going to Bark & Howl: 'Bark & Howl doesn't sell toys'.

SHOP. Jenna makes jewellery and sells it online. Explaining why Jenna spends so much time on the internet, Mark says 'Jenna has a shop'. When the Mafia considers whether they can extort protection money from Jenna by threatening to wreck her store, Don Corleone points out 'Jenna doesn't have a shop'.

LIBRARY BOOK. Jim is a book thief. He particularly likes to steal books about the history of libraries, but he does not steal from libraries. When another thief asks about

[50] In fact, both compounds appear to fare *worse* than 'carbon filter' in this regard. More particularly, 'toy shop' fails both the replacement test and the insertion test: 'a chip shop and a toy one' is as infelicitous as 'a watch maker and a shoe one', and 'toy High Street shop' and 'toy online shop' are on par with the infelicitous 'milk glass bottle. (At least according to my intuitions, it should be 'High Street toy shop' and 'online toy shop'.) Similarly, 'library book' fails the replacement test and the insertion test is inconclusive: 'a library book and a picture one' is infelicitous while 'library picture book' and 'library hardcover book' are questionable.

Jim's preferred loot, he is told: 'Jim steals library books'. When Berta worries that Jim's behaviour might bankrupt the local library, she is told: 'Jim doesn't steal library books'.

LIBRARY. Ben's friend has a large collection of history books at home, which Ben is allowed to use. The books are catalogued, but kept in cardboard boxes in the attic. There are also several old, beautifully furnished buildings full of books nearby, but Ben isn't allowed in. When asked why Ben knows so much about history, his wife replies 'Ben has access to a library'. When an artist is looking for a suitable interior for his new painting 'The Librarian', he is told: 'Ben doesn't have access to a library'.

BOOK. Kim loves her Kindle and doesn't keep any hardcopies at home. When asked what Kim spends all her money on, Keiko is told: 'Kim has lots of books'. When asked why Kim's study looks so empty, Phil is told: 'Kim doesn't have any books'.

The above CSAs still don't prove that there is *no* context-sensitive phrasal compound the constituents of which are context-insensitive. Yet they strongly suggest that the Attacker will have a very hard time finding a compound of this kind. For this reason, and because 'carbon filter' can be shown to be a phrase, the Attacker won't attempt to 'switch examples': she continues to work with 'carbon filter' and accepts that 'carbon' and 'filter' are context-sensitive.

Before moving on, three brief caveats are in order. Firstly, the above considerations are *not* intended to show that 'carbon' and 'filter' are context-sensitive: they show that one cannot plausibly hold that 'carbon filter' is context-sensitive *and* hold that its constituents are context-*in*sensitive. Given her commitment to (C2), it follows that the *Attacker* cannot plausibly deny that 'carbon' and 'filter' are context-sensitive. That 'carbon' and 'filter' are context-sensitive does not follow.

Secondly, the above considerations do *not* show that the *Defender* cannot deny that 'carbon' and 'filter' are context-sensitive. As was explained in section four, the Defender *needn't* grant (C2): she could endorse Semantic Minimalism instead and deny that 'carbon filter', and 'carbon' and 'filter', are context-sensitive. Since Semantic Minimalism is a fairly radical thesis, the Defender didn't want her reply to the Compound Argument to rely on it. Still, if the Defender really wanted to deny that 'carbon' and 'filter' are context-sensitive, she could.

Thirdly, Cappelen and Lepore argue that one must accept all CSAs if one accepts some[51]. By contrast, the above considerations merely show that one must accept the CSAs for 'carbon' and 'filter' (for 'toy' and 'shop' / for 'library' and 'book') if one accepts the CSA for 'carbon filter' (for 'toy shop' / for 'library book'). They furthermore strongly *suggest* that one must, in general, accept CSAs for the nouns 'N_1' and 'N_2' if one accepts a CSA for the (phrasal) compound 'N_1N_2'. Yet they show nothing about CSAs for verbs, adjectives, quantifiers, or any other kind of expression. Therefore, the above considerations lend no support to Cappelen and Lepore's ultimate conclusion, the claim that one must either endorse Radical Contextualism[52] or Semantic Minimalism.

With these caveats in place, I shall consider the Defender's reply to (C3#).

6. A semantic theory for 'carbon'

The Attacker now accepts that 'carbon' and 'filter' are context-sensitive. The Defender will go one step further and hold that 'carbon' and 'filter' are context-sensitive in such a way that (C3#) fails. The present section spells this claim out in detail.

[51] See Cappelen and Lepore 2005, esp. cap. 3.
[52] Radical Contextualism is the view that all, or almost all, words are context-sensitive. See e.g. Cappelen and Lepore 2005 for details.

6.1 Preliminaries

To clear the ground, recall how the Attacker used her CSA to support (C3#): having noted that (1a) and (1b) are intuitively true in c_{pipe} and c_{chem} respectively, she pointed out that 'carbon' intuitively denotes carbon in both contexts while 'filter' intuitively denotes filters in both contexts. Thus one might think that the Attacker's CSA establishes both that 'carbon filter' is context-sensitive and that the contents of 'carbon' and 'filter' do not vary across c_{pipe} and c_{chem}: one might think that the Attacker's CSA establishes both (C2) and (C3#).

It would be a mistake to think that the Attacker's CSA establishes (C3#): CSAs are always based on intuitions about truth-values, i.e. on intuitions about a semantic property of *whole sentences*. Therefore, CSAs never establish that a particular *sub-sentential* expression is context-sensitive, much less that a particular word is context-*in*sensitive[53]. Of course, if one has a CSA that shows e.g. that 'Kevin stinks' is context-sensitive, and if one *combines* this CSA with evidence that 'Kevin' is not responsible for the sentence's context-sensitivity, then one has an argument to the effect that 'stinks' is context-sensitive[54]. Still, in and of itself the CSA can only establish the context-sensitivity of the sentence. With this in mind, return to the Attacker's CSA.

When the Attacker first formulated her CSA, she assumed that 'Vauen' and 'produces' are not responsible for the context-sensitivity of 'Vauen produces carbon filters', and this assumption seems fair: we may assume that Vauen the pipe manufacturer is the only company called 'Vauen' that John and Henry and Jane and Mary know, so that it would be surprising if 'Vauen' referred to a different company in one of the two contexts. Further, although 'produces' may well be context-sensitive, it is plausible enough that it denotes the same relation in c_{pipe} and c_{chem}, viz. the relation

[53] See Dorr forthcoming, section 4.iv, for a more detailed discussion of this point.
[54] Of course, one must also rule out that 'Kevin stinks' contains a hidden variable or other covert material.

that manufacturers like Vauen bear to their products. However, that 'carbon' and 'filter' do not vary in content across c_{pipe} and c_{chem} is much more contentious: the Attacker now accepts that the two nouns are context-sensitive, i.e. that they vary in content from context to context; that they don't vary in content across c_{pipe} and c_{chem} will therefore be hard to establish. In particular, the carbon in Vauen's filters is pure graphite (or so we shall assume), while the carbon that chemists want to purify is not pure, and not necessarily graphite. Can we be *sure* that 'carbon' doesn't denote pure graphite in c_{pipe} while denoting impure carbon of some kind or other in c_{chem}? Similarly, filters for pipes are fairly small and simple devices while the filters that chemists use to purify carbon are large, sophisticated machines. Can we be *sure* that 'filter' doesn't denote small devices in c_{pipe} while denoting sophisticated machines in c_{chem}?

As the rhetorical nature of the questions suggests, I am strongly inclined to say that we can*not* be sure that 'carbon' and 'filter' do not vary in content across c_{pipe} and c_{chem}: while the Attacker may have the *intuition* that 'carbon' and 'filter' do not co-vary with 'carbon filter', her CSA does nothing to *establish* (C3#).

The dialectic situation is now as follows. The Attacker endorses (C3#), but she has given us no good reason for thinking that this claim is true. The Defender will therefore go right ahead and claim that (C3#) is false: she claims that 'carbon' and 'filter' co-vary in content with 'carbon filter'. This raises the question of how 'carbon' and 'filter' vary exactly: what are their contents at contexts at which 'carbon filter' denotes a filter that is made of carbon, as compared to the contents they have at contexts at which 'carbon filter' denotes a filter for purifying carbon? The next sub-section provides one possible answer to this question, while the following sub-section argues that the Defender has good reasons to believe that the answer is correct.

6.2 The contents of 'carbon'

To spell out the claim that 'carbon' and 'filter' co-vary in content with 'carbon filter', two background assumptions are needed. Firstly, I shall assume a fine-grained account of propositions, according to which propositions are structured entities that contain individuals, properties, relations, etc. E.g. the proposition which 'Hans is German' expresses contains Hans and the property of being German. Since the fine-grained account of propositions is widely accepted, I think that the assumption in question is fairly uncontentious[55].

Secondly I shall make an assumption about the semantics of adjectives. Some theorists hold that adjectives have properties as their contents, while others hold that adjectives express functions from properties to properties[56]. The Defender's claims can be spelled out in either framework, but it will be convenient to settle for one account or the other. I shall settle for the functional account. To illustrate it, consider the adjective 'yellow' and let *d* be the function that is its content. If *d* is applied to the property of being a giraffe, *d*'s value is the property of being a yellow giraffe. More generally, if P is a property, *d*(P) is a property that an object instantiates if and only if it instantiates both P and the property of being yellow.

Two adjectives will be particularly important to the current debate, viz. 'carbon-made' and 'carbon-purifying'. For ease of reference, let $f_{\text{c-m}}$ and $f_{\text{c-p}}$ be the functions that these adjectives have as contents[57, 58]. If $f_{\text{c-m}}$ is applied to the property of being a filter,

[55] For a discussion of the fine-grained account, see e.g. Bealer 1979, Salmon 1986 and 1989, and Soames 1985, 1987, and 2010.
[56] For further details, see Kennedy 2012.
[57] It is conceivable that 'carbon-made' and 'carbon-purifying' are adjective *phrases*, and that they express semantically complex contents instead of $f_{\text{c-m}}$ and $f_{\text{c-p}}$. However, it is beyond the scope of this paper to discuss the syntax and semantics of noun-adjective compounds like 'carbon-made' and 'carbon-purifying'. Luckily for the Defender, the success of her reply to (C3#) does not depend on any particular assumption about the syntax and semantics of 'carbon-made' and 'carbon-purifying'. For simplicity, I shall therefore assume that both expressions are words, and that they express the said functions.
[58] For simplicity, I shall assume throughout that 'carbon-made' and 'carbon-purifying' are context-insensitive, but nothing depends on this.

say, it yields the property of being a filter that is made of carbon. If $f_{c\text{-}p}$ is applied to the property of being a filter, it yields the property of being a filter for purifying carbon.

With these assumptions in place, the Defender's claims about the semantics of 'carbon' can be spelled out, thus:

(D1) For all contexts *c*, if 'carbon filter' denotes in *c* a filter that is made of carbon, then 'carbon' denotes $f_{c\text{-}m}$ in *c*.

(D2) For all contexts *c*, if 'carbon filter' denotes in *c* a filter for purifying carbon, then 'carbon' denotes $f_{c\text{-}p}$ in *c*.

If (D1) and (D2) are correct, (C3#) fails: let *c* be a context in which 'carbon filter' denotes a filter that is made of carbon, and let *c** be a context in which 'carbon filter' denotes a filter for purifying carbon. Then by (D1) and (D2), 'carbon' denotes $f_{c\text{-}m}$ in *c* while denoting $f_{c\text{-}p}$ in *c**. Thus *both* 'carbon filter' *and* 'carbon' have different contents at *c* and *c**. Since *c* and *c** were chosen arbitrarily, it follows that there is no pair of contexts c_1 and c_2 such that 'carbon' has the same content in c_1 and c_2, while 'carbon filter' has different contents in c_1 and c_2[59]. Thus (C3#) is false.

Before moving on, let me make a few clarifying remarks about (D1) and (D2). The first remark concerns the scope of the Defender's claims: while (D1) and (D2) concern only 'carbon filter', it should be clear that the Defender will endorse parallel claims for all other phrasal compounds. E.g. if 'toy shop' turns out to be a phrase, the Defender will hold that 'toy' expresses a function from properties to properties. Further, if 'toy shop' is context-sensitive (in addition to being a phrase), the Defender will hold that 'toy' co-varies in content with 'toy shop'. In that sense, the Defender's claims are fully

[59] This follows only if we assume that 'carbon filter' has only two possible contents. Yet if 'carbon filter' can have a third content – e.g. if it can denote a filter for removing carbon from exhaust gasses – the Defender will add an assumption to (D1) and (D2) that covers the contexts at which 'carbon filter' has the said content.

general. Yet since the Compound Argument features only 'carbon filter', I shall focus on this compound, and on (D1) and (D2).

The second remark concerns 'carbon'. Arguably, 'carbon filter' has a content at all contexts: at any given context, it either denotes filters that are made of carbon, or filters for purifying carbon. Given (D1) and (D2), this means that at any context, 'carbon' either expresses $f_{\text{c-m}}$ or $f_{\text{c-p}}$. Yet this is apparently false: e.g. when the pencil maker says 'That is carbon', his utterance expresses that the demonstrated sample is carbon. Hence at the context at which the pencil maker's utterance occurs, 'carbon' denotes *carbon*, rather than $f_{\text{c-m}}$ or $f_{\text{c-p}}$. I will return to this problem below, but shall set it aside for now: for now, grant that (D1) and (D2) are consistent with the claim that 'carbon' denotes *carbon* at some contexts.

The third remark concerns 'filter'. (D1) and (D2) together yield that the content of 'carbon' varies with the content of 'carbon filter'. If this is correct, the Attacker's third premise fails. Yet the premise also fails if the content of 'filter' varies with the content of 'carbon filter'. Hence one could imagine a reply to the Compound Argument that replaces (D1) and (D2) with similar statements for 'filter': 'filter' denotes *this-and-this* when 'carbon filter' denotes a filter that is made of carbon, and *that-and-that* when 'carbon filter' denotes a filter for purifying carbon. I don't mean to suggest that statements of this kind are less plausible than (D1) and (D2), or that they don't amount to a valid reply to the Compound Argument. Personally, I find it more natural to suppose that it is 'carbon' that co-varies in content with 'carbon filter'. Therefore, I won't consider any accounts on which 'filter' co-varies with 'carbon filter'.

The next remark concerns a tenet encountered earlier, viz. the tenet that 'carbon filter' contains a hidden variable. Prima facie, this tenet may seem very similar to the combination of (D1) and (D2). However, there are at least four important differences: (i) That 'carbon filter' contains a hidden variable is a *syntactic* claim about the *phrase*

'carbon filter', while (D1) and (D2) are *semantic* claims about the *word* 'carbon'. (ii) While (D1) and (D2) imply that 'carbon' co-varies in content with 'carbon filter', the syntactic tenet does not imply this. Hence while (D1) and (D2) are inconsistent with (C3#), the syntactic tenet is consistent with the Attacker's third premise. (iii) The syntactic tenet is inconsistent with (the second part of) premise (C1), while (D1) and (D2) are consistent with (C1). (iv) While (D1) and (D2) together imply that 'carbon' is context-sensitive, the syntactic tenet does not imply this. In sum, (D1) and (D2) are quite different from the tenet that 'carbon filter' contains a hidden variable. Therefore, the objections from section three against the syntactic tenet do not apply to (D1) and (D2).

The final remark concerns 'carbon-made' and 'carbon-purifying'. Above I *assumed* that these adjectives respectively express $f_{c\text{-}m}$ and $f_{c\text{-}p}$. *If* that assumption is correct, the Defender endorses (D1) and (D2). If it turns out, however, that $f_{c\text{-}m}$ and $f_{c\text{-}p}$ aren't the contents of 'carbon-made' and 'carbon-purifying', the Defender will adjust her tenets accordingly. E.g. if it turns out that 'carbon-made' and 'carbon-purifying' express properties – $P_{c\text{-}m}$ and $P_{c\text{-}p}$, say – then the Defender will hold that 'carbon' expresses $P_{c\text{-}m}$ at all contexts at which 'carbon filter' denotes a filter that is made of carbon, while expressing $P_{c\text{-}p}$ at all contexts at which 'carbon filter' denotes a filter for purifying carbon. To put the matter more generally, what contents 'carbon' expresses at the said contexts *doesn't matter* (much) to the Defender: what matters is that the content of 'carbon' co-varies with the content of 'carbon filter'.

With these clarifying remarks in place, I turn to showing that the Defender has good reasons to endorse (D1) and (D2).

6.3 Supporting (D1) and (D2)

If (D1) and (D2) are correct, John's and Mary's utterances of (1a) and (1b) are true, in accordance with out intuitions. Further, (D1) and (D2) together imply that (C3#) is false. Hence (D1) and (D2) yield a semantic account for 'carbon filter' which explains the relevant linguistic intuitions, while reconciling the context-sensitivity of 'carbon filter' with Compositionality. This alone already makes it quite plausible for the Defender to endorse (D1) and (D2). Yet as I shall now show, the Defender has another, more specific reason for endorsing (D1) and (D2). I shall focus on (D1), but everything I say applies mutatis mutandis to (D2).

To start, note that in a framework of structured propositions, the following principle is a very plausible addition to Compositionality.

CONSTITUENCY: If *k* is a constituent of proposition *p*, and if sentence S expresses *p* in context *c*, then some constituent of S has *k* as content in *c*.

Hence since the Defender has adopted a framework of structured propositions, she endorses CONSTITUENCY also.

Next, let c_{D1} be a context in which 'carbon filter' denotes a filter that is made of carbon. Then (1a) = 'Vauen produces carbon filters' expresses in c_{D1} that Vauen produces filters that are made of carbon. Thus in c_{D1}, (1a) and (14) arguably have the same content[60].

(14) Vauen produces carbon-made filters.

The proposition that (14) expresses at c_{D1} can be represented thus:

P_{14}: $\langle\langle V, \langle f_{c\text{-}m}, \text{FILTER}\rangle\rangle, \text{PRODUCE}\rangle$

[60] I shall defend this claim presently.

Here 'V' represents Vauen and PRODUCE is the producing-relation. Further, recall that $f_{\text{c-m}}$ is, by assumption, the function that the adjective 'carbon-made' expresses. When applied to FILTER (i.e. the property that 'filter' expresses at c_{D1}), it returns the property of being a filter that is made of carbon.

Now, if (1a) and (14) really express the same content at c_{D1}, then P_{14} is the content of (1a) at c_{D1}. Hence by CONSTITUENCY, each constituent of P_{14} is, at c_{D1}, the content of some constituent of (1a). More fully, (1a) has a constituent that has V as content at c_{D1}, a constituent that has $f_{\text{c-m}}$ as content at c_{D1}, a constituent that has FILTER as content at c_{D1}, and a constituent that has PRODUCE as content at c_{D1}. Very plausibly, V, FILTER and PRODUCE are respectively the contents of 'Vauen', 'filters', and 'produces' at c_{D1}. Therefore, $f_{\text{c-m}}$ must be the content of 'carbon' at c_{D1}. Since c_{D1} was chosen arbitrarily, it follows that $f_{\text{c-m}}$ is the content of 'carbon' at all contexts at which 'carbon filter' denotes a filter that is made of carbon. In other words, (D1) is true.

Using a sentence like (15), one can furthermore run a parallel argument to the effect that 'carbon' denotes $f_{\text{c-p}}$ in all contexts in which 'carbon filter' denotes a filter for purifying carbon.

(15) Vauen produces carbon-purifying filters.

In sum, the argument just presented shows that 'carbon' varies in content in exactly the way in which (D1) and (D2) claim it varies. Thus, the argument gives the Defender a strong reason to endorse (D1) and (D2).

6.4 Structured vs. unstructured contents

The above argument assumes that (1a) and (14) express the same content at c_{D1}; call this the Same-Content Assumption, or SCA. In this sub-section, I shall explain why it is

plausible (for the Defender) to endorse SCA. For this purpose, consider the following view, on which SCA fails.

> 'Carbon filter' has an *unstructured* entity as its content, at all contexts: at some contexts, such as c_{D1}, its content is the property of being a filter that is made of carbon; at other contexts, its content is the property of being a filter for purifying carbon. Therefore, the contents that 'carbon' and 'filter' have at c_{D1} are not constituents of the content that 'carbon filter' has at c_{D1}, nor of the proposition that (1a) expresses at c_{D1}. Hence while $f_{c\text{-}m}$ and FILTER are constituents of P_{14}, they are not constituents of the proposition that (1a) expresses at c_{D1}. And while the latter proposition has as constituent the property of being a filter that is made of carbon, the former proposition does not have this property as a constituent. Thus, (1a) and (14) do not express the same content at c_{D1}.

If 'carbon filter' has a property as its content at every context, then it is plausible that *all* compounds express a property at every context. Thus the proposal at hand is that for all compounds κ and contexts c, κ expresses a property at c. Call this the Property View. The question is why the Defender finds the Property View less plausible than SCA.

One disadvantage of the Property View is that it gives rise to a fairly complicated taxonomy of compounds. On the Property View, both verbal and phrasal compounds express the same kind of content, viz. a property. Hence from a semantic point of view, phrasal compounds are grouped together with the verbal compounds according to the Property View: both are semantically simple. However, while phrasal compound like 'carbon filter' have constituents, verbal compounds don't have constituents. Hence from a syntactic point of view, verbal and phrasal compounds are *not* grouped together: the former are syntactically simple while the latter are syntactically complex. Thus the

Property View is committed to a fairly complex taxonomy which places verbal and phrasal compounds in the same semantic taxon, but in different syntactic taxa.

By contrast, the Defender's SCA gives rise to a fairly simple taxonomy of compounds, viz. to one in which phrasal compounds are grouped together *both* syntactically *and* semantically. On her view, all phrasal compounds have a structured content at all contexts, of the form of $\langle f, P \rangle$ – f being a function from properties to properties and P being a property. Hence on her view, syntactically complex compounds are also semantically complex, and syntactically simple compounds are also semantically simple. Thus as desirable, simple theories come with simple taxonomies, the Defender's view has an advantage over the Property View.

One might object that there is another kind of simple taxonomy. Suppose that *all* expressions have unstructured entities as their contents (at all contexts). Then the syntactic simple/complex distinction will cut across the semantic simple/complex distinction, but this cut will be completely systematic: *no* expression will be both syntactically and semantically complex. Hence that phrasal compounds are semantically simple but syntactically complex is fully predicted by the taxonomy and hence no detriment to its simplicity or desirability.

While it may be true that the described taxonomy is as simple and desirable as the Defender's, the underlying theory is not currently available: the Defender assumed a fine-grained theory of propositions, which rules out that all expressions have unstructured entities as their contents: the contents of sentences are structured. Of course, the proponent of the Property View could reject the fine-grained theory of propositions and adopt a possible-worlds theory of propositions. Yet the question is why *the Defender* finds SCA more plausible than the Property View. Thus, showing that the

Property View is plausible *if one rejects one of the Defender's tenets* won't convince the Defender that the Property View is more plausible than SCA[61].

In sum, SCA allows for a simpler taxonomy of compounds than the Property View, and is hence to be preferred.

There is another reason for preferring SCA to the Property View. Note that the brief argument in (16a) appears to be logically valid. It seems to be on par with the logically valid argument in (16b), rather than with the logically invalid argument in (16c).

(16a) Vauen produces carbon filters. Thus, Vauen produces filters.

(16b) BMW produces expensive cars. Thus, BMW produces cars.

(16c) Mercedes produces cabriolets. Thus, Mercedes produces cars.

Now, if 'carbon filter' expresses the same kind of content as 'expensive car', viz. a content of the form of <function, property>, then (16a) is indeed on par with (16b). By contrast, if 'carbon filter' expresses the same kind of content as 'cabriolet', viz. a property, then (16a) is on par with (16c), contrary to our intuitions. Hence while SCA yields the intuitively correct result, the Property View does not. Thus, the Defender prefers SCA to the Property View.

The proponent of the Property View might object that (16a) is in fact *in*valid. Prima facie, (16a) and the arguments in (17) are instances of the same inference pattern.

(17a) Perkin is an impostor king. Thus, Perkin is a king.

(17b) Jane bought an imitation pearl. Thus, Jane bought a pearl.

[61] Adopting a possible-worlds theory of propositions may not be a viable option for the proponent of the Property View anyway. Ultimately, the point of the Property View is to convince the Defender that SCA is false, i.e. that (1a) and (14) don't express the same proposition at c_{D1}. However it is difficult to imagine a possible world in which it is true (false) that Vauen produces filters that are made of carbon, but false (true) that Vauen produces carbon-made filters.

Hence if (16a) is valid, so are the arguments in (17). However, the objection continues, the arguments in (17) are clearly invalid. Thus, (16a) is also invalid and on par with (16c), contrary to SCA and in accordance with the Property View.

The Defender has reason to deny that (16a) and the arguments in (17) instantiate the same inference pattern. Roughly, (I_1) is the inference pattern of (16a) while (I_2) is the pattern that (16c) instantiates.

(I_1) S φs N_1N_2s. ∴S φs N_2s.

(I_2) S φs Ns. ∴S φs N*s.

Now, if 'imposter king' and 'imitation pearl' are phrases, then (17a) and (17b) are indeed instances of (I_1). However, the Defender has reason to hold that the two compounds in question are words: both compounds fail the modification test and the replacement test; the insertion test is negative for 'imposter king' and inconclusive for 'imitation pearl'[62]. Thus as 'imposter king' and 'imitation pearl' are words, (17a) and (17b) are instances of (I_2), rather than of (I_1), and the objection fails.

In sum, the Defender has two good reasons for preferring SCA to the Property View. In turn, if SCA is true, (D1) and (D2) can be supported with the above argument. Thus, the Defender has a convincing reason for accepting (D1) and (D2).

7. Objections

In closing, I would like to discuss three objections that one might raise against the Defender's account.

[62] Modification: in e.g. 'convincing impostor king' and 'obvious imitation pearl' the adjectives modify the entire compound and not only the dependent noun. Replacement: 'a soldier king and an impostor one' is infelicitous, as is 'a button pearl and an imitation one'. Insertion: 'impostor soldier king' is infelicitous, and 'imitation button pearl' is questionable.

7.1 Names

To formulate the first objection, note that many compounds have a proper name as their dependent: ‘Shakespeare admirer’, ‘London expert’, ‘Ford employee’, ‘Wagner opera’, ‘Stanwell pipe’, ‘Thatcher biography’, ‘Elvis concert’. Now suppose that Ben utters ‘Every Shakespeare admirer has visited Stratford-upon-Avon’. The Defender will hold that ‘Shakespeare’ has a function as its content at the context of Ben’s utterance, call it f_S. When applied to the property of being an admirer, it yields the property of being an admirer of Shakespeare’s. However, Kripke taught us that the content of a proper name is always the name-bearer – Shakespeare, in the case at hand. Thus, the Defender’s account fails because it contradicts Kripke’s results.

Reply. Although the Defender’s account deviates from Kripke’s, it still preserves Kripke’s insights. For instance, consider ‘Shakespeare might not have been the author of Hamlet’. Kripke points out, and his account implies, that this sentence is true: in a world in which John Doe, say, authored Hamlet, Shakespeare is still Shakespeare, but not the author of Hamlet. Similarly, ‘Shakespeare admirers might not have admired the author of Hamlet’ is true according to the Defender’s account. For, f_S maps the property of being an admirer to the property of being an admirer of *Shakespeare’s*, and that latter property is instantiated only by people who admire Shakespeare himself: in a world in which John Doe wrote Hamlet, people who admire the author of Hamlet don’t automatically instantiate the property in question. Thus, although the Defender’s account deviates from Kripke’s, the deviation is inconsequential.

To reinforce this point, recall that Montague treats names as generalised quantifiers. Therefore, Montague’s account also contradicts Kripke’s claim that the content of a name is always the name-bearer. However, to my knowledge, nobody has ever objected to Montague’s account on those grounds. The reason is the same as before: although the

account deviates from Kripke's, it preserves all of Kripke's insights[63]. In sum, if one wants to object to the Defender's account that it deviates from Kripke's, one must raise the same objection against Montague. Neither instance of the objection is compelling.

7.2 'Carbon' means *carbon*!

Suppose John utters 'Bill bought some carbon'. Presumably, 'carbon filter' has a content relative to the context in which John's utterance occurs: it either denotes a filter for purifying carbon or a filter that is made of carbon. Hence by (D1) and (D2), 'carbon' either has the same content as 'carbon-made' in John's context, or the same content as 'carbon-purifying'. Consequently, the Defender's account predicts that John's utterance expresses either that Bill bought some(thing) carbon-made, or that Bill bought some(thing) carbon-purifying. Yet John's utterance expresses neither: it clearly expresses that Bill bought some *carbon*. Thus, the Defender's account makes a false prediction about John's utterance.

Reply. The problem is not specific to the Defender's account. E.g. Contextualism about 'know' faces the same challenge. According to this view, every context is associated with an epistemic standard. At contexts that are associated with a strict epistemic standard, 'know' expresses a relation that an agent bears to a proposition *p* only if the agent has very good evidence for *p*. At contexts that are associated with a more lenient epistemic standard, 'know' expresses a more inclusive relation: in order to

[63] To return to the above example, consider again 'Shakespeare might not have been the author of Hamlet'. On Montague's account this sentence expresses, roughly, that the property of having authored Hamlet (AUTH) might not have instantiated the second-order property of being instantiated by Shakespeare (SHAKE). Further, AUTH instantiates SHAKE if and only if Shakespeare instantiates AUTH. Hence in a world in which John Doe authored Hamlet, AUTH does not instantiate SHAKE. Thus 'Shakespeare might not have been the author of Hamlet' is true on Montague's account.

bear that relation to *p*, the agent doesn't need particularly strong evidence for *p*. Let's refer to the two relations in question as k_{strong} and k_{weak}[64].

Now, since every context is associated with an epistemic standard, it follows that for every context *c*, 'know' either has k_{weak} or k_{strong} as content at *c*. Yet when John utters 'Bill knows who Barack Obama is', 'know' apparently has neither k_{weak} nor k_{strong} as its content. His utterance expresses, roughly, that Bill has a certain amount of identifying information about Obama; questions about evidence don't enter into the picture. Thus, Contextualism about 'know' faces the same challenge as the Defender.

The Contextualist about 'know' can respond to the challenge by reformulating her account: she can claim that the account isn't a theory about the semantics of 'know' *per se*, but only a theory about 'know' *as it occurs in* 'knowing that'-constructions. Alternatively, the account only concerns the semantics of certain *sentences*, viz. of sentences of the form of 'S knows that *p*'. Along these lines, DeRose writes that Contextualism about 'know' is "a theory according to which the truth-conditions of […] sentences of the form 'S knows that p' and 'S does not know that p' and related variants of such sentences […] vary in certain ways"[65]. And although DeRose does not say it explicitly, his statement suggests that 'knowing-who'-constructions aren't "related variants" of sentences of the form of 'S knows that p'[66]. Thus, the Contextualist meets the challenge at hand by restricting her account.

Returning to the Defender, she can meet the challenge in much the same way. In particular, she can claim that her account is not a theory about the semantics of 'carbon' *per se*; it is only a theory about 'carbon' *as it occurs in* 'carbon filter' (and other compounds). Alternatively, her account only concerns sentences in which 'carbon filter'

[64] As should be obvious, this is a simplified version of Contextualism about 'know': most Contextualists hold that there are more than two relations that 'know' can express.
[65] DeRose 2009, p. 2.
[66] DeRose later compares Contextualism about 'knowing-that'-constructions to Contextualism about 'knowing-who'-constructions. The theory he gives for the latter constructions is similar to, but clearly distinct from his theory for 'knowing-that'-constructions (See DeRose 2009, appendix to cap. 2.)

occurs, but not sentences like ‘Bill bought some carbon’, in which ‘carbon’ occurs in isolation. Let me propose two ways of spelling this out.

The first proposal focusses on the syntax of ‘carbon filter’. To start, note that ‘carbon’ is used as a modifier in (1a) = ‘Vauen produces carbon filters’, while being used as an NP head in ‘Bill bought some carbon’. In other words, note that ‘carbon’ has more than one grammatical function. Now, according to the current proposal, this fact must be taken into consideration when assigning a content to ‘carbon’. In particular, given any context *c*, there is no such thing as *the* content of ‘carbon’ at *c*. Rather, ‘carbon’ has (at least) *two* contents relative to any context, viz. a modifier-content and a head-content, as it were. Its modifier-content is always either $f_{\text{c-m}}$ or $f_{\text{c-p}}$; its head-content is always the property of being carbon.

To apply this proposal to sentences, let *s* be a sentence and suppose that ‘carbon’ is used in *s* as a modifier. Then when *s* is evaluated a given context *c*, ‘carbon’ contributes the modifier-content that it has at *c* to the proposition that *s* expresses at *c*. Similarly, if ‘carbon’ is used in *s* as NP head, it contributes its head-content to the proposition that *s* expresses at *c*. Since ‘carbon’ functions as a head in ‘Bill bought some carbon’, it contributes its head-content to the proposition that the sentence expresses. Hence the sentence expresses that Bill bought some carbon, and John’s sentence-utterance expresses this proposition also, as desired.

Depending on one’s ontology of words, the current proposal will find wide application. According to one ontology, e.g. the verb ‘to cook’ and the noun ‘cook’ are distinct words that are merely homonymous (in some of their forms). Yet according to another ontology, there is just one word ‘cook’ in English: like ‘carbon’, it has more than one grammatical function, and is used as a verb in some sentences and as a noun in others. If this latter ontology is correct, it stands to reason that ‘cook’ has more than one content at any context, viz. a verbal content and a nominal content. Its verbal content

may be the property of preparing food, while its nominal content is the property of being a chef. If 'cook' is used as a verb (noun) in a given sentence *s*, it contributes its verbal content (nominal content) to the proposition that *s* expresses.

To summarise the first proposal, the Defender's account only concerns 'carbon' as it occurs in 'carbon filter' in the sense that (D1) and (D2) are only statements about the modifier-content of 'carbon', but not statements about its head-content.

Let me present the second proposal, which appeals to the distinction between sentences-in-context and sentence-utterances. To start, there are *three* kinds of contexts according to this proposal: those at which 'carbon filter' denotes a filter that is made of carbon, those at which 'carbon filter' denotes a filter for purifying carbon, and those at which 'carbon filter' denotes neither. (One might stipulate that 'carbon filter' denotes the empty set at the third kind of context.) Let's refer to these contexts, respectively, as *m*-contexts, *p*-contexts, and 0-contexts. Further, $f_{\text{c-m}}$ is the content of 'carbon' at all *m*-contexts, while $f_{\text{c-p}}$ is its content at all *p*-contexts, in accordance with (D1) and (D2). However, at 0-contexts, 'carbon' simply denotes carbon.

With *m*-, *p*-, and 0-contexts in place, the proposal is this: although one can *evaluate* the *sentence* 'Bill bought some carbon' at *m*-contexts and at *p*-contexts (and at 0-contexts, of course), no *utterance* of 'Bill bought some carbon' ever occurs in an *m*- or *p*-context: all of the sentence's utterances occur in 0-contexts, so that each of its utterances expresses that Bill bought some carbon[67]. Likewise, one can evaluate (1a) at 0-contexts, but each utterance of (1a) occurs either in an *m*-context or in a *p*-context. As should be obvious, the proposal implies correctly that John's utterance of 'Bill bought some carbon' expresses that John bought some carbon.

[67] One can remain neutral about what 'Bill bought some carbon' expresses at *m*- and *p*-contexts: one could hold that its character is simply not defined at these contexts, or that the sentence expresses that Bill bought something carbon-made / carbon-purifying.

One might object that it is *ad hoc* to claim that utterances of 'Bill bought some carbon' never occur in *m*- or *p*-contexts, and that utterances of (1a) never occur in 0-contexts. Yet I don't think that this is a convincing objection. For, what context an utterance occurs in depends – among other things – on what was uttered: E.g. if one says 'By basketball standards, Ella is short', one's utterance plausibly occurs in a context in which basketball standards for tallness are operative[68], so that 'short' expresses a property that even fairly tall people instantiate. Further, it is at least conceivable that one's utterance occurs in this context *because one uttered* 'by basketball standards'[69]. In turn, if this proposal is correct, the Defender can apply it to (1a) and 'Bill bought some carbon': utterances of (1a) always occur in *m*- or *p*-contexts *because one utters* 'carbon filter' when uttering (1a); and utterances of 'Bill bought some carbon' always occur in 0-contexts because in uttering the sentence, one utters 'carbon' in isolation.

To summarise the second proposal, the Defender's account only concerns 'carbon' as it occurs in 'carbon filter' in the sense that (D1) and (D2) are only statements about the contexts in which the utterances of (1a) occur, but not statements about all contexts at which (1a) can be evaluated.

7.3 Multiple occurrences of 'carbon'

The third objection tries to strengthen the second. Suppose Jane utters 'Frank put some carbon into the carbon filter'. Even if the Defender's account only concerns 'carbon' as it occurs in 'carbon filter', her account concerns the sentence that Jane utters. For, 'carbon' is used as a modifier in the sentence in question, and Jane utters 'carbon filter' in uttering the sentence. Hence according to the Defender, 'carbon'

[68] I am assuming for a moment that 'short' is context-sensitive.
[69] Incidentally, some Contextualists about 'know' make a similar claim when talking about disagreements: see e.g. DeRose 2009, cap. 4.

denotes either $f_{c\text{-}m}$ or $f_{c\text{-}p}$ at the context in which Jane's utterance occurs. Consequently, the Defender's account predicts that Jane's utterance expresses either that Frank put some(thing) carbon-made into the carbon-made filter, or that Frank put some(thing) carbon-purifying into the carbon-purifying filter. Yet Jane's utterance expresses neither: it expresses that Frank put some *carbon* into the carbon-made / carbon-purifying filter. Thus, the Defender's account makes a false prediction about Jane's utterance.

Reply. While the objection identifies an interesting problem[70], the problem is not specific to the Defender's account. To see this, consider the sentences in (18).

(18a) The cartographer played an important part in charting the northern part of the jungle.

(18b) The general gave an order to place an order with Amazon.

(18c) The border for her dress was designed near the Russian border.

(18d) When the glazier showed Ben the bill for replacing the glass, he dropped his glass.

(18e) Peter caught a cold from staying in the cold for too long.

(18f) Ella broke her foot near the foot of a mountain.

Each of these sentences gives rise to the same problem that the objector raises against the Defender's account. Focussing on (18a), ordinary utterances of this sentence express P_1 rather than P_2 or P_3.

P_1 that the cartographer played an important *role* in charting the jungle's northern *area*.

P_2 that the cartographer played an important *role* in charting the jungle's northern *role*.

P_3 that the cartographer played an important *area* in charting the jungle's northern *area*.

Yet it is not at all easy to explain why utterances of (18a) normally express P_1: according to classical accounts of context-sensitivity[71], a given sentence-utterance *u*

[70] A particular version of this problem is discussed in detail in the first paper in this dissertation.
[71] See e.g. Kaplan 1989.

expresses that *p* iff the uttered sentence expresses that *p* at the context in which *u* occurs. In turn, what content a given sentence *s* expresses relative to a given context *c* depends on what contents *s*'s constituents have at *c*. Finally, standard assumptions about the compositional structure of (18a) imply that (18a) expresses P_2 at all contexts at which 'part' means *role*, while expressing P_3 at all contexts at which 'part' means *area*. Thus classical accounts of context-sensitivity predict, falsely, that no utterance of (18a) expresses P_1.

Let's speak of the *Multiple Occurrence Problem*, or MOP, to mean the problem that the sentences in (18) exhibit. If the Defender's account is correct, Jane's utterance of 'Frank put some carbon into the carbon filter' is another instance of the MOP. Yet that is all it is: another *instance* of a rather *general* problem. And one can hardly fault the Defender for not solving the general problem. At the same time, whatever the correct solution to the general problem may be, it will be applicable to Jane's utterance – exactly because her utterance is an instance of the MOP and not a problem that is specific to the Defender's account. In sum, the Defender can reply to the current objection by referring the objector to whatever the correct solution to the general problem is.

8. Conclusion

The Compound Argument poses an interesting challenge for the proponent of Compositionality. Unless one is inclined towards Semantic Minimalism, one cannot plausibly deny that 'carbon filter' and similar compounds are context-sensitive. Given Compositionality, this means that one must develop a semantic theory for 'carbon' or 'filter' (or both), according to which 'carbon' or 'filter' co-vary in content with 'carbon filter'. I showed that such a theory can be developed by assuming that 'carbon' has an

adjectival content when it occurs in 'carbon filter'. I further argued that the proponent of Compositionality has good reasons to endorse this assumption. I thus conclude that context-sensitive compounds do not pose a threat to Compositionality.

References

Austin, J. L. (1979). 'Other Minds', in J. L. Austin, Philosophical Papers, 3rd edn., Oxford University Press 1979, 76 – 116.

Bauer, L. (1998). 'When is a sequence of two nouns a compound in English?', *English Language and Linguistics* **2**, 65 – 86.

Bealer, G. (1979). 'Theories of Properties, Relations and Propositions', *Journal of Philosophy* **76**, 634 – 648.

Bell, M. J. (2005). 'Against nouns as syntactic premodifiers in English noun phrases', *Working Papers in English and Applied Linguistics* **11**, 1 – 48.

Cappelen, H. and Lepore, E. (2002). 'Indexicality, binding, anaphora, and *a priori* truth', *Analysis* **62**, 271 – 81.

Cappelen, H. and Lepore, E. (2005). 'Insensitive Semantics. A Defense of Semantic Minimalism and Speech Act Pluralism', Blackwell.

Chomsky, N. (1980). 'Rules and Representations', Columbia University Press.

Cohen, J. and Rickless, S. C. (2009). 'Binding arguments and hidden variables', *Analysis* **67**, 65 – 71.

Cresswell, M. (1973). 'Logics and Languages', London: Methuen.

Davies, M. (1983). 'Idiom and Metaphor', *Proceedings of the Aristotelian Society* **83**, 67 – 86.

DeRose, K. (2009). 'The Case for Contextualism. Knowledge, Skepticism, and Context, Volume 1', Oxford University Press.

Dorr, C. (forthcoming). 'Transparency and the Context-Sensitivity of Attitude Reports', forthcoming in M. Garcia-Carpintero and G. Martí (eds.), 'Empty Representations: Reference and Non-existence', Oxford University Press 2014.

Farkas, D. (1997). 'Evaluation Indices and Scope', in A. Szabolcsi (ed.), 'Ways of Scope Taking', Kluwer 1997, 183 – 215.

Giegerich, H. (2004). 'Compound or phrase? English noun plus noun constructions and the stress criterion', *English Language and Linguistics* **8**, 1 – 24.

Hawthorne, J. (2006). 'Testing for Context-Dependence', *Philosophy and Phenomenological Research* **73**, 443 – 450.

Kaplan, D. (1989a). 'Demonstratives', in J. Almong, J. Perry, and H. Wettstein (eds.), 'Themes from Kaplan', Oxford University Press 1989, 481 – 563.

Katz, J. (1973). ‘Compositionality, Idiomaticity, and Lexical Substitution’, in S. Anderson and P. Kiparsky (eds.), ‘A Festschrift for Morris Halle’, New York: Holt, Rinehart, and Winston 1973, 357 – 76.

Kennedy, C. (2012). ‘Adjectives’, in G. Russell and D. Graff Fara (eds.), ‘The Routledge Companion to Philosophy of Language’, Routledge 2012, 328 – 341.

Lees, R. D. (1963). ‘The grammar of English nominalizations’. Bloomington: Indiana University and The Hague: Mouton.

Levi, J. (1978). ‘The Syntax and Semantics of Complex Nominals’, New York: Academic Press.

Lewis, D. K. (1970). ‘General semantics’, *Synthese* **22** 18 – 67.

Montague, R. (1973). ‘The proper treatment of quantification in ordinary English’, in R. Thomason (ed.), ‘Formal Philosophy: Selected Papers of Richard Montague’, Yale University Press 1974, 247 – 270.

Montague, R. (1974). ‘English as a formal language’, in R. Thomason (ed.), ‘Formal Philosophy: Selected Papers of Richard Montague’, Yale University Press 1974, 188 – 221.

Payne, J. and Huddleston, R. (2002). ‘Nouns and noun phrases’, in R. Huddleston and G. K. Pullum (eds.), ‘The Cambridge grammar of the English language’, Cambridge University Press, 323–524.

Recanati, F. (2002). ‘Unarticulated Constituents’, *Linguistics and Philosophy* **25**, 299 – 345.

Salmon, N. (1986). ‘Frege's Puzzle’, MIT Press.

Salmon, N. (1989). ‘Tense and Singular Propositions’, in J. Almong, J. Perry, and H. Wettstein (eds.), ‘Themes from Kaplan’, Oxford University Press 1989, 391 – 392.

Siegel, M. E. (1976) ‘Capturing the Adjective’, Ph.D. thesis, University of Massachusetts at Amherst.

Soames, S. (1985). ‘Lost Innocence’, *Linguistics and Philosophy* **8**, 59 – 71.

Soames, S. (1987). ‘Direct Reference, Propositional Attitudes and Semantic Content’, *Philosophical Topics* **15**, 47 – 87.

Soames, S. (2010). ‘What is Meaning?’, Princeton University Press.

Stanley, J. (2000). ‘Context and logical form’, *Linguistics and Philosophy* **23**, 391 – 434.

Stanley, J. (2002). ‘Nominal Restrictions’, in G. Peter and G. Preyer (eds.), ‘Logical Form and Language’, Oxford University Press 2002, 365 – 390.

Stanley, J. (2005). ‘Knowledge and Practical Interests’, Oxford University Press.

Szabó, Z. G. (2006). ‘Sensitivity Training’, *Mind and Language* **21**, 31 – 38.

Unger, P. (2002). ‘Philosophical Relativity’, Oxford University Press.

Wheeler, S. (1972). ‘Attributives and their modifiers’, *Noûs* **6**, 310 – 334.