

The myth of occurrence-based semantics*

Bryan Pickel and Brian Rabern

University of Edinburgh

- forthcoming in *Linguistics and Philosophy* -

Abstract

The principle of compositionality requires that the meaning of a complex expression remains the same after substitution of synonymous expressions. Alleged counterexamples to compositionality seem to force a theoretical choice: either apparent synonyms are not synonyms or synonyms do not syntactically occur where they appear to occur. Some theorists have instead looked to Frege's doctrine of "reference shift" according to which the meaning of an expression is sensitive to its linguistic context. This doctrine is alleged to retain the relevant claims about synonymy and substitution while respecting the compositionality principle. Thus, Salmon (2006) and Glanzberg and King (2020) offer occurrence-based accounts of variable binding, and Pagin and Westerståhl (2010c) argue that an occurrence-based semantics delivers a compositional account of quotation. Our thesis is this: the occurrence-based strategies resolve the apparent failures of substitutivity in the same general way as the standard expression-based semantics do. So it is a myth that a Frege-inspired occurrence-based semantics affords a genuine alternative strategy.

The principle of compositionality is a cornerstone of contemporary formal semantics. A semantic theory is compositional for a language when the meaning of a complex expression is determined by the meanings of the simpler expressions from which it is syntactically derived. Minimally, this requires that the meaning of a complex remains the same after substituting synonymous expressions while leaving all else the same. That is:

*For comments and discussion, thanks to Micheal Glanzberg, Nathan Klinedinst, Peter Pagin, and two anonymous referees for *Linguistics and Philosophy*. Earlier versions of this paper were presented at the *Context and variables* workshop at All Souls College in Oxford (2017), the *What is Said–What is Meant* workshop in Berlin (2016), and the 2016 Mid-Atlantic Philosophy of Language Workshop in Morgantown – the paper benefited from the questions and objections at these events. It would have, no doubt, benefited from Glanzberg's comments at the 2020 Pacific APA, but this conference was cancelled due to the pandemic.

COMPOSITIONALITY: If ϕ_β differs from ϕ_α only by substitution of β for α , and β and α are synonymous, then ϕ_α and ϕ_β are synonymous.

Semanticists have been largely successful in assimilating problematic cases into compositional semantic theories (see, e.g., Portner and Partee 2008).

Yet, certain recalcitrant cases remain, e.g. variable or pronoun binding, quotation, attitude reports, etc. In these cases, we appear to have non-synonymous expressions ϕ and ψ that differ only by the substitution of synonymous expressions α and β . Proponents of compositional semantics are faced with a standard diet of options in response to these apparent violations of the principle of compositionality.

DENY SYNONYMY: While the complex expressions ϕ and ψ do differ only by the substitution of α and β , and perhaps there is some dimension of meaning along which α and β are synonymous, there is actually another dimension of meaning according to which α and β are not synonymous.

DENY SUBSTITUTION: While α and β are synonymous, and perhaps it appears that ϕ and ψ are syntactically related by the mere substitution of α and β , the expressions ϕ and ψ do not actually differ only by the substitution of α and β .

Some find that this strict diet does not satisfy. The apparent counterexamples to compositionality force a theoretical choice: *either* the apparent synonyms are not synonyms *or* the synonyms do not syntactically occur where they appear to occur. Those with wider tastes look to an earlier tradition stemming from Gottlob Frege for dealing with apparent counter-examples to the compositionality of what he called *Bedeutung*, or reference.

According to Frege, “the truth value of a sentence remains unchanged when an expression is replaced by another having the same reference” (Frege 1892/1952: 64). But Frege observed that apparently co-referential terms such as ‘the capital of Scotland’ and ‘Edinburgh’ do not substitute *salva veritate* into various linguistic contexts. Standard examples include (i) quotation contexts, (ii) propositional attitude ascriptions, and (iii) modal contexts. The sentences in the pairs (1), (2), and (3) differ by the substitution of ‘the capital of Scotland’ for ‘Edinburgh’, but the first sentence in each pair may be true while the second sentence in each pair may be false.

- (1) ‘Edinburgh’ begins with the fifth letter of the English alphabet.
‘The capital of Scotland’ begins with the fifth letter of the English alphabet.

- (2) Pierre said that Edinburgh is beautiful.
Pierre said that the capital of Scotland is beautiful.¹
- (3) Edinburgh might not be Edinburgh.
The capital of Scotland might not be Edinburgh.

Frege proposed a radical response. Though the expressions co-refer, they may also “shift” their reference in embedded contexts. In the case of quotation:

... a word standing between quotation marks must not be taken as having its ordinary reference. (Frege 1892/1952: 58-59)

In the case of propositional attitude ascriptions such as reported indirect speech or modals, the reference is the customary sense or cognitive significance of the embedded expression.

In reported speech, words are used indirectly or have their indirect reference. We distinguish accordingly the customary from the indirect reference of a words; and its customary sense from its indirect sense. The indirect reference of a word is accordingly its customary sense. (Frege 1892/1952: 59)

So according to Frege, the referent of an expression is sensitive to the linguistic environment it is embedded in. Occurrences of the expressions ‘Edinburgh’ and ‘the capital of Scotland’ co-refer when they are unembedded. But, the expressions acquire different referents when they are embedded in quotation, propositional attitude ascriptions, or modal contexts.²

Some have seen in Frege’s remarks, a doctrine of *occurrence-based semantics* which provides an attractive alternative to the standard diet of options.³ According to this doctrine, an expression has meaning through its *occurrences*. And different occurrences of the same expression type may differ in meaning. An occurrence of an expression in a

¹Pierre may believe that the capital of Scotland is Glasgow.

²Frege’s 1879 theory of identity in *Begriffsschrift* could also be understood as an occurrence-based view (as an anonymous referee suggested). For example when Frege says “[signs] suddenly display their own selves as soon as they are combined by the sign for identity” (Frege 1879/1967: §8), this sounds highly reminiscent of what Frege later says about quotational contexts in Frege (1892/1952). See a recent discussion and reconstruction of Frege’s early view of identity in Pardey and Wehmeier (2019).

³In addition to Frege’s doctrine of indirect reference Salmon (2006) also relates occurrence-based semantics to Frege’s infamous *context principle*. Yet, the context principle is primarily put to use in the *Grundlagen* in relation to the “contextual definition” of number, and not as a general principle guiding the semantics of indirect contexts.

complex expression is an expression type paired with a “linguistic context” in the complex expression, where a *linguistic context* is essentially a particular position within the syntactic structure.⁴

It is thought that this doctrine can retain the relevant instances of *SYNONYMY* but relativized to an occurrence: unembedded occurrences of ‘Edinburgh’ and ‘the capital of Scotland’ may be genuinely synonymous. But the meaning of these expressions systematically shifts when they are embedded in various linguistic contexts. Because the meaning of an expression is relativized to a linguistic context, proponents of occurrence-based semantics must reformulate an occurrence-relative compositionality principle such as the following.⁵

OCCURRENCE-BASED COMPOSITIONALITY: If ϕ_β differs from ϕ_α only by substitution of β for α in linguistic context c , and β and α are synonymous in c , then ϕ_α and ϕ_β are synonymous (in an unembedded context).

It is thought that this alternative compositionality principle provides a systematic enough foundation for semantics, and avoids the choice forced by the strict diet of standard semantics. Proponents of occurrence-based semantics promise liberation from the constraints of standard semantics. Thus, Glanzberg and King (2020) defend a classical picture of semantics, whereby the semantic value of a sentence is a proposition, against a threat from the compositional semantics of binding—they do so by positing various meaning adjustments that are “triggered by syntactic environments”. Salmon (2006), similarly, offers an occurrence-based alternative for variable-binding, which promises to avoid unpalatable consequences of the standard approaches. And Pagin and Westerståhl (2010c) argue that—in contrast to the standard semantics—occurrence-based semantics delivers a compositional (in the occurrence-based sense) and *straightforward* account of quotation

⁴Pagin and Westerståhl gloss the basic idea as follows: “take a well-formed complex term $s[u]$ and knock out the constituent u . What remains, $s[\dots]$ is the linguistic context of that occurrence of u in $s[u]$, the environment of the argument place” (Pagin and Westerståhl 2010c: 394). They formally model this as a sequence of derivation rules encoding the path in the derivation tree that starts from the root and ends with the relevant node (Pagin and Westerståhl 2010c, §4.2). There are various choices for formally modeling linguistic context depending on the specific syntactic framework, but since our key point is general, we needn’t make a strong commitment. For further discussion of linguistic contexts (and occurrences of expressions therein) see Pagin and Westerståhl (2010a: §3.8), Glanzberg and King (2020: 15-19), and Salmon (2006: 416-20).

⁵The principle of occurrence-based compositionality is essentially an informal statement of what Pagin and Westerståhl define as LC-Subst(μ, C) (2010a: 400), and if rendered substitutionally Glanzberg and King’s *Strong Type Adjusting Compositionality* is getting at the same idea (Glanzberg and King 2020: 18). See also Salmon (2006: 430) and Pagin and Westerståhl (2010a: 260).

allowing for cases of synonymous expressions embedded in not-synonymous quotational constructions.

But, unfortunately, occurrence-based semantics does not deliver the results it advertises. We will argue that full articulation of the theoretical commitments of occurrence-based semantics demonstrates that the approach does not present a genuine alternative to the standard diet of options. In fact, it collapses into one of the the standard strategies for resolving apparent violations of compositionality; DENY SYNONYMY OF DENY SUBSTITUTION.

In §1, we examine how to best develop Frege’s observations into a systematic occurrence-based semantics. In §2, we examine Pagin and Westerståhl’s occurrence-based semantics as developed in their theory of quotation, arguing that it collapses into one of the standard responses and thus does not meet their own desiderata. In §3, we show that the occurrence-based semantics for variable binding in first-order logic developed in Salmon (2006) is not a genuine alternative, and that for similar reasons, the type-adjusting approach provided by Glanzberg and King (2020) for the semantics of natural language binding also collapses into one of the standard options. We conclude with a methodological lesson.⁶

1 Fregean reference shifts and occurrence-based semantics

Above, we described three pairs (1), (2), and (3) which differ merely by the substitution of the co-referential terms ‘the capital of Scotland’ and ‘Edinburgh’. The first sentence in each pair may be true while the second, false. Frege explained substitution failures of these types by arguing that the particular uses (in the relevant sentences) of ‘the capital of Scotland’ and ‘Edinburgh’, in fact, have different referents. The customary referents of ‘the capital of Scotland’ and ‘Edinburgh’ are the same city. But as they occur embedded under quotation marks, propositional attitude ascriptions, and modal operators, respectively, Frege thinks that these expressions have distinct referents. Consider again the pair (1):

- (1) ‘Edinburgh’ begins with the fifth letter of the English alphabet.
 ‘The capital of Scotland’ begins with the fifth letter of the English alphabet.

According to Frege, the uses of ‘the capital of Scotland’ and ‘Edinburgh’ in (1) refer to themselves, and thus such uses are not co-referential. Frege tells an analogous story about

⁶The lessons here apply more generally to the “switcher semantics” developed in, e.g., of Pagin and Glüer (2006), Pagin (2019), and Packalén (2016).

propositional attitude and modal contexts, whereby in such contexts words refer to their customary senses. In this way Frege insists that an expression can have different referents in its different occurrences in different linguistic contexts.

Proponents of occurrence-based semantics believe this Fregean idea can be implemented so as to expand our standard diet of options concerning apparent failures of compositionality. We believe otherwise. Following Kaplan (1964: 22-23) we distinguish two methods for implementing Frege's proposal.⁷ There are those employing *the method of direct discourse* and those employing *the method of indirect discourse*. The method of direct discourse uses distinct expressions in the formalism in order to explain apparent substitution failures such as (1). The method of indirect discourse uses the same expression, but offers more complicated semantic rules which allow the expression to make different truth conditional contributions in its different occurrences, to explain the substitution failures (Kaplan 1964: 27-8). Properly articulating each approach requires some care. But doing so reveals that neither approach expands the standard diet of options.

1.1 The method of direct discourse

The method of direct discourse corresponds to an utterly banal way in which an expression might have different meanings in different linguistic contexts and thereby lead to substitution failures. The expression 'bank' might be used as a synonym for 'financial institution'. But now consider the fact that the first sentence of the pair (4) may be true while the second sentence is false. Thus, (4) is an apparent violation of the principle of compositionality.

- (4) I parked my boat on the bank.
 I parked my boat on the financial institution.

In particular, the first sentence may be true without the second sentence being true because 'bank' may be used to mean something other than *financial institution*, viz. the land that lines the river. That is, 'bank' is *ambiguous*. It has two meanings. In a properly regimented language, it would be regimented by two different expressions. On the relevant readings of the sentences in (4), they do *not* differ by the substitution of synonymous expressions. So, the sentences in (4) are not a genuine violation of the principle of compositionality.

⁷See also the discussion in Klement (2002: chapter 4).

The method of direct discourse explains the differences in truth-value of the sentences in the pairs (1), (2), and (3) just as we have explained the difference in truth-value of the sentences in (4). According to Kaplan, “[t]his then is Frege’s treatment of intermediate contexts—obliquity indicates ambiguity” (Kaplan 1968: 184). To analyse the ambiguity, we introduce distinct symbols with different meanings, just as we might differentiate the two uses of ‘bank’ by writing a subscript, e.g. ‘bank₁’ and ‘bank₂’. As Kaplan (*ibid.*) says, Frege’s approach has the effect of “indicating (some would say, postulating) ambiguities where others have seen only opacity.” Positing ambiguity when faced with an apparent compositionality failure is, of course, among our standard options, namely DENY SUBSTITUTION. Thus this would seem to foreclose the possibility of a novel occurrence-based semantics arising from Frege’s insight. If occurrence-based semantics is meant to provide a genuine alternative to the standard diet of options, it can’t be derived from this interpretation of the Fregean approach.⁸

Some of Frege’s own remarks in a letter to Russell do suggest the method of direct discourse.

To avoid ambiguity, we ought really to have special signs in indirect speech, though their connection with the corresponding signs in direct speech should be easy to recognize. (Frege 1902/1980: 153)

Because he treats a word such as ‘Edinburgh’ as referring to different things depending on whether it is embedded, Frege sees the word as ambiguous across these uses. The ambiguity is to be resolved in our formalism by introducing different symbols corresponding to each use.⁹

But a subtle reading of this passage raises questions about Frege’s intent. Having observed that Frege suggests that the terms embedded in indirect speech and those in direct speech should have a connection, Kripke (2008: 190) asks “are they independent symbolic expressions or are they not?” Kripke continues:

⁸See Goodman and Lederman (2019) for a recent discussion and criticism of Frege’s ambiguity strategy. And see Pickel and Rabern (in review) for a reply.

⁹Frege doesn’t actually implement this idea (nor any other way of treating indirect contexts) into his official formalism, but in the letter to Russell he uses the device of underlining to distinguish customary and indirect reference. Church’s formalization of Frege explicitly adopts the method of direct discourse, where he uses subscripts to disambiguate (see his *Logic of Sense and Denotation* in Church 1951). See also *Introduction to Mathematical Logic* where Church speaks of “eliminating the oblique uses of names by introducing special names to denote the senses which other names express” (Church 1956: 8). See also Klement (2002: chapter 4).

And of course though 'London' primarily refers to a great city in England, it is also the name of a city in Ontario (Canada), and for all I know other cities. [. . .] According to Kaplan, Frege's theory of direct and indirect quotation . . . is that the ambiguity in these contexts is similar. In direct and indirect quotation contexts phrases do not have their usual denotation. So far so good, and no doubt this *is* Frege's theory. [footnote 29] But it is significant that these are not ordinary ambiguities or unusual references unlike . . . the 'London' example. For in Frege's theory of indirect quotation and the like, the unusual references are determined entirely from the context. In the other examples the unusual reference can often be *guessed* from the context, but it is not determined by any semantic rule. (Kripke 2008: 189-90)

Thus according to Kripke, Frege's theory should not be assimilated to ordinary ambiguity—instead it should be understood as something more systematic and rule-governed. Those who are inclined to emphasize this difference between ordinary ambiguities and Fregean "reference shift" must develop a semantics which can systematically predict the *shifty* behaviour of an expression across various linguistic contexts.

1.2 The method of indirect discourse

If the Fregean doctrine of "reference shift" points to something more subtle and systematic how should it be understood? The method of indirect discourse points the way forward. Rather than semantically analyzing the different contributions of an expression in its different occurrences by introducing two expressions which univocally make these different contributions, the method of indirect discourse complicates the semantics for the single expression.

A heavy handed way of complicating the semantics would be to introduce different interpretation functions that are operative for different occurrences. Consider, for example, the occurrences of 'Edinburgh' in (5) and (6).

- (5) 'Edinburgh' refers to Edinburgh.
- (6) "Edinburgh" refers to 'Edinburgh'

One would like to say that 'Edinburgh' occurs four times in these two sentences: once unquoted, twice in single quotes, and once in double quotes. One might introduce three

interpretation functions: $\llbracket \cdot \rrbracket_{\$}$ which is operative in unquoted contexts, $\llbracket \cdot \rrbracket_{\mathcal{E}}$ which is operative in single quoted contexts, and $\llbracket \cdot \rrbracket_{\mathcal{E}}$ which is operative in double quoted contexts. On such a view, the single expression ‘Edinburgh’ can make three different truth conditional contributions in its different occurrences. It can contribute $\llbracket \text{Edinburgh} \rrbracket_{\$}$, which is a city; or $\llbracket \text{Edinburgh} \rrbracket_{\mathcal{E}}$, which is a name of a city, or $\llbracket \text{Edinburgh} \rrbracket_{\mathcal{E}}$, which is the quote name of the name of a city.

One might hope that this sort of approach—which seems to make good on Frege’s claim that uses of names under (single) quotation marks contribute themselves—provides the promising occurrence-based alternative to the standard diet of options. But we deny that it provides such an alternative. In particular, insofar as the approach explains substitution failures, it does so by denying that the substituted expressions are synonymous.

To see this consider the fact that standard cases of ambiguity such as ‘bank’ can also be treated using different interpretation functions. So, rather than introducing distinct expressions to explain the ambiguity in ‘bank’, the method of indirect discourse would introduce two interpretation functions $\llbracket \cdot \rrbracket_{\text{R}}$ and $\llbracket \cdot \rrbracket_{\text{F}}$. Each interpretation function gives a disambiguation of ‘bank’, where $\llbracket \text{bank} \rrbracket_{\text{F}}$ is a financial institution but $\llbracket \text{bank} \rrbracket_{\text{R}}$ is the lining of a river. Now consider again the apparent substitution failure in (4). Insofar as the sentences ‘I parked my boat on the bank’ and ‘I parked my boat on the financial institution’ differ in truth-value, it is because they are being evaluated with respect to the interpretation function $\llbracket \cdot \rrbracket_{\text{R}}$.

$$\llbracket \text{I parked my boat on the bank} \rrbracket_{\text{R}} \neq \llbracket \text{I parked my boat on the financial institution} \rrbracket_{\text{R}}$$

But this interpretation function distinguishes these two sentences only because it yields different meanings for the parts. Namely, $\llbracket \text{bank} \rrbracket_{\text{R}} \neq \llbracket \text{financial institution} \rrbracket_{\text{R}}$. Thus, insofar as the sentences in (4) differ in truth-value, they do so because their components differ in meaning.

The comparison with ambiguity brings out another problem. As formulated thus far, the variation in meanings is bound to strike one as unsystematic. It is not governed by *semantic rules*. There doesn’t seem to be any interesting relation between interpretations such as $\llbracket \cdot \rrbracket_{\$}$, which is operative in unquoted contexts, and $\llbracket \cdot \rrbracket_{\mathcal{E}}$, which is operative in single quoted contexts. This brings the same sorts of ills as ambiguity. For instance, without some sort of systematic relation among the meaning functions, one will seemingly have to learn them separately just as one learns the different lexical entries corresponding to ‘bank’ separately. The fact that these different meaning functions apply in different

linguistic contexts doesn't help distinguish this treatment from an ambiguity analysis, since the different lexical entries corresponding to an ambiguous word such as 'can' will be determined by whether it occurs as a common noun or modal verb. The problem is exacerbated when one reflects upon the fact that quotation is recursive: one can always quote a quoted expression. So we need infinitely many interpretation functions, one for each level of quotation embedding. (Similar remarks may apply to analogous treatments of variable binding and attitude embeddings.)

Instead one wants to say that each expression is governed by a *single* semantical rule that allows it to make different truth conditional contributions in different linguistic contexts. The method of indirect discourse requires articulating this semantic rule for an expression. For example, on such an approach, the expression 'Edinburgh' is governed by a single linguistic rule determining that the expression refers to Edinburgh while unembedded and to the expression itself when embedded in a single quotation marks. On this implementation of the method of indirect discourse, the meaning of 'Edinburgh' has then been *intensionalized*. Rather than being a city, the meaning of 'Edinburgh' is given by a function from linguistic contexts to its values in these contexts. That is, the meaning of an expression will be given by a single interpretation function $\llbracket \cdot \rrbracket$. Given an expression such as 'Edinburgh', the meaning $\llbracket \text{Edinburgh} \rrbracket$ will itself be a function from linguistic contexts into occurrence meanings. In certain linguistic contexts c , $\llbracket \text{Edinburgh} \rrbracket(c)$ will be a city. But in another context c^* , $\llbracket \text{Edinburgh} \rrbracket(c^*)$ will be the expression 'Edinburgh'.

This version of the method of indirect discourse does not expand the options for responding to substitution failures either. The failure of 'Edinburgh' and 'the capital of Scotland' to substitute in (1) would be explained by the fact that these expressions themselves have different meanings. It is therefore an application of DENY SYNONYMY.

1.3 Occurrence-based semantics

Advocates of occurrence-based semantics nevertheless insist that Frege's remarks point the way to a novel solution to problems of substitution. They do so by insisting that occurrences—expressions in linguistic contexts—are the fundamental bearers of meaning. One way to implement this would be to frame the semantic axioms using a two-placed interpretation function (Pagin and Westerståhl 2010c). The function takes both an expression and a linguistic context as arguments and outputs an occurrence meaning.

$$\llbracket \cdot \rrbracket: (\text{EXPRESSIONS} \times \text{LINGUISTIC CONTEXTS}) \longrightarrow \text{OCCURENCE MEANINGS}$$

An expression such as ‘Edinburgh’ will have a single lexical entry which for any linguistic context c assigns a meaning to that expression in context.

Advocates of occurrence-based semantics insist that this approach carries an important pay-off. There is no such thing as the semantic value of ‘Edinburgh’, \llbracket Edinburgh \rrbracket , or the semantic value of ‘the capital of Scotland’, \llbracket the capital of Scotland \rrbracket . There are only the meanings of these expressions at some particular linguistic context c : \llbracket Edinburgh \rrbracket^c and \llbracket the capital of Scotland \rrbracket^c . ‘Edinburgh’ and ‘the capital of Scotland’ are synonymous as they occur in contexts such as ‘Edinburgh is rainy’ and ‘the capital of Scotland is rainy’, but not in the context of quotation such as those in (1).

This approach is underspecified at a crucial point: what—if anything—is the semantic value of an expression type. On one view, each expression-context pair has its own lexical entry. On this view, ‘Edinburgh’ as it occurs in ‘Edinburgh is the capital of Scotland’ has a different lexical entry from ‘Edinburgh’ as it occurs in “‘Edinburgh’ begins with the fifth letter of the English alphabet”. On the alternative view, these various meanings in context are subsumed under a more general lexical entry for the expression type. On this view, there is a single semantic value for ‘Edinburgh’ that determines its occurrence meanings across the various contexts.

Unfortunately for the occurrence-based semanticist, neither option—viewing the meanings of expressions in context as basic or as subsumed under a larger semantics for expression types—expands the range of possible replies to puzzles arising from substitution. Taking each expression in context to have its own semantic axiom is implausibly unsystematic. It again raises the spectre of ambiguity. On the other hand, taking the various occurrence meanings to follow from the meaning of the expression itself simply intensionalizes the meaning and therefore is a version of DENY SYNONYMY.

In what follows we will examine occurrence-based responses to two substitution puzzles. These responses purport to be outside of the standard diet options. However, we show that once the proposals are made precise they fall in line with options from the standard diet. That is, they are ultimately variants of DENY SUBSTITUTION OR DENY SYNONYMY.

2 Pure quotation

As we have intimated, a kind of stress test for compositionality is the problematic case of quotation—what happens when language is turned upon itself? Assume that the name

‘Edinburgh’ occurs as a syntactic constituent of *this* sentence (the one you are currently reading), as the orthography suggests. Thus, the application of quote marks is a genuine syntactic operation—that is, a quote phrase is a complex expression derived from applying quotes to an expression.¹⁰ Let’s represent an expression α embedded under quotes in the object language as follows: *quote* α . Quoted expressions refer to expressions and thus the semantics should yield the following:

$\llbracket \textit{quote Edinburgh} \rrbracket = \text{‘Edinburgh’}$

$\llbracket \textit{quote the capital of Scotland} \rrbracket = \text{‘the capital of Scotland’}$

The problem with compositionality is that ‘Edinburgh’ and ‘the capital of Scotland’ co-refer, that is $\llbracket \text{Edinburgh} \rrbracket = \llbracket \text{the capital of Scotland} \rrbracket$. By the principle of compositionality, if two expressions such as ‘*quote* Edinburgh’ and ‘*quote* the capital of Scotland’ differ only by the substitution of synonymous terms, then the complex expressions themselves are synonymous. It would follow that $\llbracket \textit{quote Edinburgh} \rrbracket = \llbracket \textit{quote The capital of Scotland} \rrbracket$, which is false.

The problem here is not simply due to the assumption that the semantic value of a name or definite description is its referent, the problem generalizes. Take any distinct but synonymous expressions α and β . By assumption the expressions have the same semantic value, $\llbracket \alpha \rrbracket = \llbracket \beta \rrbracket$, but since they are distinct expressions their quotations must differ in value: $\llbracket \textit{quote } \alpha \rrbracket \neq \llbracket \textit{quote } \beta \rrbracket$. Yet by compositionality, $\llbracket \textit{quote } \alpha \rrbracket = \llbracket \textit{quote } \beta \rrbracket$, since these expressions differ only by the substitution of synonyms. Thus, if quotation is a genuine syntactic operation, and there are distinct but synonymous expression types in the language, then quotation is not compositional. The obvious thing to do at this point is call into question some of the assumptions that lead to the failure of compositionality.

The standard diet of options provides two choices: DENY SYNONYMY OR DENY SUBSTITUTION. To deny DENY SYNONYMY one would deny that the purported dimension of meaning whereby $\llbracket \alpha \rrbracket = \llbracket \beta \rrbracket$ exhausts the meaning of α and β —there is a richer notion of meaning according to which α and β differ. But since this claim will have to be generalized for any pair of simple expressions α and β this amounts to claiming that there are no two lexically distinct (quotable) expressions α and β which have the same meaning. Thus, the semantic

¹⁰The key assumption here is that there is a syntactic rule which when applied to an expression α yields the quote name of α . So alternatively we could talk about the derivation rule QUOTE and the result of applying the rule to an expression QUOTE(α). Compositionality is threatened, since it requires a rule that maps the semantic value of α to the semantic value of QUOTE(α).

theory will be compositional but only trivially so—the antecedent of COMPOSITIONALITY is false for any distinct α and β .

Due to the absurdly fine-grained meanings such a strategy must posit and the resulting trivialization of compositionality a more popular approach traditionally has been to deny that quotation is a syntactic operation, in others words deny that the word ‘Edinburgh’ occurs as a syntactic constituent of the following sentence:

‘Edinburgh’ begins with the fifth letter of the English alphabet.

In terms of our standard diet this strategy is to DENY SUBSTITUTION. This type of view has been endorsed in one form or another by many theorists, e.g. it was endorsed by Tarski 1936/1956 (cf. Quine 1940: 26), when he said that expressions surrounded by quotes

... may be treated like single words of a language, and thus like syntactically simple expressions. The single constituents of these names—the quotation marks and the expressions standing between them—fulfil the same function as letters and complexes of successive letters in single words. Hence they can possess no independent meaning. (159)

While a popular view, theorists have been dissatisfied with the idea that quotations are syntactically simple expressions, e.g. Davidson complains as follows:

If quotations are structureless singular terms, then there is no more significance to the category of quotation mark names than to the category of names that begin and end with the letter ‘a’ (‘Atlanta’, ‘Alabama’, ‘Alta’, ‘Athena’, etc.). On this view, there is no relation, beyond an accident of spelling, between an expression and the quotation mark name of that expression. (Davidson 1979: 30)

It seems that an approach which accounts for the non-accidental nature of these connections would be preferable.¹¹

¹¹Davidson famously put forward his own alternative, which agrees that ‘Edinburgh’ does not occur as a syntactic constituent of “ ‘Edinburgh’”, whereby quotation marks are demonstratives (roughly equivalent to ‘that expression’): “The singular term is the quotation marks, which may be read ‘the expression a token of which is here’ [...] the inscription inside does not refer to anything at all, nor is it part of any expression that does. Rather it is the quotation marks that do all the referring, and they help to refer to a shape by pointing out something that has it.” (Davidson 1979: 90). Note that one could agree with Davidson that quotation marks

As we have said, Pagin and Westerståhl promise that occurrence-based semantics provides an attractive alternative. In particular, they argue one can offer what they call a “straightforward” semantics for quotation without having to DENY SUBSTITUTION OF DENY SYNONYMY. They call these desiderata (a) and (b), respectively.¹²

A straightforward account of (the use of quote marks in) pure quotation is one which (a) takes the quoted phrase to be a syntactic constituent of the quoting phrase (the quoted phrase surrounded by quote marks), and (b) allows at least one case of two syntactically distinct and quotable expressions having the same semantic interpretation (meaning). (Pagin and Westerståhl 2010c: 382)

Pagin and Westerståhl (2010b) reformulate the interpretation relation $\llbracket \cdot \rrbracket$ “so that it takes a context type as additional argument, where we have one default context type for ordinary use, and a distinct one for quotation, and similarly for the semantic operations ” (275).¹³

Pagin and Westerståhl model a linguistic context as a sequence of derivation rules encoding the path in the derivation tree that starts from the root and ends with the relevant node. So, for example, consider these three expressions:

- (i) Edinburgh
- (ii) ‘Edinburgh’
- (iii) “ ‘Edinburgh’ ”

In (i)-(iii) the name ‘Edinburgh’ occurs in three different linguistic contexts. Where QUOTE is the derivation rule for pure quotation it first occurs in (i) unquoted, the null context $\langle \rangle$, and then in (ii) quoted in the context $\langle \text{QUOTE} \rangle$, and finally in (iii) it occurs in quotes within quotes, so the context $\langle \langle \text{QUOTE} \rangle, \langle \text{QUOTE} \rangle \rangle$. According to Pagin and Westerståhl’s occurrence-based approach the three occurrences of ‘Edinburgh’ each have a different semantic value. The occurrence in $\langle \rangle$ has a city as its value, the occurrence in $\langle \text{QUOTE} \rangle$ has

are demonstratives (‘the expression which is my syntactic sister’) and insist that the expression ‘Edinburgh’ occurs as a syntactic sister to the quotes. This contextualist account is worth considering. But we don’t take any official stand on the syntax and semantics of quotation in this paper, since our aim is more metatheoretical. For discussion of the various issues involved see Cappelen and Lepore (2007).

¹²Pagin and Westerståhl explain that “(b) is the very weak requirement that one’s notion of meaning allows some non-trivial synonymies” (382, footnote 1).

¹³In Pagin and Westerståhl (2010b), they develop a related proposal which makes use of multiple interpretation functions, a simple version of which we rejected in §1.2. But Pagin and Westerståhl (2010c) develop the proposal further, making use of the strategy of relativizing $\llbracket \cdot \rrbracket$ to context types.

the name of a city as its value, and the occurrence in $\langle\langle\text{QUOTE}\rangle, \langle\text{QUOTE}\rangle\rangle$ has the name of a name of a city as its value.

Since here we are only concerned with quotational linguistic contexts we will model $\langle\rangle, \langle\text{QUOTE}\rangle, \langle\langle\text{QUOTE}\rangle, \langle\text{QUOTE}\rangle\rangle$, etc., as $0, 1, 2, \dots$, respectively. So to present their proposal, let a number $n \geq 0$ indicate that an expression is embedded under n -many quotation operators. The relativized notion of compositionality says that each expression is assigned a meaning, in a linguistic context, and that the meaning of a complex expression (in context) is a function of the meanings of its constituents in their respective contexts. One may allow that two unquoted expressions agree in meaning, for instance that $\llbracket\text{Edinburgh}\rrbracket^0 = \llbracket\text{the capital of Scotland}\rrbracket^0$, while denying that enclosing them in quotes results in expressions that agree in meaning, $\llbracket\text{quote Edinburgh}\rrbracket^0 \neq \llbracket\text{quote the capital of Scotland}\rrbracket^0$. The reason is that a name in quotations must be evaluated in a different linguistic context from the name unquoted.¹⁴ So it is possible that $\llbracket\text{quote Edinburgh}\rrbracket^0 \neq \llbracket\text{quote the capital of Scotland}\rrbracket^0$, because the embedded expressions ‘Edinburgh’ and ‘the capital of Scotland’ must be evaluated at a different linguistic context. They are evaluated relative to the linguistic context of embedding under single quotation marks, modelled as 1. And, $\llbracket\text{Edinburgh}\rrbracket^1 \neq \llbracket\text{the capital of Scotland}\rrbracket^1$.

Pagin and Westerståhl advertise their view as a weakening of the principle of compositionality which implements Frege’s occurrence-based semantics.

... allowing for linguistic context dependence ... we can obtain a semantics which is compositional in the generalized sense [...] The new semantics ... will be very close to adapting Frege’s view in Frege 1892 that quotation provides a context in which expressions refer to themselves. (Pagin and Westerståhl 2010b: 275)

Since compositionality applies at the level of occurrences, Pagin and Westerståhl (2010c) say that they can deliver a “straightforward” account of quotation that does not DENY SUBSTITUTION OR DENY SYNONYMY, but which is nonetheless systematic. We take this to mean that even if two expressions α and β have the same meanings in contexts with no quotation marks and these very expressions may occur inside of quotation marks, the result of encasing them in single quotation marks need not yield expressions with the same meaning.

¹⁴As Pagin and Westerståhl (2010c: 397) say: “The first main idea for an extension of a semantic framework to handle linguistic context dependence is that when a semantic function is applied to a term s , the value depends on the relevant context type. The second main idea is that the context type of an immediate subterm of a term s may be different from the context type of s itself.”

This account treats the interpretation function, $\llbracket \cdot \rrbracket$, as a two-placed relation that takes an expression and a linguistic context to yield a meaning. A *prima facie* reading of their view would have Pagin and Westerståhl (2010c) saying that the expressions ‘Edinburgh’ and ‘the capital of Scotland’ simply do not have meanings, only their *occurrences* do. On this *prima facie* reading, there is no single semantic axiom giving the meaning of an expression such as ‘Edinburgh’ in all contexts. Rather, there are a sequence of functions that give an expression various meanings: $\llbracket \cdot \rrbracket^0, \llbracket \cdot \rrbracket^1, \llbracket \cdot \rrbracket^2 \dots$. We have $\llbracket \text{Edinburgh} \rrbracket^0 = \llbracket \text{the capital of Scotland} \rrbracket^0$ but $\llbracket \text{Edinburgh} \rrbracket^1 \neq \llbracket \text{the capital of Scotland} \rrbracket^1$.

Pagin and Westerståhl (2010c: 386) insist that is possible to deliver an occurrence-based compositional account of quotation in which (a) expressions genuinely occur inside of quotations and (b) there are nontrivial synonyms.¹⁵ Unfortunately, while this naive reading of their account is compositional it is not *straightforward* since it does not respect clause (b). That is, it does not allow at least one case of two syntactically distinct and quotable *expressions* (expressions!) to have the same meaning. Given this, the account is not a genuine alternative to the standard diet of options, since the proposal collapses into DENY SYNONYMY, and in fact it ends up denying that there are any synonymous expression pairs whatsoever.

The deeper problem is that this *prima facie* reading simply revisits the ambiguity approach considered above in §1.2. There we considered approaches that introduce multiple interpretation functions $\llbracket \cdot \rrbracket_{\$}$, $\llbracket \cdot \rrbracket_{\mathcal{E}}$, and $\llbracket \cdot \rrbracket_{\in}$ to explain substitution failures. The problem with these approaches was that interpretation functions would have to be individually learned rather than subsumed under a general semantic axiom giving the meaning of the expression. This *prima facie* reading of Pagin and Westerståhl (2010c)—on which the various interpretation functions $\llbracket \cdot \rrbracket^0, \llbracket \cdot \rrbracket^1, \llbracket \cdot \rrbracket^2 \dots$ —are not subsumed under a single semantic axiom suffers the same ills. The various meanings of an expression such as ‘Edinburgh’ are connected only by the fact that they are assigned to the word ‘Edinburgh’ just as the various meanings of ‘bank’ are connected only by being assigned to the word ‘bank’.¹⁶

As before, what is wanted is a single semantic rule assigning a meaning to an expression

¹⁵“Our claim in this paper is that with a corresponding notion of linguistic context, the straightforward account of quotation can be made compositional, in the general, contextual sense” (*ibid.*).

¹⁶Pagin and Westerståhl (2010c) contrast their own view with the expressions-in-context (or ambiguity) approach (see their appendix 1). They insist: “While there are good reasons to let the lexicon distinguish two or more meanings of *bank*, for example, there are no good reasons why the lexicon should care about quotation” (409). They insist that their account has a single meaning assignment μ that evaluates an expression t in a context c , whereas the alternative ambiguity account has disparate lexical entries, which are not governed by a systematic rule.

such as ‘Edinburgh’ all at once, in all different linguistic contexts. But this single semantic rule would allow us to reconstruct a notion of expression meaning. The general *rule* would assign the expression an occurrence meaning relative to a linguistic context. That is, one can understand a novel sentence by understanding the expressions it contains and, in particular, by knowing how their meanings vary in linguistic context. Adapting Fregean terminology, Dummett calls this general semantic feature of an expression its *one common sense*.

The sense of a word cannot vary from context to context, but is a property of the word itself apart from any context: for it is by knowing the sense of the constituent words, independently of their occurrence in *this* sentence, that we understand the sentence. If the sense of a word varied from context to context, this would have to be according to some general rule, if we were to understand the sentence in which it occurs: and then this general rule would in reality constitute the *one* common sense which the word possessed. (Dummett 1973: 268)

Synonymous expressions, then, would agree in Dummettian sense.

Effectively, facts about expression synonymy can be extracted from facts about expression-in-context synonymy by requiring that synonymous expressions are synonymous at every context. This common semantic value of the expression—equivalent to *curry*($[[\cdot]]$) which we write as $[[\cdot]]_*$ —is a function which takes a linguistic context and outputs the meaning of occurrences of the expression in the context.

$$[[\cdot]]_*: \text{EXPRESSIONS} \longrightarrow (\text{LINGUISTIC CONTEXTS} \longrightarrow \text{OCCURRENCE MEANINGS})$$

But understood this way the approach does not differ from the standard diet of options. The two-place function $[[\cdot]]$ may determine that expressions such as ‘Edinburgh’ and ‘the capital of Scotland’ have the same meaning at one linguistic context, but different meanings at another linguistic context. This is simply because the expressions have semantic values, respectively $[[\text{Edinburgh}]]_*$ and $[[\text{the capital of Scotland}]]_*$, that differ *tout court*. But as an account of substitution failure, this amounts to DENY SYNONYMY, the meaning of an expression at a linguistic context does not exhaust its semantic contribution.

Indeed, Pagin and Westerståhl (2010c: 400) do reconstruct a notion of synonymy in exactly these terms, which they call “total synonymy”. Two expressions exhibit total synonymy just in case they are synonymous in every context. This is just to say that two

expressions α and β are totally synonymous when $\llbracket \alpha \rrbracket_* = \llbracket \beta \rrbracket_*$. But then it follows that there are no distinct, but totally synonymous expressions! Pagin and Westerståhl admit just this when they say

if there is a quotation context type in E , then there will be no non-trivial total synonymy pairs: a term s is totally synonymous only with itself. (*ibid.*: 400)

So Pagin and Westerståhl agree that their view allows a definition of expression synonymy which entails that there are no (non-trivial) *total* synonyms. This undermines their desiderata for a “straightforward” account of quotation: there is not “at least one case of two syntactically distinct and quotable expressions having the same semantic interpretation”.

Pagin and Westerståhl (2010c: 399-400) do offer weaker notions of expression synonymy. One notion *designated synonymy* stipulates a designated linguistic context, say 0, and holds that two expressions are “synonymous” just in case $\llbracket \alpha \rrbracket^0 = \llbracket \beta \rrbracket^0$.¹⁷ But this seems to simply stipulate that quotation calls for special treatment. In our minds, this undermines the point of the exercise. If all that is wanted is that the expressions agree semantically in the unquoted contexts, then the view could be implemented even by someone who has a notion of meaning assigned by the interpretation $\llbracket . \rrbracket_*$. Even though ‘Edinburgh’ and ‘the capital of Scotland’ disagree in meaning in the absolute sense, $\llbracket \text{Edinburgh} \rrbracket_* \neq \llbracket \text{the capital of Scotland} \rrbracket_*$, they agree in meaning relative to unquoted contexts. That is, $\llbracket \text{Edinburgh} \rrbracket_*(0) = \llbracket \text{the capital of Scotland} \rrbracket_*(0)$.

Consider an analogous position. In Kaplan’s (1989a) logic of demonstratives, the meaning of an expression is its character, a function from extralinguistic contexts to contents. Thus, the meaning of the first person pronoun ‘I’ is the agent of the context. On this view, ‘I’ has a different meaning from any proper name, since proper names have the same content in every context. However, one might put forward the view that the first person pronoun ‘I’ and one’s own name exhibit designated synonymy, since they agree in content whenever one utters them. Does this show that there is a deep semantic connection between the first person pronoun and one’s name? We think not. Rather, we take the example to illustrate that the notion of designated synonymy is not significant for the theory of meaning.

¹⁷The other notion, *use synonymy*, requires that the two expressions mean the same in any non-quotation context, which again is context 0 in our model.

The general lesson here is that in order to account for the unified behaviour of an expression—our ability to understand diverse sentences that contain it, the expression must be assigned semantic properties by the semantic theory. In the occurrence-based approaches, the relevant semantic properties of an expression are that it has various meanings in various linguistic contexts. But then it will be natural to ask whether two expressions are synonymous or not. If the expressions can occur in the same linguistic contexts and yet differ in value in some of those contexts, then they cannot be synonymous. Occurrence-based semantics, so construed, is not a genuine alternative to the standard options for diagnosing failures in compositionality.

3 Variable binding

Variable binding is another salient stress test for compositionality. The worry begins with the thought that a free variable contributes an object relative to an assignment. So two variables ‘ x ’ and ‘ y ’ might have the same semantic value relative to an assignment σ : $\llbracket x \rrbracket^\sigma = \llbracket y \rrbracket^\sigma$. It would follow that the variable contributes that object to sentences containing it, when assessed relative to the same assignment. Consider the pair (7).

- (7) (a) $x \neq x$
 (b) $x \neq y$

Assuming that $\llbracket x \rrbracket^\sigma = \llbracket y \rrbracket^\sigma$, it follows that $\llbracket (7a) \rrbracket^\sigma = \llbracket (7b) \rrbracket^\sigma$. But the result of substituting ‘ x ’ for ‘ y ’ or (7a) for (7b) does not always result in a sentence with the same truth-value relative to assignment σ . For example, consider the following:

- (8) (a) $\exists x x \neq x$
 (b) $\exists x x \neq y$

(8a) and (8b) differ by the substitution of ‘ x ’ for ‘ y ’ and of (7a) for (7b). Nonetheless, (8a) is false on any assignment, but (8b) is true on any assignment (provided there are at least two things). So $\llbracket (8a) \rrbracket^\sigma \neq \llbracket (8b) \rrbracket^\sigma$. Thus, in this case ‘ x ’ and ‘ y ’ are synonymous (in the relevant sense) and (8b) is the result of substituting ‘ y ’ for ‘ x ’ in (8a), but the complex expressions (8a) and (8b) are *not* synonymous (in the relevant sense). For these reasons Soames (2011) says bluntly: “Compositionality even fails for Tarskian quantification”

(129). This apparent compositionality failure threatens foundational presuppositions of formal semantics.

The standard diet of options as applies to this case as well. The DENY SYNONYMY strategy insists that even though $\llbracket x \rrbracket^\sigma = \llbracket y \rrbracket^\sigma$, the variables ‘ x ’ and ‘ y ’ are not fully synonymous, since their semantic values are more fine-grained. In particular, the expressions are taken to have as their semantic values functions from assignments into objects: $\llbracket x \rrbracket = \lambda\sigma.\sigma(x)$ and $\llbracket y \rrbracket = \lambda\sigma.\sigma(y)$. As a result, $\llbracket x \rrbracket \neq \llbracket y \rrbracket$. In this way one can provide a compositional semantics for quantification.¹⁸

The other prominent strategy is to DENY SUBSTITUTION. This strategy insists that the complex expressions ‘ $\exists x x \neq x$ ’ and ‘ $\exists x x \neq y$ ’ don’t in fact differ by the substitution of (the synonymous) expressions ‘ x ’ and ‘ y ’. One traditional version of this strategy would outright deny that the variables are constituents of the quantified formulae. For example, on Frege’s view “variables” are simply typographic parts of the disconnected quantifier sign serving only to link the quantifier to the open spaces in predicates.¹⁹ For Frege variables are not expressions of the formal language, and thus talk of variables having sense or being synonymous is simply incoherent.²⁰

But one need not go so far. In particular one might think that variables are genuine syntactic units and that ‘ x ’ and ‘ y ’ might be synonymous, but nevertheless DENY SUBSTITUTION. On this proposal, the synonymous expressions written as ‘ x ’ and ‘ y ’ in ‘ $x \neq y$ ’ must be distinguished from the expressions with the same shape which occur embedded in ‘ $\exists x x \neq x$ ’ and ‘ $\exists x x \neq y$ ’. To put the matter another way, a mark such as ‘ x ’ is *ambiguous*. The mark corresponds to distinct lexical entries. As Salmon (2006) characterizes this position, “the lower-case letter ‘ x ’ (*qua* variable) ambiguously represents two distinct expressions: ‘ x ’-bound and ‘ x ’-free.”²¹ So on this account the formulae ‘ $\exists x x \neq x$ ’ and ‘ $\exists x x \neq y$ ’ do not differ by the substitution of *synonymous* expressions (namely, ‘ x ’-free and ‘ y ’-free), but rather by the substitution of their non-synonymous homonyms (‘ x ’-bound and ‘ y ’-bound).²²

¹⁸See Janssen (1997) and so-called cylindrified models. See also Lewis (1970), and more recent discussion in Rabern (2012, 2013) and Yli-Vakkuri (2013).

¹⁹According to Frege, a quantifier sign ‘ $\exists y \dots y \dots$ ’ is saturated by a predicate ‘ $\forall x x = \dots$ ’ to form the quantified expression ‘ $\exists y \forall x x = y$ ’ (see Frege 1893/2013: §8). Likewise, for the Bourbaki collective who would render a formula such as ‘ $\exists x(Fx \wedge Gx)$ ’ as a variable-less structure consisting of symbols, gaps, and wires connecting symbols to gaps (see Bourbaki 1954).

²⁰See Wehmeier (2018) and Pickel and Rabern (manuscript) “Against Fregean Quantification” for discussion.

²¹Salmon does not endorse this position, instead he attributes it to Jason Stanley.

²²Kaplan (1989a: 489-90) explicitly adopts this view for free and bound pronouns of natural language.

Thus, given the standard options one may DENY SYNONYMY: the semantic values of the variables ‘ x ’ and ‘ y ’ and of the sentences (7a) for (7b) are in fact distinct (even though they agree on assignment σ). Or, one may DENY SUBSTITUTION: (8a) and (8b) do not differ by the substitution of synonyms. Faced with this choice some cast about for a further alternative in occurrence-based semantics. We now turn to examining the occurrence-based approaches developed in Salmon (2006) and Glanzberg and King (2020). We argue that neither present a genuine alternative. (Note that a careful reading of Salmon suggests that he would not disagree with this conclusion, whereas Glanzberg and King require their approach to not collapse into a standard approach given their theoretical motives.)

3.1 Salmon’s theory of bondage

Salmon (2006) self-consciously develops an occurrence-based semantics for variables in the tradition of Frege.

In developing an occurrence-based semantics of variable-binding, I take my cue from Frege’s theory of indirect (oblique, ungerade) contexts. [. . .] Quantifiers are variable-binding operators. Like ‘believes that’, variable-binding operators induce the variables they bind to undergo semantic shift, but a shift of a different sort from intensional or “indirect” (oblique) operators. (Salmon 2006: 420-421)

To outline Salmon’s picture we will follow his terminology and speak of an expression’s “designatum” instead of its meanings or semantic value, but nothing turns on this. Salmon distinguishes the designatum of an occurrence of an expression (at an assignment) unembedded under a variable-binder from the designatum of an occurrence of an expression embedded under a variable-binder (at an assignment).²³ Thus Salmon distinguishes the designatum of a free occurrence of a variable from the designatum of a bound occurrence of a variable.

In response to the apparent compositionality failure above, Salmon insists on the following: The free occurrences of variables ‘ x ’ and ‘ y ’ in (7b) designate the same individual relative to σ . But in (8b), the occurrences of the variables under the quantifier take on

Kaplan (1989b) also states that “pronouns are lexically ambiguous, having both an anaphoric and a demonstrative use” (572).

²³Here we will only focus on occurrences of variables, in particular, even though Salmon’s theory applies to all expressions generally (both simple and complex).

designata distinct from their free designata (and also from each other). That is why (8b) has a different designatum from (8a).

Salmon provides a systematic, rule-governed theory about the designata of a variable in its different occurrences. For Salmon, the basic semantics relativizes an expression to a linguistic context, which Salmon models as a sequence of variables corresponding to the sequence of quantifiers the expression is embedded under. Letting β be a variable, the designatum of β when unembedded (modelled as the empty sequence \emptyset) relative to an assignment σ is the value of the function σ as applied to the variable.

$$\llbracket \beta \rrbracket^{\emptyset, \sigma} = \sigma(\beta)$$

So the free occurrences of variables ' x ' and ' y ' might have the same designatum under some assignment σ . When a variable β is embedded under some variable-binding operator, such as the quantifier ' $\exists \alpha$ ', the variable shifts its designatum. Relative to the context of being embedded under a binder indexed with the variable α , the variable β designates a function that takes an individual i and outputs the value of a free occurrence of β at an assignment that differs from σ at most by assigning i to α .

$$\llbracket \beta \rrbracket^{\alpha, \sigma} = \lambda i. \llbracket \beta \rrbracket^{\emptyset, \sigma[\alpha/i]}$$

As a result, the occurrences of ' x ' and ' y ' embedded under the quantifier in (8b) designate $\lambda i. \sigma[x/i](x)$ and $\lambda i. \sigma[x/i](y)$, respectively. The former is the identity function $\lambda i. i$, while the latter is the constant function which takes a member of the domain and outputs $\sigma(y)$. These are different functions, so relative to σ and the linguistic context of being embedded under ' $\exists x$ ', the variables ' x ' and ' y ' have different designata. Salmon generalizes this treatment to a variable embedded under n -quantifiers (modeled as a sequence of n variables).²⁴

So far so good. This account is in working order. However, those looking to Salmon for an alternative to the standard diet of responses to apparent compositionality failures will be disappointed. Given that there is a linguistic context where distinct variables ' x ' and ' y ' have distinct designata, it follows that ' x ' and ' y ' have distinct (total) semantic values. So, much like the discussion in the previous section, this approach collapses into DENY SYNONYMY. For the record, we do not think of our stance as a challenge to anything Salmon says. Salmon himself says that the various meanings assigned by the occurrence-based semantics are in fact semantic values of the expressions themselves.

²⁴The recursive clause for the designata of a variable's occurrences, following (Salmon 2006: 427) may be stated as: $\llbracket \beta \rrbracket^{\langle \alpha_1, \dots, \alpha_n, \alpha_{n+1} \rangle, \sigma} = \lambda i. \llbracket \beta \rrbracket^{\langle \alpha_1, \dots, \alpha_n \rangle, \sigma[\alpha_{n+1}/i]}$.

The two approaches [occurrence-based and expression-based], though different, are not intrinsically in conflict. [...] Frege’s occurrence-based semantics in fact assigns semantic values both to expressions and their occurrences, even while honoring his Context Principle. His notions of customary designatum, indirect sense, doubly indirect designatum, and the like, are semantic values of the expression itself. The customary designatum is the designatum of the expression’s occurrences in “customary” settings, that is, its occurrences that are in extensional position and not within the scope of a variable-binding operator. (Salmon 2006: 432)

Thus for Salmon the semantic value of an *expression* seems to be the function that takes a linguistic context c and outputs its designatum in c . Then the result is in accord with our conclusion that the expressions ‘ x ’ and ‘ y ’ have different semantic values. So the proposal ultimately denies that these expressions are synonymous and thereby does not expand on the standard diet of responses.²⁵

3.2 Glanzberg and King on type adjusting

Glanzberg and King (2020) are concerned to defend a traditional account of the relationship between compositional semantics and the objects of assertion and belief—in primarily *natural* as opposed to *formal* languages.²⁶ The worry they are addressing starts from the observation that natural language contains expressions of the same syntactic type as complete sentences that can be non-vacuously embedded under quantifiers.²⁷ Consider

²⁵Given this one might think that the Salmon-style semantics—with its relativization to linguistic contexts—is unnecessarily complicated. If the concern is simply to provide a compositional treatment of the language of first-order logic, then why not just resort to the relatively simpler strategy of raising all the semantic types to be assignment sensitive (e.g. let $\llbracket x \rrbracket = \lambda\sigma.\sigma(x)$ and $\llbracket y \rrbracket = \lambda\sigma.\sigma(y)$)? We agree, but take it that Salmon is motivated by issues that are separable from compositional semantics *per se*. Salmon seems primarily concerned with the semantic *content* of an expression (or occurrence thereof)—in the sense of its corresponding constituent in a structured proposition—and thinks that “occurrence-based semantics illuminates”, since it “upholds intuitions about what is actually mentioned” (Salmon 2006: 418). In particular, Salmon is concerned with the contribution a bound occurrence of a variable in a sentence makes to the proposition expressed by the sentence. Any such account must eventually confront the *antimony of the variable*, see Fine (2003). Salmon (2006: 426 footnote 14) sketches his solution in terms of his bondage semantics; see also the related solution in Pickel and Rabern (2016).

²⁶This recent work builds in certain ways on King (2003) and Glanzberg (2011). See also Cappelen and Hawthorne (2009) and Schaffer (forthcoming).

²⁷The problem here is an instance of the more general problem that sentences embed under various shifty operators and thus require semantic values that are richer than propositions (Lewis 1980). See Ninan (2010) and Rabern (2012, 2017) for elaboration. See also literature on so-called “operator arguments”, e.g., Fritz et al. (2019).

the fact that according to many approaches to natural language syntax, any sentence which contains a quantifier in its direct object position such as (9) has an underlying *LF* where the quantifier has moved outside the scope of the sentence but leaves a co-indexed trace as in (10).

(9) John saw everyone.

(10) Everyone₁ [John saw *t*₁]

The structure (10) contains two immediate constituent expressions ‘Everyone₁’ and the “open sentence” (11).

(11) John saw *t*₁

Following the standard semantics of Heim and Kratzer (1998), Glanzberg and King relativize the semantic value of this open sentence to an assignment (à la Tarski). In particular, the semantic value of (11) relative to an assignment σ will be the proposition that John saw σ_1 (though in their presentation they simplify and treat this as a truth-value).

But this treatment raises obvious compositionality difficulties, because two open sentences may express the same proposition relative to an assignment, yet embed differently in larger quantified constructions. Contrast (11) with (12).

(12) John saw *t*₂.

These two open sentences express the same proposition relative to any assignment σ for which $\sigma_1 = \sigma_2$. But if the proposition expressed by a sentence is its compositional semantic value, then expressions differing by the substitution of (11) for (12) should have the same semantic values at assignment σ . But these expressions do not embed the same way under a quantifier such as ‘Everyone₁’.

(10) Everyone₁ [John saw *t*₁]

(13) Everyone₁ [John saw *t*₂]

(10) is true at σ just in case John saw everyone, but (13) is true at σ just in case John saw σ_2 . So they do not have the same semantic values at σ . (There is an obvious worry here that ‘Everyone₁ [John saw *t*₂]’ does not result from movement in any natural language

sentence. We use this example to simplify a point which could be made by contrasting ‘Everyone₁ [Someone₂ [*t*₁ saw *t*₂]]’ and ‘Everyone₁ [Someone₂ [*t*₂ saw *t*₁]]’, which may result from movement in the natural language sentences ‘everyone saw someone’ and ‘someone saw everyone’, respectively.²⁸)

The standard diet of responses are obviously available here. One might DENY SYNONYMY by holding that the expressions (11) and (12) are not genuinely synonymous. On this approach, the meaning of an open sentence is given by the function it determines from assignments to propositions. Two sentences are synonymous if they determine the same function. Alternatively, one might DENY SUBSTITUTION, by arguing that (11) does not occur in (10). The latter option would be a direct challenge to natural language syntax. So Glanzberg and King offer the hope for a third way, which at a certain level of abstraction resembles Salmon’s theory of binding considered above. Glanzberg and King insist that occurrence-based semantics has the virtue of retaining the classical view on the relationship between semantic value and content.

Glanzberg and King aim to vindicate the thought that unembedded sentences (11) and (12) have the same semantic values at σ . So they endorse $\llbracket \text{John saw } t_1 \rrbracket^\sigma = \llbracket \text{John saw } t_2 \rrbracket^\sigma$. Moreover, they endorse some version of the thesis that the meaning of a complex expression at an assignment is a function of the meanings of the expressions from which it derives at that assignment. So they need to explain how it is possible that $\llbracket \text{Everyone}_1 \llbracket \text{John saw } t_1 \rrbracket^\sigma \rrbracket^\sigma \neq \llbracket \text{Everyone}_1 \llbracket \text{John saw } t_2 \rrbracket^\sigma \rrbracket^\sigma$.

Their explanation appeals to “type adjusting”. In particular, they treat the indexed quantifier ‘Everyone₁’ as looking to combine with a predicate value, a function from term meanings to sentence meanings. When embedded under such an operator, they insist that a sentence shifts its meaning to a “type adjusted meaning”. In particular, the type adjusted meaning of a sentence γ at assignment σ when embedded under a quantifier numerically indexed by α is the function which takes an individual i from the domain and outputs the value of γ assessed at the assignment $\sigma[\alpha/i]$, which assigns i to the numerical index α .

Appealing to this type shifting rule allows Glanzberg and King to retain the basic thought that the semantic value of a complex expression at an assignment is a function of the semantic values of its syntactic components at that very same assignment. The trick is

²⁸One could also make a similar point by appealing to cases with bound *pronouns* (instead of traces), e.g. ‘every man thinks that he is clever’, but this raises further issues concerning how to treat free and bound pronouns that we wish to avoid here. See Rabern (2013) and Yli-Vakkuri (2013) for arguments along these lines.

that the various *occurrences* of an expression may have distinct semantic values (even when assessed relative to the same assignment). The relevant semantic values for the principle of compositionality are their “type adjusted meanings” and not their “normal meanings”. So in the context of ‘Everyone₁ [John saw t_1]’, the type adjusted meaning of ‘John saw t_1 ’ relative to σ is $\lambda i. \llbracket \text{John saw } t_1 \rrbracket^{\sigma[1/i]}$. By way of contrast, in the context of ‘Everyone₁ [John saw t_2]’, the type adjusted meaning of ‘John saw t_2 ’ relative to σ is $\lambda i. \llbracket \text{John saw } t_2 \rrbracket^{\sigma[1/i]}$. As a result, the type adjusted meanings of (11) and (12) are distinct at σ and so make different contributions as they occur in (10) and (13).²⁹

So for Glanzberg and King, the “type adjusted semantic values” of a sentence embedded under a quantifier with index α at assignment σ will be a function that takes an object i and outputs the sentence’s unembedded semantic value at $\sigma[\alpha/i]$. This corresponds exactly to the shifted semantic value of a sentence embedded under a single quantifier posited by Salmon, which we discussed above. So we purportedly have a single univocal expression which takes on different semantic values relative to σ depending on its linguistic context.

Does this proposal expand the standard diet of options? Recall from our discussion above that even an occurrence-based semantics must have a notion of *expression synonymy*, whereby two expressions are synonymous when they share meanings in their various occurrences. In the context of Glanzberg and King, two sentences would be synonymous *only if* for any index α and assignment σ , embedding them under a quantifier indexed by α produces the same type adjusted meaning relative to σ . Glanzberg and King seem to agree.

... type adjusted meanings of expressions should be defined up from the meanings originally assigned to those expressions, which in our case will be the semantic values of these expressions, relative to parameters. We should be able to get the type adjusted meanings from the meanings assigned by our semantic theory. (Glanzberg and King 2020: 17)

Clearly, the type adjusted meaning of an expression at an assignment σ such as (11) is not a function of its unembedded semantic value at σ , since the unembedded semantic value of (11) at σ is the same as the unembedded semantic value of (12) at σ , but their type adjusted meanings differ.

²⁹That is, $\lambda i. \llbracket \text{John saw } t_1 \rrbracket^{\sigma[1/i]} \neq \lambda i. \llbracket \text{John saw } t_2 \rrbracket^{\sigma[1/i]}$.

Noting this, Glanzberg and King observe that the expression meaning assigned to (11) cannot be its semantic value at σ . Rather, they say:³⁰

... [W]e should really think of [an expression] γ as having a *family of meanings relative to σ* corresponding to different choices of what gets assigned to [numerical index] α . As a result, we cannot think of $\llbracket \gamma \rrbracket^\sigma$ for a specific choice of σ as capturing the meaning of γ in the intuitive sense. For of course, it may be that $\llbracket \gamma \rrbracket^\sigma = \llbracket \psi \rrbracket^\sigma$ even though for another choice of assignment function σ^* : $\llbracket \gamma \rrbracket^{\sigma^*} \neq \llbracket \psi \rrbracket^{\sigma^*}$ and γ and ψ intuitively mean different things (' α skis' vs. ' α surfs'). Hence, it is only the family of meanings that in any intuitive sense captures the meaning of γ . [...] we need more than the meaning assigned to γ by our meaning assignment ($=\llbracket \gamma \rrbracket^\sigma$). In the case of [type adjusted meanings due to quantifiers], what is needed is what we've called the family of meanings of γ relative to σ . . . (Glanzberg and King 2020: 17)

So in order to predict the type adjusted meaning on the basis of the expression meaning—which is required for keeping their theory “compositionally kosher”—Glanzberg and King posit a notion of expression meaning which is determined by a family of meaning assignments. In particular, for each index α , there will be a meaning assignment to γ returning $\lambda i. \llbracket \gamma \rrbracket^{\sigma[\alpha/i]}$. Two expressions γ and ψ are synonymous relative to σ , when the meaning assignment for γ and ψ at σ agree relative to any index α .

On this approach the overall meaning assigned to an expression is given by its family of meanings. Two expressions agree in meaning at an assignment when they determine the same function from numerical indices to type adjusted meanings. Since the meaning assignments to (11) and (12) will differ for some choices of α , they have different families of meanings relative to σ . So they are not synonymous in the relevant sense. As a result, Glanzberg and King should not be seen as expanding the standard diet of options. Rather, they DENY SYNONYMY in response to the apparent failure of compositionality.

Given this they seem to be abandoning their preferred picture whereby sentential meaning is identified with the objects of assertion and belief. The *total* semantic value of a sentence (relative to an assignment) is not a proposition. A sentence has a family of meaning values—a family of type adjusted meanings—that are appealed to do the compositional work. Thus, the semantic value assigned to a sentence is given by its family of meanings. And two sentences agree in semantic value when they have the same family

³⁰We uniformly replace their notation with our own.

of meanings—that is, two sentences agree in semantic value when they determine the same function from linguistic environments to “type adjusted meanings”. Thus, the semantic value of a sentence (relative to an assignment) is not a proposition, instead it is a collection of occurrence meanings where *one* of those meanings is a proposition. So their view rejects the classical picture: the semantic values of sentence types are not propositional.³¹

4 Conclusion

We have argued that it is a myth that Frege-inspired occurrence-based semantics affords a genuine alternative strategy to the standard expression-based semantics. Occurrence-based approaches fall squarely into the standard general strategies for resolving compositionality failures.

When Frege’s remarks are construed in terms of the method of direct discourse, we argued that it amounts to positing ambiguity when faced with an apparent violation of compositionality. Thus it clearly collapses into the standard DENY SUBSTITUTION option. If instead the Fregean doctrine of “reference shift” is construed in terms of the method of indirect discourse as positing more sophisticated semantic rules, we argued that it amounts to enriching the semantic values. While this method has the virtues of systematicity and the avoidance of brute ambiguities, we argued that it ultimately collapses into the other standard option: DENY SYNONYMY.

Our two case studies, Pagin and Westerståhl (2010c) on quotation and Salmon (2006) and Glanzberg and King (2020) on variable binding, support our claim. The lesson in each case was a general one that would apply to occurrence-based treatments of any linguistic construction. On the occurrence-based approach (without brute ambiguities) a single expression has a diverse family of occurrence meanings—for each linguistic context a meaning from the family is employed to do the compositional work. We insist that the *semantic value* of an expression should be understood as the value assigned to the expression type that accounts for its contribution to truth-conditions across all the linguistic contexts in which it can occur. Lewis said it well when he said:

To be a semantic value is to be a big enough package of information. A semantic value worthy of the name must carry all the information that will be needed to generate

³¹See Stanley (1997), Ninan (2010), Yalcin (2014), and Rabern (2012, 2017) for discussion of the alternative picture which has roots in Dummett (1973), Evans (1979), and Lewis (1980).

other semantic values. Anything that we need to bundle together many of to get a big enough package is *ipso facto* not an adequate semantic value. (Lewis 1986: 44)

Thus, when a semantic theory is construed in an occurrence-based way *the* total semantic value assigned to an expression is given by its entire family of occurrence meanings. Two expressions agree in semantic value when they have the same family of meanings (viz. when they determine the same function from linguistic contexts to occurrence meanings). Thus, positing a difference in occurrence meaning requires a difference in meaning *tout court*. The strategy, then, ultimately amounts to *enriching* semantic values. And while enriching semantic values in order to avoid failures of compositionality may be theoretically motivated in many cases, it is, of course, nothing new.

References

- Bourbaki, N.: 1954, *Éléments de mathématique: Théorie des ensembles*, Paris, Hermann.
- Cappelen, H. and Hawthorne, J.: 2009, *Relativism and Monadic Truth*, Oxford University Press.
- Cappelen, H. and Lepore, E.: 2007, *Language turned on itself: the semantics and pragmatics of metalinguistic discourse*, Oxford University Press.
- Church, A.: 1951, A formulation of the logic of sense and denotation, in P. Henle, H. Kallen and S. Langer (eds), *Structure, method and meaning: Essays in honor of H.M. Sheffer*, Liberal Arts Press: New York.
- Church, A.: 1956, *Introduction to mathematical logic*, Vol. 13, Princeton University Press.
- Davidson, D.: 1979, Quotation, *Theory and decision* **11**(1), 27–40.
- Dummett, M.: 1973, *Frege: Philosophy of Language*, London: Gerald Duckworth.
- Evans, G.: 1979, Reference and Contingency, *The Monist* **62**, 161–189.
- Fine, K.: 2003, The role of variables, *The Journal of Philosophy* **100**(12), 605–631.
- Frege, G.: 1879/1967, Begriffsschrift, a formula language, modeled upon that of arithmetic, for pure thought, in J. van Heijenoort (ed.), *From Frege to Gödel: A source book in mathematical logic, 1879-1931*, Harvard University Press, pp. 1–82.

- Frege, G.: 1892/1952, On Sense and Reference, in P. Geach and M. Black (eds), *Translations from the philosophical writings of Gottlob Frege*, New York: Philosophical Library.
- Frege, G.: 1893/2013, *Gottlob Frege: Basic Laws of Arithmetic*, Oxford University Press.
- Frege, G.: 1902/1980, Frege to Russell (28-12-1902), in G. Gariel, H. Hermes, F. Kambartel, C. Thiel, A. Veraart, B. McGuinness and H. Kaal (eds), *Gottlob Frege: Philosophical and Mathematical Correspondence*, Oxford: Basil Blackwell.
- Fritz, P., Hawthorne, J. and Yli-Vakkuri, J.: 2019, Operator arguments revisited, *Philosophical Studies* **176**(11), 2933–2959.
- Glanzberg, M.: 2011, More on operators and tense, *Analysis* **71**(1), 112–123.
- Glanzberg, M. and King, J. C.: 2020, Binding, compositionality, and semantic values, *Philosophers' Imprint* **20**(2).
- Goodman, J. and Lederman, H.: 2019, Sense, reference and substitution, *Philosophical Studies* .
- Heim, I. and Kratzer, A.: 1998, *Semantics in Generative Grammar*, Blackwell Publishers.
- Janssen, T.: 1997, Compositionality, in J. van Benthem and A. ter Meulen (eds), *Handbook of Logic and Language*, Amsterdam: Elsevier.
- Kaplan, D.: 1964, *Foundations of intensional logic*, PhD thesis, UCLA.
- Kaplan, D.: 1968, Quantifying in, *Synthese* **19**(1/2), 178–214.
- Kaplan, D.: 1989a, Demonstratives, in J. Almog, J. Perry and H. Wettstein (eds), *Themes from Kaplan*, Oxford University Press, pp. 481–563.
- Kaplan, D.: 1989b, Afterthoughts, in J. Almog, J. Perry and H. Wettstein (eds), *Themes from Kaplan*, Oxford University Press, pp. 565–614.
- King, J.: 2003, Tense, modality, and semantic values, *Philosophical perspectives* **17**(1), 195–246.
- Klement, K. C.: 2002, *Frege and the Logic of Sense and Reference*, Routledge.

- Kripke, S. A.: 2008, Frege's theory of sense and reference: Some exegetical notes, *Theoria* **74**(3), 181–218.
- Lewis, D.: 1970, General semantics, *Synthese* **22**(1), 18–67.
- Lewis, D.: 1980, Index, context and content, in S. Kanger and S. Ohman (eds), *Philosophy and Grammar*, Amsterdam: Reidel, pp. 79–100.
- Lewis, D. K.: 1986, *On the Plurality of Worlds*, Wiley-Blackwell.
- Ninan, D.: 2010, Semantics and the objects of assertion, *Linguistics and Philosophy* **33**(5), 335–380.
- Packalén, S.: 2016, *Content and Composition. An Essay on Tense, Content and Semantic Value*, PhD thesis, Stockholm University.
- Pagin, P.: 2019, Belief Sentences and Compositionality: Notional Part, *Journal of Semantics* **36**(2), 241–284.
- Pagin, P. and Glüer, K.: 2006, Proper names and relational modality, *Linguistics and Philosophy* **29**(5), 507–535.
- Pagin, P. and Westerståhl, D.: 2010a, Compositionality I: Definitions and Variants, *Philosophy Compass* **5**(3), 250–264.
- Pagin, P. and Westerståhl, D.: 2010b, Compositionality II: Arguments and Problems, *Philosophy Compass* **5**(3), 265–282.
- Pagin, P. and Westerståhl, D.: 2010c, Pure quotation and general compositionality, *Linguistics and Philosophy* **33**(5), 381–415.
- Pardey, U. and Wehmeier, K. F.: 2019, Frege's *Begriffsschrift* theory of identity vindicated, in E. Lepore and D. Sosa (eds), *Oxford Studies in Philosophy of Language*, Vol. 1, Oxford University Press, pp. 122–147.
- Pickel, B. and Rabern, B.: 2016, The antinomy of the variable: A Tarskian resolution, *Journal of Philosophy* **113**(3), 137–170.
- Pickel, B. and Rabern, B.: in review, Frege and saving substitution.

- Pickel, B. and Rabern, B.: manuscript, Against Fregean Quantification.
- Portner, P. and Partee, B.: 2008, *Formal semantics: the essential readings*, Blackwell Publishing.
- Quine, W.: 1940, *Mathematical Logic*, Harvard University Press.
- Rabern, B.: 2012, Against the identification of assertoric content with compositional value, *Synthese* **189**(1), 75–96.
- Rabern, B.: 2013, Monsters in Kaplan’s logic of demonstratives, *Philosophical Studies* **164**(2), 393–404.
- Rabern, B.: 2017, A bridge from semantic value to content, *Philosophical Topics* **45**(2), 181–207.
- Salmon, N.: 2006, A theory of bondage, *The Philosophical Review* **115**(4), 415.
- Schaffer, J.: forthcoming, Confessions of a schmentencite: Towards an explicit semantics, *Inquiry: An Interdisciplinary Journal of Philosophy*.
- Soames, S.: 2011, True at, *Analysis* **71**(1), 124–133.
- Stanley, J.: 1997, Rigidity and content, in Heck (ed.), *Language, thought, and logic: Essays in honor of Michael Dummett*, Oxford University Press, pp. 131–156.
- Tarski, A.: 1936/1956, The concept of truth in formalized languages, *Logic, semantics, metamathematics* pp. 152–278.
- Wehmeier, K. F.: 2018, The proper treatment of variables in predicate logic, *Linguistics and Philosophy* **41**(2), 209–249.
- Yalcin, S.: 2014, Semantics and metasemantics in the context of generative grammar, in A. Burgess and B. Sherman (eds), *Metasemantics: New Essays on the Foundations of Meaning*, Oxford University Press.
- Yli-Vakkuri, J.: 2013, Propositions and compositionality, *Philosophical Perspectives* **27**(1), 526–563.