

9-2018

Unarticulated Constituents and Theories of Meaning

Jesse Rappaport
The Graduate Center, CUNY

[How does access to this work benefit you? Let us know!](#)

Follow this and additional works at: https://academicworks.cuny.edu/gc_etds

 Part of the [Philosophy of Language Commons](#), and the [Philosophy of Mind Commons](#)

Recommended Citation

Rappaport, Jesse, "Unarticulated Constituents and Theories of Meaning" (2018). *CUNY Academic Works*.
https://academicworks.cuny.edu/gc_etds/2885

This Dissertation is brought to you by CUNY Academic Works. It has been accepted for inclusion in All Dissertations, Theses, and Capstone Projects by an authorized administrator of CUNY Academic Works. For more information, please contact deposit@gc.cuny.edu.

UNARTICULATED CONSTITUENTS AND THEORIES OF MEANING

by

JESSE RAPPAPORT

A dissertation submitted to the Graduate Faculty in Philosophy in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2018

© 2018
JESSE RAPPAPORT
All Rights Reserved

ABSTRACT

Unarticulated Constituents and Theories of Meaning

by

Jesse Rappaport

Advisor: Stephen Neale

This work is an investigation into a phenomenon introduced by John Perry that I call ‘totally unarticulated constituents.’ These are entities that are part of the propositional content of a speech act, but are not represented by any part of the sentence uttered *or* of the thought that is being expressed - that is, they are fully unarticulated. After offering a novel definition of this phenomenon, I argue that totally unarticulated constituents are attested in natural language, and may in fact be quite common. This raises fatal problems for a prominent theory of underspecification defended by Jason Stanley, according to which all context-sensitivity (including unarticulated constituents) can be traced to covert variables in the syntax. I then use these findings to draw out important lessons for the philosophy of language, including a rejection of a long-standing Gricean issue known as the “meaning-intention problem.” I also explore the dialectic between Paul Grice’s intention-based semantics and Ruth Millikan’s teleosemantics, arguing that Millikan’s perception-based response to the problem of underspecification is untenable unless it is modified to give prominence to the speaker’s intentions.

Unarticulated Constituents and Theories of Meaning

by

Jesse Rappaport

This manuscript has been read and accepted for the Graduate Faculty in Philosophy in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

Date

David Rosenthal, Ph.D.

Chair of Examining Committee

Date

Nickolas Pappas, Ph.D.

Executive Officer

Supervisory Committee:

Stephen Neale, Ph.D.

David Rosenthal, Ph.D.

David Papineau, Ph.D.

THE CITY UNIVERSITY OF NEW YORK

Preface

The catalyst for this work is John Perry's 1986 article, 'Thought Without Representation.' In that essay, Perry engages in a thought experiment to motivate the claim that subjects can produce utterances that are, in some particular sense, "about" a given entity, even though there is no part of that sentence which designates that entity. Furthermore, what is striking about Perry's cases is that they suggest that the same phenomenon can occur at the level of *thought* – that is, that a thought can be about (or, "concern," in Perry's terminology) a given entity, even though there is no part of the thought that designates that entity. To give it a name, I say that such cases exhibit "totally unarticulated constituents."

Although Perry's essay was influential, I believe the consequences of the phenomenon he introduced have still not been fully grasped by those working in the philosophy of language. Thus, one of the goals of this text is to defend the existence of totally unarticulated constituents in natural language, and to explore what this means for theories of underarticulation, and for theories of meaning more generally. In particular, I argue that totally unarticulated constituents present fatal problems for a theory of underarticulation defended by Jason Stanley, according to which all context-sensitivity can be traced to covert variables in the syntax.

Thus, the narrow theoretical focus of this work is the topic of unarticulated constituents. More broadly, however, I will use unarticulated constituents as a theoretical tool to explore the conflicts between Gricean intention-based semantics and teleosemantic theories of natural language, such as that offered by Ruth Millikan. I focus on Millikan because she has

made a serious attempt to develop a teleosemantic theory of natural language, and to address the problem that underarticulation poses for such theories.

I will argue that Millikan's attempt to account for underarticulation by appealing to analogous facts about perception fails, and thus Millikan's teleosemantics for natural language is inadequate as it stands. However, I claim that it can be repaired by incorporating some insights from Grice. In addition, there are some aspects of Millikan's theory that can be used to improve our understanding of intention-based semantics. Thus, I believe a reconciliation between Grice and Millikan is called for.

There is no single thesis that this entire work is intended to defend. However, I do have a number of theoretical aims, most of which I have already alluded to:

- (i) To clarify the notion of unarticulated constituent, and to provide a definition that is extensionally correct.
- (ii) To defend the existence of totally unarticulated constituents in natural language.
- (iii) To draw out the consequences of totally unarticulated constituents for various theories of underarticulation (and, most saliently, to argue that they refute Jason Stanley's position).
- (iv) To draw out the consequences of totally unarticulated constituents for a problem in philosophy of language known as the "meaning-intention problem."
- (v) To clarify the dialectic between Paul Grice's intention-based semantics and Ruth Millikan's teleosemantics.
- (vi) To critique Millikan's response to the problem of underarticulation.

(vii) To suggest a way to reconcile Grice and Millikan.

In brief, I intend to defend the existence of totally unarticulated constituents, and then explore the consequences of this phenomenon for theories of underarticulation, and for the conflict between intention-based semantics and teleosemantics. The progression of the text over the following chapters follows this line of development.

This work would not have been possible without the support of friends, family, and faculty. I am very grateful to everyone whom I have conversed with about these topics, and especially to Stephen Neale, my advisor, and the members of my committee. I am also grateful for having had the opportunity to present these ideas at conferences over the years, especially in Dubrovnik, Croatia, and at the *Context, Cognition, and Communication* conference in Warsaw, Poland, where I presented part of this work to John Perry.

Table of Contents

1. How to Define ‘Unarticulated Constituents’	1
2. Totally Unarticulated Constituents	28
3. Is a Uniform Account Possible?	68
4. Is There a Meaning-Intention Problem?	87
5. A Gricean Twist on Teleosemantics	111

1. How to Define ‘Unarticulated Constituents’

I. Introduction

Suppose a child steps out of the house, notices that there is snow falling, and exclaims:

(1) It’s snowing!

In performing this utterance, the child has managed to say something about the weather. More specifically, supposing she lives in New York City, she has managed to say something about the weather in New York. Even more specifically, she has managed to say something about the weather in New York at the time of her utterance. Thus, in addition to saying something about the weather, the child has said something about New York, and also about the current time. That’s a lot of referring to do in uttering such a short sentence!

We can tell from the linguistic information in the sentence uttered that the child is talking about the weather – she uses the verb ‘snow.’ However, there doesn’t appear, at first glance, to be anything in the sentence uttered that tells us that the statement made is about the weather *in New York*. A speaker in Los Angeles could use the same sentence to assert that it’s snowing in Los Angeles. Thus, if we take appearances at face value, it seems that the child has succeeded in expressing a thought about New York (or, about her general location), despite the fact that she has not *referred* to New York *with* any part of the sentence uttered. In other words, New York appears to be an *unarticulated*

constituent of the utterance – it is a constituent of the proposition meant by the speaker, but it is not denoted by any of the verbally articulated components of the utterance.¹

Weather predicates provide a striking example of unarticulated constituents, but the general phenomenon is quite pervasive. Consider the following sentences:

- (2) Alice is tall.
- (3) It's three o'clock.
- (4) The mayor is on TV again.
- (5) Bob believes that Superman can fly.
- (6) Carol knows that her car is parked in the garage.

In each of these cases, what is meant seems to depend on factors that are not explicitly mentioned or referenced in the utterance:

- (2a) Alice is tall (*for a three-year-old*).
- (3a) It's three o'clock (*P.M. Eastern Standard Time*)
- (4a) The mayor (*of Toronto*) is on TV again.

¹ In this text, I will assume a structured-proposition view of propositional content (cf. (King 2007), (Soames 1987), (Salmon 1986)), mainly for ease of exposition. Intuitively, propositions can involve certain entities, in the straightforward sense that the proposition that Theatetus is flying now involves Theatetus and the time of utterance; the truth conditions of that proposition depend on how things are with Theatetus and the present time. It is easiest to capture this intuition on the structured-proposition view, since we can say that Theatetus is literally a constituent of the proposition *that Theatetus is flying now*. On a possible-worlds conception of propositions, propositions lack internal structure, thus it is harder to express the intuitive relation of object-dependence in the theory's own terms. What is really important for the phenomenon that I am interested in, however, is the intuitive sense in which propositions can involve certain entities, a fact which seems to hold independently of one's views on the metaphysics of propositions.

- (5a) Bob believes that Superman can fly (*under the mode of presentation of a man in a red cape*).
- (6a) Carol knows that her car is parked in the garage (*relative to the standards of knowledge currently in force*).

Indeed, for each of these cases, philosophers have proposed unarticulated constituents as responses to the context-sensitivity evidenced.

Although the phenomenon of unarticulated constituents should have some intuitive force, Adam Sennet (2011) has shown that providing a careful definition is more difficult than at first appears. In this chapter, I will attempt to provide a definition of *unarticulated constituent* that responds to Sennet's arguments.

II. The Definition of *Unarticulated Constituent*

A. Ostensive Definition

Thus far, I have attempted to get a grip on the issue of unarticulated constituents by way of example, or ostensive definition. Although such examples are illuminating, is this theoretically acceptable, or do we need to offer necessary and sufficient conditions for being an unarticulated constituent?

Sennet (2011) argues that we should not be satisfied with ostensive definition.

The ostensive definition merely isolates a phenomenon. The whole point of the explicit definition is to provide a theoretical tool to explain this phenomenon. Without the explicit definition, all we have is a label for the phenomenon, not a distinctive addition to our toolbox for characterizing the slippage with which we began.²

Sennet claims that without an explicit definition, we cannot offer an explanatory theory of unarticulated constituents, since all we would have is a descriptive “label.” Presumably, the tacit assumption is that a theory cannot explain a phenomenon if that phenomenon cannot be defined in the terms of the theory.

Although I believe that an adequate definition of ‘unarticulated constituent’ *can* be achieved, this assumption is rather dubious. Newtonian physics can help us understand why a chair falls if it is kicked, and indeed, to some degree, what a chair *is*, since a chair is a physical object. Nonetheless, one cannot define ‘chair’ in the language of Newtonian physics (or, for that matter, in any other language, it would seem). ‘Chair’ being a vague, family-resemblance notion, ostensive definition is probably the best that one can do. It isn’t the case, then, that we need X to be definable in terms of theory T in order for theory T to do explanatory work with respect to X.

However, the extent to which such tolerance of ostensive definition is acceptable depends, in large measure, on the kind of work you want the definiendum to do. Chairs do not themselves figure in any larger explanatory project. However, philosophers have enlisted unarticulated constituents to play an explanatory role in virtually every domain of contemporary philosophy. Although every

² (Sennet 2011, 426).

theory must have its primitives, the more primitives it has, the less explanatory it will appear. Thus, I agree with Sennet's criticism that many in the philosophy literature have been too willing to put unarticulated constituents to work without relying on any precise, descriptively adequate definition.

In the following section, I will attempt to develop a first-pass definition of 'unarticulated constituent' and consider Sennet's objections to existing treatments of the issue.

B. The Need for Structure

i. Some Existing Definitions

As indicated above, being an unarticulated constituent is a matter of standing in a certain relation to a representation – an unarticulated constituent is a constituent of a proposition that is unarticulated with respect to a particular vehicle, or mode of articulation. Intuitively, we want to say something like: X is an unarticulated constituent of a proposition p , relative to a representation R , just in case R is about X , even though R does not fully or explicitly denote or represent X .

Let us consider some attempts to make the intuitive conception of unarticulated constituents more precise.

According to Francois Recanati (2002), unarticulated constituents “are part of the statement made even though they correspond to nothing in the uttered sentence.”³ He also states that they are “part

³ (Recanati 2002, 300).

of the intuitive truth-conditions of the utterance.”⁴ So, for Recanati, *X* is an unarticulated constituent of some statement just in case it is part of the intuitive truth-conditions of that statement, and *X* does not “correspond” to any part of the sentence uttered. Thus, we have a notion of unarticulatedness as “representation without correspondence.” As we will see, there is considerable difficulty in spelling out the notion of correspondence required by this definition.

Stephen Neale (2007) defines unarticulated constituents in a more speech-act-theoretic way.

According to Neale, if I utter,

(7) It’s raining,

“I could be expressing the proposition that it’s raining in Reykjavík. Yet there is no part of sentence [(7)] that I would be using to refer to Reykjavík.”⁵ On Neale’s view, *X* is an unarticulated constituent of a proposition, relative to some speech act, just in case *X* is a part of the proposition meant, or expressed, but there is no part of the sentence *with which* the speaker refers to *X*. Thus, unarticulatedness is a matter of “representation without semantically encoded referential meaning intentions.”

One of the primary figures in the debate, John Perry (1998), says, regarding a sentence like (7):

In this case, I say that the place is an unarticulated constituent of the proposition expressed by the utterance. It is a constituent, because, since rain occurs at a time in

⁴ (Recanati 2002, 300).

⁵ (Neale 2007, 254).

a place, there is no truth-evaluable proposition unless a place is supplied. It is unarticulated, because there is no morpheme that designates that place.⁶

Here, Perry suggests that *X* is an unarticulated constituent of a proposition expressed by an utterance just in case *X* is a part of the proposition expressed, but there is no part of the sentence that designates *X*.

In this passage, Perry also seems to endorse the inference from (i) it is metaphysically necessary that rain takes place at a time and a place, to (ii) there is no truth-evaluable proposition regarding that event unless a place is supplied. However, this type of inference is rather dubious. It is metaphysically necessary of all physical events that they take place at a time and a place – there is nothing special about rain in this regard. However, we cannot infer from (iii) it is metaphysically necessary that if Alice sat down, then she sat down at some location, to (iv) there is no truth-evaluable proposition stating that Alice sat down unless a place is supplied. Perry seems to be correct in stating that *It's raining* is not truth-evaluable without a location, but wrong to infer this from the stated metaphysical necessity.

Perry also writes that “the starting point of investigation” ought to be “the question of how there can be a constituent in the proposition, with no corresponding expression in the utterance.”⁷ Here, again, we are faced with a notion of *correspondence* that needs to be explained.

Finally, let us consider Crimmins’s (1992) description:

⁶ (Perry 1998, 9).

⁷ (Perry 1998, 9).

An unarticulated constituent of the content of a statement is an item that is used by the semantics as a building block of the statement's content but is such that there is no (overt) expression in the sentence that supplies the object as its content. In a semantics that takes propositions to be structures containing objects and properties, an unarticulated constituent is simply a propositional constituent that is not explicitly mentioned: it is not the content of any expression in the sentence⁸

This description is similar to the others, but differs insofar as it entails that the unarticulated constituent is processed semantically, rather than as a result of free enrichment.

ii. The Problem of Structure

While each of these theorists seems to be getting at the same core phenomenon, Sennet (2011) observes that such characterizations do not even achieve descriptive, or extensional, adequacy. In particular, they fail due to sentences like:

(8) I'm in New York and it's raining.

Each of the analyses offered above state that an unarticulated constituent is a constituent of the proposition intuitively expressed to which there is *no* corresponding part in the sentence uttered. But, as seen in (8), the phenomenon arises also in cases where the unarticulated constituent is referred to elsewhere in the sentence, as in:

⁸ (Crimmins 1992, 16).

(8a) I'm in New York and it's raining (*in New York*).

Intuitively, (8) seems to involve the same kind of phenomenon we have been addressing; the speaker states something about the weather *in New York* despite the fact that she does not explicitly say that it's raining *in New York*. Indeed, the equality of the two locations is in no way required. Suppose that Alice lives in a city where it has never snowed. One day, she decides to leave town and visit New York. While she's there, she calls her brother in her hometown and hears that it's snowing for the first time ever:

(9) I'm in New York and it's snowing!?! I can't believe it.

(8) appears to demonstrate a clear case of an unarticulated constituent, but the definitions offered above do not apply to it, since there *is* a corresponding part in the sentence that denotes New York, or with which the speaker refers to New York: namely, the occurrence of *New York*.

What we need then, according to Sennet, is an account of unarticulated constituents that relates not simply to utterances or sentences as a whole, but to *occurrences* of expressions within a sentential structure. However, extending the definition in such a manner incurs significant cost, so he claims. In particular, it requires spelling out the notion of correspondence, so that we can say that *X* is an unarticulated constituent of the proposition expressed just in case there is no *corresponding* part of the sentence, even if there is some other (non-corresponding) part that refers to *X*. Applying this to (8), we could say that New York is articulated with respect to the predicate *being in* (there is a part of the structure of the sentence that corresponds to New York's position in the structured proposition),

but is unarticulated with respect to the predicate *raining* (there is no corresponding part in the sentence).

Sennet argues that this is a large burden for the defender of underspecification to take on, since it requires

...positing sufficient sentence structure to coordinate the positions in a proposition with similarly individuated parts of the sentence. This requires a fair bit of structure to be imposed on the sentence and, arguably, a fair bit of it to be hidden structure: after all, we need enough structure to point out that there is a spot in the sentence that correlates with, say, the second occurrence of [New York] in [(8)] that doesn't contain ['New York'].⁹

It isn't clear what Sennet is taking issue with in this passage. He seems to be arguing that in order to hold that X does not have a corresponding position in a sentence S , we have to posit the relevant (imaginary?) structure in S and observe that it does not contain an expression denoting X . But the lack of a corresponding structure does not imply the existence of some alternative structure of the same complexity.

Furthermore, the defender of unarticulated constituents is not opposed to hidden syntactic structures *per se* – rather, he or she is (or ought to be) opposed to invocations of hidden structure solely on the basis of philosophical speculation, or to resolve a philosophical issue. Thus, there

⁹ (Sennet 2011, 425).

seems to be nothing preventing one from explicating unarticulated constituency in terms of lack of corresponding position in a structure.

Stephen Neale (2007) acknowledges the “weakness” of his initial characterization in light of Sennet’s criticisms, and opts for a strategy similar to the one discussed by Sennet: “Any patch would seem to require taking into account *position* or *occurrence*, pairing *portions of X*” (where there are such) with *constituents* of $\langle \dots a \dots \rangle$.”¹⁰ Neale favors the following sort of view:

a is an unarticulated constituent of $\langle \dots a \dots \rangle$ relative to X only if there is a natural expansion $E(X)$ of X with the following property: if $E(X)$ rather than X had been used to express $\langle \dots a \dots \rangle$ the tokening $E(X)$ would have a portion corresponding to and representing *a*.”¹¹

Thus, an entity, X , is an unarticulated constituent of an utterance of a sentence S just in case X does not have a corresponding position in S , but X *does* have a corresponding position in a (the?) “natural expansion” of S .

In this passage, Neale offers only a necessary condition for being an unarticulated constituent. It isn’t clear how the analysis could be improved to provide sufficient conditions as well. It also isn’t clear how much work is being done by the phrase “natural expansion.” Consider:

(10a) It’s raining.

¹⁰ (Neale 2007, 85).

¹¹ (Neale 2007, 85).

(10b) It's raining in Los Angeles.

Suppose (10a) is uttered in New York, and is used to mean that it's raining in New York. Is (10b) a “natural expansion” of (10a)? Let's suppose that it is. Then, Neale's analysis entails: if (10b) rather than (10a) had been used to express the proposition that it's raining *in New York*, (10b) would have a portion corresponding to and representing New York. But this is patently false. So (10b) must not be a natural expansion of (10a). But then we are left with the feeling that the natural expansions are just the expansions that intuitively capture our impressions of unarticulated constituents, in which case the analysis is unexplanatory.¹²

In sum, it isn't clear how the defender of unarticulated constituents who wishes to pursue a Neale-style strategy should spell out the notion of correspondence necessary for the definition. Doing so seems to require at least positing a notion of propositional structure that closely corresponds to sentential structure, a costly though not totally implausible hypothesis. But even granted such a fine-grained notion of structure, it isn't obvious how to capture the required correspondence, or isomorphism, of structure, since the appeal to a “natural expansion” does not appear tenable.

III. Defining *Unarticulated Constituent* Without Correspondence

¹² Neale (2016) discusses the issue in more recent work, especially the need for taking account of occurrences of expressions, but does not appear to draw any substantive conclusions.

I claim that there is a way of capturing unarticulated constituency without positing any additional structure, or appealing to counterfactual utterances, and it is capable of achieving *nearly* full descriptive adequacy.

(UC1) X is an unarticulated constituent of a representation R just in case the proposition expressed by R predicates F of X , but there is no part of R that semantically encodes the predication of F to X .

It is important to note that by ‘predication’ here, I have in mind something like ‘function application.’ If a speaker says, *It’s not raining*, and means thereby that it’s not raining in New York, then one might want to say that the speaker is not predicating being in New York of a rain event. In one sense, ‘predicating F of X ’ involves actually saying that X is F . Since *It’s not raining* in fact *denies* that X is F , it does not involve predication in that sense. However, in a broader sense, we can think of predication along the same lines as function application. For instance, if I say, *It’s not the case that John is tall*, then even though I am denying that John is tall, I am still semantically applying *tall* to *John* in the semantic composition – they are combined via function application. I think it’s reasonable to think of this as predication, and it is this sense I will employ in this text.

Let us observe how (UC1) handles the problematic case discussed above:

(8) I’m in New York and it’s raining.

Intuitively, (8) contains an unarticulated constituent due to the presence of *it’s raining* without an explicit location, however previous definitions have failed to capture this fact: (8) is ruled out as an

instance of unarticulated constituency because New York *is* articulated in the sentence. However, it follows from (UC1) that New York is an unarticulated constituent of an utterance of (8), since the proposition expressed by that utterance predicates rain *in* New York, but there is no part of the sentence uttered that semantically encodes the predication of rain at the location of New York.

Similar reasoning applies to Sennet's (2011) examples:

(11) Brooke comes to Vancouver when it snows.

Intuitively, (11) contains an unarticulated instance of the propositional constituent Vancouver, in virtue of the fact that it quantifies over snowing events *in* Vancouver, but it doesn't not explicitly say *when it snows in Vancouver*. (UC1) works for this example as well. The proposition expressed by an utterance of (11) quantifies over snowing events in Vancouver, but there is no part of the sentence uttered that semantically encodes quantification over snowing events in Vancouver.

There is one kind of case in which (UC1) fails to make the correct predictions: namely, those cases in which the relevant predication is duplicated:

(12) It's raining in New York, and it sure is raining!

(13) If it's raining in New York, then it's raining.

(14) Either it's raining in New York, or it's *not* raining.

(UC1) does not count (12)-(14) as containing unarticulated constituents, because the relevant predication (that there is a raining event *in* New York) *is* semantically encoded by another part of the

sentence. On the other hand, these utterances are somewhat pragmatically defective. (14) is acceptable, but may be a case of syntactic ellipsis, in which case both predications are encoded (but one is deleted). It is unclear whether to treat these sentences as genuine counter-examples.

However, there is a way to improve our formulation.

(UC2) X is an unarticulated constituent of a representation R just in case there is a (possibly improper) part (or, sub-clause) of R that expresses a proposition that predicates F of X , but *that part* does not have any parts that semantically encode predication of F to X .

This successfully captures (12)-(14). Consider (12). (12) counts as containing an unarticulated constituent because (12) contains a part (namely, the second clause *it sure is raining*) that expresses a proposition that predicates F of X , but that clause does not have parts that semantically encode predication of F to X .

Thus, we have managed to achieve descriptive adequacy, and the only structure we require is enough to separate the two clauses of (12). This is surely no burden at all.

Clearly, with this formulation, I am committing myself to a form of compositional semantics. For I am assuming that “predicating F of X ” is something that both propositions and sentences can do.

How can this be so?

Considering propositions as ordered pairs, we can understand predication of F of X as a proposition's containing somewhere, as a constituent, the pair $\langle F, X \rangle$. On the sentential side, we do

not directly have the act of predication from syntax. However, a compositional semantics allows us to state, e.g., that a phrase like $X \text{ is } F$ predicates F of X because it encodes function application of F to X , which is a formal semantic notion of predication.

Nevertheless, although my account of unarticulated constituency is committed to the existence of substantial compositional structure, it does not require any notion of correspondence, like the one discussed by Sennet. Thus, we have evaded one of the primary difficulties witnessed above.

In the next section, I will compare my approach to an alternative one developed within a Millikanian framework.

IV. Representation and Unarticulated Constituents for Millikan

A. Background on Millikan's Theory

Ruth Millikan's (1984, 2005) teleosemantic theory of representational content for biological signs and natural language is complex and many-faceted. Here, I will discuss her theory of natural language meaning, focusing on her distinction between satisfaction conditions and semantic mapping functions.

According to Millikan, language is a complex of *natural conventions*. A natural convention is a *reproduced pattern of behavior* that is reproduced in large part due to "weight of precedent." Some unpacking: X is a reproduction of Y just in case X resembles Y in some respects and the properties of X are

counterfactually dependent on the properties of *Y*. For example, a photocopy is a reproduction of a photograph because it resembles the photograph in important respects, and if the photograph had been different (had it been a picture of a tree, say), then the photocopy would have differed as well.

This describes a reproduction. But not all reproduced behavior counts intuitively as a convention.

For instance, suppose you show me a new way to hold chopsticks that makes it easier to eat sushi. If I adopt your method, then my behavior is a reproduction of yours. However, it wouldn't be a *conventional* reproduction, since it lacks the intuitively necessary property of *arbitrariness*. A convention ought to be a pattern of behavior that is reproduced simply because *it* (rather than some equally serviceable alternative type of behavior) has been produced in the past, and not because it is intrinsically superior to the alternatives. For instance, it's conventional that Americans drive on the right and British people drive on the left. It's *not* conventional that American and British drivers check their mirrors before beginning a journey – even if this behavior has been learned (and hence reproduced), it is not being reproduced solely, or mainly, due to “weight of precedent.”

Language conventions, then, according to Millikan, are natural conventions. Their meaning is determined by their *stabilizing function*. In general, the etiological function of some trait or behavior is the effect that that trait has which explains the continued reproduction of the trait. For example, activating my turn signal both informs other drivers of my vehicular intentions, and produces a rhythmic clicking sound in the car – however, while the former is a *function* of my turn-signaling behavior, the latter is merely an incidental side-effect. This follows from the fact that if my activating the turn signal failed to produce the former effect (because, say, my tail lights have gone out), while continuing to produce the latter effect, I would not persist in reproducing the behavior.

In the context of linguistic communication, where we require the cooperation or coordination of both speaker and hearer, Millikan speaks of the “stabilizing function” of a linguistic form as the function which serves to stabilize the continued participation of speakers and hearers by serving their interests sufficiently often: “Because the conventional function of a linguistic form will remain stable only if it continues to serve the interests of both speakers and hearers often enough, I call it a ‘stabilizing function.’”¹³

What, then, are the stabilizing functions of language forms? For Millikan, indicative sentences have the function of causing true beliefs in hearers, while imperative sentences have the function of guiding the hearer to perform some action. In a normal situation, if I say, *I’m hungry*, and you form the true belief that I’m hungry, then the sentence *I’m hungry* has fulfilled its proper function. Any further purposes that I, as a speaker, may have (such as getting you to bring me some food) are *my* purposes, but not the purposes of the sentence, whose function is simply to engender a true belief about my hunger.

What makes it the case, though, that *I’m hungry* has the stabilizing function of causing you to believe that I’m hungry, rather than causing you to believe that I’m tired? Well, so the story goes, my linguistic and cognitive faculties are designed in such a way that there is a “Normal”¹⁴ mechanism by which, upon hearing and parsing an utterance, I form a belief as a result. This belief orients my behavior in a way that, *ceteris paribus*, is appropriate given the sentence’s truth condition. Thus, the

¹³ (Millikan 2005, 58).

¹⁴ If a trait has a function of producing some beneficial effect, then there will be a small set of conditions under which, or mechanisms with which, the trait will produce that effect. These typical conditions for the performance of a proper function are the Normal conditions. (Note that the Normal conditions themselves need not be typical, i.e., it may be rare that the Normal conditions obtain.) Following Millikan, I will capitalize ‘Normal’ when using it in this biological sense.

truth conditions of an indicative sentence depend on the content of the belief that is formed, by way of Normal mechanisms, as a result of hearing that sentence.

This gives us a framework for determining the truth conditions of a sentence. However, Millikan does not hold that the meaning of a sentence, or any other representation, is exhausted by its truth conditions, or satisfaction conditions. Part of what we think of as the meaning of a representation depends on *how* the representation maps onto, or determines, its truth conditions. Millikan refers to this as the representation's *semantic mapping function*.

Whereas satisfaction conditions determine the conditions under which an utterance of a sentence is true or false, the semantic mapping function is a matter of the degree of structural isomorphism between representation and state of affairs. Coincidentally, Millikan often explains this notion by way of weather predicates. She asks us to compare the sentences (i) *It's raining* and (ii) *Rain is falling here now*. According to Millikan, "The truth-conditions of 'It's raining' and of 'Rain is falling here now' are the same, but the semantic mapping function is different."¹⁵ This is because (ii) has more meaningful structure than (i).

Meaningful structure is posited in the representational vehicle when one observes that transformations in the vehicle produce systematic changes in the state of affairs represented: "The semantic-mapping function is given by rules according to which significant transformations of the sentence that conserve its syntactic form yield different truth- or satisfaction-conditions."¹⁶ For instance, we could apply some transformations (substitution of words of the same syntactic

¹⁵ (Millikan 2005, 64).

¹⁶ (Millikan 2005, 63).

category) to a sentence like (ii), in order to obtain: (iii) *Rain is falling in New York now*; and, (iv) *Rain was falling here yesterday*. Such transformations produce systematic, predictable changes in the state of affairs represented. Substituting *in New York* for *here* in (ii) results in a state of affairs that is similar to the state of affairs represented by (ii) except there is rain falling in New York, rather than at the location of utterance. It is the possibility of such systematic correspondence in transformations (i.e., isomorphism) that justifies positing meaningful structure in a representation.

For these same reasons, we are led to deny the existence of any meaningful part of (i) *It's raining* that represents the location of rain, even though for each utterance, the truth conditions that the sentence maps on to depend on the location of utterance.¹⁷ *It's raining* might have the same truth conditions as *It's raining in New York* on one occasion, and *It's raining in Los Angeles* on another; however, since these changes in meaning are not reflected in any changes in the vehicle, we may conclude that the location is not captured by the semantic mapping function.

B. Millikan on Underarticulation

As we can see from the preceding discussion, Millikan's theory of semantic mapping is, in effect, an account of articulatedness: "The semantic mapping of a sentence articulates it, placing it in a logical space of contrasting possibilities. Its truth-condition is not, as such articulated."¹⁸ Let us attempt to clarify Millikan's views on articulation.

¹⁷ We will, for now, assume (wrongly) with Millikan that *It's raining* pertains (always, or by default) to the location of utterance.

¹⁸ (Millikan 2005, 64).

In the just-cited passage, Millikan suggests that it is *sentences* that are articulated (with respect to the logical space that they are placed in by their semantic mapping functions). She also states: “‘Many drops of water are presently precipitating from the atmosphere and landing close to this place’ also has the same truth-condition but is articulated by yet another semantic-mapping function.”¹⁹ This seems to differ from the standard usage, according to which it is *constituents* (of a sentence/utterance/thought) that are articulated. On this formulation, a *sentence* is *articulated by* a semantic mapping function.

Let us attempt to formulate a definition of ‘unarticulated constituent’ in this Millikanian framework. We will deviate from Millikan in speaking of unarticulated constituents, rather than sentences.

(UCM1) X is an unarticulated constituent of a sentence S just in case X is part of the state of affairs represented, but there is no part of S that maps onto X in accordance with the semantic mapping function.

Obviously, in order to accommodate Sennet’s counter-examples, we would need to modify this in order to apply to clauses, as suggested above:

(UCM2) X is an unarticulated constituent of a sentence S just in case there is a (possibly improper) part of S that maps onto a state of affairs involving X , but there is no part of that part that maps onto X .

¹⁹ (Millikan 2005, 64).

This formulation should be descriptively adequate. However, it is inconsistent with Millikan's other views.

The semantic mapping function is a kind of correspondence relation between sign and world. Thus, Millikan seems to require some notion of correspondence. Indeed, she is an open defender of a picture, or correspondence, theory of meaning, giving a biological twist on a kind of Wittgensteinian (Tractarian) position. If we assume such a picture theory, then it seems like it should be easy to characterize *unarticulated constituent*: an unarticulated constituent is present if there is something in the represented scene that doesn't appear in the picture. However, Millikan's view is more nuanced, since she claims that correspondence at the sentence level is fundamental, while correspondence of parts is at best a derived notion.

In *Language, Thought, and Other Biological Categories* (1984), Millikan raises the problem of the unity of the proposition: if 'Theatetus flies' expresses the proposition that Theatetus flies, and the latter is an entity composed of Theatetus and the property of flying, then all we seem to be left with is a pair consisting of a thing and a property: <Theatetus, flying>. What's missing is the *application* of the property *to* the entity. Thus, in addition to Theatetus and *flying*, we seem to need to mention the *instantiates* relation, so that it is clear that Theatetus is *instantiating* flying, as in: <instantiating, <Theatetus, flying>>. However, now we are off on a regress. For, once again, it appears that we are left with a mere collection of relations, properties, and things (three now, rather than two), and nothing actually *doing* the relating; thus, another instantiation relation will have to be invoked, and another one, *ad infinitum*. As Millikan describes the problem:

‘Theatetus’ corresponds to Theatetus; ‘flies’ corresponds to flies (flying); the relation between ‘Theatetus’ and ‘flies’ corresponds to a real or instantiated relation (instantiated, e.g., between Theatetus and walking). Again, everything corresponds – *except* ‘Theatetus flies.’²⁰

Millikan attributes this problem to an atomistic ontology and a reference-based semantics: if you start with parts of things, you will never be able to put them together. As a response, she claims that correspondence holds fundamentally at the sentence level: sentences, when produced under Normal circumstances, correspond to the state of affairs that they map on to – this state of affairs is their “real value.” A term, on the other hand, “is supposed to appear in the context of a sentence. That is, this condition is a Normal condition for its proper performance. Second, qua in the context of a sentence, it is supposed to correspond to or map onto something – its referent.”²¹ Thus, since it is only in the context of a sentence that a term can have a real value, or, can refer, the correspondence between term and referent should be viewed as parasitic on the correspondence between sentence and state of affairs: “It is, in the *first* instance, the *sentence* that corresponds, *hence* ‘Theatetus.’”²² Furthermore, there is no unity of the proposition problem for states of affairs: a state of affairs is (supposed to be) a complete picture, not a collection of parts that need to be “glued” together somehow. Thus, by rejecting atomism, we avoid the problem of the unity of the proposition. “Somehow, we must begin by correlating *sentences* with world affairs, correspondence of words with things comings after”²³

²⁰ (Millikan 1984, 102).

²¹ (Millikan 1984, 104).

²² (Millikan 1984, 106).

²³ (Millikan 1984, 107).

Given this background, we should still be able to say (ignoring Sennet's criticisms) that X is an unarticulated constituent of a proposition, relative to a sentence S , just in case X is a part of the state of affairs that S maps onto (or is supposed to map onto) but there is no part of S that corresponds to, or maps onto, X . Since term reference is legitimate, though parasitic or derived, this formulation should be unobjectionable. However, Millikan also endorses a radically pragmatist, structuralist view of ontology.

First, we have already seen that she is committed to the existence of things called states of affairs. Unfortunately, she does not expound much on what these are. Nonetheless, states of affairs are represented by way of isomorphisms between structures:

The governing idea here is that, in the first instance at least, it is *transformations* of the icon that correspond to *transformations* of the real value – *operations* upon the icon that correspond to *operations* upon the real value – not elements of the icon that correspond to elements of the real value.²⁴

She goes on to develop this structuralist picture:

Whatever is considered as subject to a set of transformations is as such 'articulate' in a certain way. It is not articulated into parts but into invariant and variant aspects. What remains unchanged under all transformations in a transformation set is the

²⁴ (Millikan 1984, 107).

invariant aspect of a thing relative to that set; what changes under these transformations are variant aspects.²⁵

Thus, so far we have primitive states of affairs, and then articulations of these states of affairs by representations, according to a principle of isomorphism.

Millikan pushes this picture even further by endorsing a radically pragmatist view of articulations:

Each of these sets of transformations [i.e., different articulations of the same state of affairs] may articulate the same affair in a different way, and there is no reason to suppose in advance that any of these articulations would have to be *the* elemental or ideal articulation of that affair.²⁶

Thus, despite pursuing a highly naturalistic project, Millikan endorses a rather extreme and speculative structuralist-pragmatist view of ontology. Since my interest at the moment is not in ontology, but simply in definitions, I will not attempt to critique this view

What does all of this have to do with unarticulated constituents? Given Millikan's structuralist-pragmatist ontology, (UCM2) is seen to be inadequate because there is no ultimate fact of the matter about whether or not some entity X is a part of some state of affairs. States of affairs are articulated into parts (or variant and invariant "aspects" – it isn't clear what this means) in relation to a

²⁵ (Millikan 1984, 107).

²⁶ (Millikan 1984, 109).

representation; and since there is no privileged mode of articulation, there can be no concrete fact of the matter regarding whether or not *X* is a part of some state of affairs.

One potential solution would be for Millikan to modify (UCM2) so that instead of speaking of states of affairs containing or “involving” objects *tout court*, it speaks of states of affairs involving objects according to a highly explicit, canonical articulation:

(UCM3) *X* is an unarticulated constituent of a sentence *S* just in case there is a (possibly improper) part of *S* that maps onto a state of affairs that (according to a highly explicit, canonical articulation) involves *X*, but there is no part of that part that maps onto *X*.

Thus, New York is an unarticulated constituent of *It's raining*, as uttered in New York, because there is an explicit, canonical articulation of this state of affairs that makes explicit the location: *It's raining now in New York*. The uttered sentence corresponds to that state of affairs, but there is no part of the sentence that corresponds to New York.

Perhaps this fix will be satisfactory to anyone who is attracted to Millikan's pragmatist structuralism. However, I will not pursue such a project; I fail to see how it produces any metaphysical gains, and I do not find a purely transformation-based structuralist ontology to be coherent. Thus, I will take talk of propositional constituents as literally true. Indeed, in more recent work, Millikan does not make frequent appeal to her earlier ontological commitments.

V. Conclusions

In this chapter, I have reviewed various attempts to respond to Sennet's arguments regarding the difficulty of spelling out a definition, or even an extensionally adequate characterization, of 'unarticulated constituent.' Sennet's counterexamples force the theorist to begin to speak of structures, of occurrences of entities within propositions and of phrases within sentences. Doing so also necessitates a notion of correspondence. We have looked at two attempts to elaborate such a notion: on Stephen Neale's explicit approach, correspondence is restored at the level of some counterfactual utterance; on Millikan's view, correspondence is ensured first at the sentential level, and only derivatively for terms. However, her highly irregular ontology makes it difficult to pursue this position to the end.

In the end, I have claimed that I can offer a descriptively adequate definition of 'unarticulated constituent' without assuming correspondence of parts, and with minimal structural assumptions:

(UC2) X is an unarticulated constituent of a representation R just in case there is a (possibly improper) part of R that expresses a proposition that predicates F of X , but *that part* does not have any parts that semantically encode predication of F to X .

The only correspondence here is between a clause (or a sub-clause) and the proposition it maps onto. I will adopt this definition going forward.

With this definition in hand, in the next chapter I shall proceed to examine various attempts to *explain* the phenomenon of unarticulated constituents.

2. Totally Unarticulated Constituents

I. Degrees of Unarticulation

Unarticulated constituents exist when there is a lack of correspondence between a vehicle of representation and the state of affairs represented. In particular, I have argued that with unarticulated constituents, there is a part of the vehicle that encodes predicating F of X (or that is used to predicate F of X), but there is no part of that part that encodes reference to X. In the case of linguistic utterances, the phenomenon occurs whenever a speaker uses a sentence S (or a sentential part of S) to predicate F of X, but the surface structure of S (or the relevant part of S) does not contain any parts that encode reference to X.

Here, we have characterized unarticulated constituents by attending to the “endpoints” of the process of utterance production. At the origin, we have the *message* – the thought that the speaker intends to communicate to her addressee. (Following the Language of Thought (LOT) hypothesis,²⁷ I will assume that a thought is a structured mental representation.) And at the end, there is the perceptible behavior that is produced in service of performing the speech act. We have observed that unarticulated constituents involve a lack of correspondence between these two points, but, since most theorists recognize a number of representational levels between these, there is strong debate about where exactly the under-articulation gets “resolved.” For it is possible that one of the intermediate stages of linguistic communication *does* have a part that refers to some entity, X, even though its predecessor does not.

²⁷ Cf. (Fodor 1975).

Thus, theorists posit different *degrees* of unarticulation – thinking about the hearer’s perspective, some argue that the unarticulation is resolved at a rather shallow level in interpretation, while for others it is resolved later on (or, perhaps, not at all). Here, I will discuss various positions on the degree of unarticulation that exists.

A. Logical Form

Jason Stanley advocates the radical position that unarticulation is fully resolved in the syntax, specifically, at the level of “logical form.” The opening sentence of his book *Language in Context* states its main thesis: “all truth-conditional effects of extra-linguistic context can be traced to logical form”(2007, 30).

It is difficult to get a precise reading on what Stanley means by “logical form.” In “Context and Logical Form,” he writes: “the Logical Form of a sentence is something like the ‘real structure’ of that sentence,” where this structure is revealed through empirical inquiry. He contrasts this with a *revisionary* conception of logical form as providing a regimentation of ordinary language that eliminates its imperfections. Thus, one may infer that Stanley has in mind either what syntacticians would call deep structure, or LF.²⁸ He also cites Harman (1972), who defends the identification of deep structure and logical form.

What is important about Stanley’s view, in any case, is that unarticulation is resolved in the *syntax*. That is, although they are not visible in the surface structure, every case of unarticulated constituents

²⁸ There is disagreement about whether LF is a distinct level of representation.

involves unpronounced, hidden variables at some deep syntactic level: “What I will defend is the claim that all truth-conditional context-dependence results from fixing the values of contextually sensitive elements in the real structure of natural language sentences”(31).

This is an extremely bold position, as Stanley acknowledges. If even one genuine counter-example is produced, the claim is refuted. Nevertheless, Stanley relies on two main arguments in defense of his view.

The first line of argument involves the observation that unarticulated constituents can be given a dependent-variable-style reading, as in:

(1) Every time John lights a cigarette, it rains.

With this sentence, the typical intended meaning is that every time John lights a cigarette, it rains *wherever John is at that time*. In other words, the location of the rain seems to depend on the quantifier *every time John lights a cigarette*, as in:

(1*) For every time t such that John lights a cigarette at t , it rains at t at *John's location at t* .

But how could a genuine unarticulated constituent acquire a bound variable reading? Stanley argues that an intentionalist cannot account for this, because bound variables are theoretical posits known primarily to specialists, but one needn't be a specialist to utter (1) with the dependent reading. Thus, since for the intentionalist, reference is determined by the speaker's intentions, they cannot posit that speakers are really referring to bound variables, since many speakers lack the very concept.

How, then, can an intentionalist account for a bound reading, as in (1)? Stanley claims that without the presence of a covert variable in the syntactic structure, we have no understanding of how a bound reading could be achieved.

This argument is deployed in a number of contexts by Stanley. It is an instance of a broader class of criticisms of intentionalist theories known as *meaning-intention* problems, in which the critic argues that the intentionalist is offering an explanation of linguistic facts by positing entities that ordinary speakers lack a concept of, thus conflicting with their commitment to the primacy of speakers' intentions as the source of meaning. The meaning-intention problem was famously deployed by Stephen Schiffer against the hidden indexical theory of attitude ascriptions, and it is a pressing challenge for the intentionalist. I will have more to say on the meaning-intention problem later on.

Stanley's second primary form of argument against the defenders of free enrichment as an explanation of unarticulated constituents is that the notion of free enrichment lacks explanatory power. Taking the word "free" literally, it seems to overgenerate readings. Thus, "free enrichment" must somehow be constrained – but such constraints have never been produced.

Consider, for instance, the following kind of sentence:

(2) Everyone likes Sally,

with the intended reading being equivalent to:

(3) Everyone in my class likes Sally.

The proponent of free enrichment will say that the basic propositional form encoded by (2) is pragmatically “enriched” so as to be interpreted as equivalent to (3). Cashing out the metaphor in the terms of Relevance Theory, if we assume that the encoded meaning of *everyone* is simply the unrestricted universal person-quantifier, then the hearer will decode that meaning and find that it fails to achieve his expectation of Relevance, and thus the propositional “skeleton” will be modulated in such a way that *everyone* is read as restricted to the people in the speaker’s class (presumably, the topic of conversation).

Thus, on a contextualist view, the linguistically encoded meaning is often modulated, even before the processing of implicatures takes place (i.e., at the level of “what is said,” or explicature). And this process is supposed to be “free” to the extent that the ways in which meanings are modulated is not mandated by any linguistic property, but rather is simply dependent on the speakers’ communicative intentions. Similarly, for hearers, the presence of such modulation is not signaled by anything syntactic, but is guided by their search for Relevance, or some other pragmatic mechanism.

Stanley argues that if such a process really existed, it would vastly overgenerate readings. For instance:

4. Everyone likes Sally.
5. Everyone likes Sally and himself.
6. Everyone likes Sally and his mother.

If the encoded propositional forms can be modulated and enriched in the way posited by the explanation of (2), why can (4) not be enriched to mean something equivalent to (5) or (6)? It does not seem that such readings are available (except, perhaps, by implicature, which is distinct from free enrichment). So, it must be the case the free enrichment is not *entirely* free, but instead is somehow constrained. The problem is, without an account of these constraints, appealing to free enrichment seems completely *ad hoc*, and no one seems to have much of an account to offer.

Relevance Theorists claim that explicatures must be a “development of the logical form” encoded by the sentence uttered. Thus, presumably, they would argue that (3) is an acceptable development of the logical form encoded by (2), but (5) and (6) are not developments of the form encoded by (4). However, a precise account of the notion “development of the logical form” has not been offered, and thus it is vulnerable to the same criticism stated above.

Stanley’s position, then, is that there is no tenable alternative to the syntactic view of unarticulated constituents, since the main competitor, free enrichment, is found to be non-explanatory. First, it fails to provide an account of bound-variable readings of unarticulated constituents. Second, it demands an account of the constraints that exist on free enrichment, on pain of over-generating readings, but none have been formulated. Since (a) many people grant the existence of phonologically null syntactic elements (covert variables) in general, (b) we have an explanatory account of bound variable readings as actually involving bound variables, and (c) the alternative free enrichment account lacks explanatory power, Stanley concludes that the radical position that all context-sensitivity can be traced to elements in logical form is the correct one.

B. Free Enrichment

The main competitor to Stanley's view in the literature is that unarticulation is resolved through free enrichment – *non-mandatory* modulations of encoded meaning that affect central truth-conditional meaning. According to many contextualist theorists, syntax encodes a set of semantic properties, which specify a proto-propositional “skeleton” or “blueprint,” that frequently (or perhaps always) fails to specify truth conditions. Once this propositional blueprint has been derived, the hearer begins to recover the speaker's intended meaning by modulating the blueprint and searching for implicatures.

The process of free enrichment was not invented simply to provide an analysis of weather reports. Rather, it is motivated by consideration of a wide variety of linguistic data that prompt the observation that what the speaker directly, centrally means often seems to depart (sometimes dramatically) from what appears to be the encoded meaning of the sentence. Furthermore, it departs in ways that seem fairly unsystematic and not constrained by linguistic material, but rather depend simply on the speaker's communicative intentions.

Unfortunately, as we have seen, the very feature that makes free enrichment a powerful strategy when it comes to explaining linguistic phenomena also threatens to rob it of explanatory power – the more “free” it is, the wider variety of phenomena it will be able to apply to, but also the more *ad hoc* it will appear when trying to prevent over-generating readings. This is a serious challenge to the proponent of free enrichment.

What about Stanley's argument about bound-variable readings? In response to this line of argument, Recanati considers two conceptions of free enrichment – “semantic” and “syntactic.” On the

semantic conception, “free enrichment has been construed as a mechanism through which some constituent in the truth-conditional interpretation of an utterance is contextually supplied without being articulated in the sentence.” (Confusingly, Recanati uses the term *interpretation* to denote the mind-independent state of affairs/truth conditions themselves.) The bound-variable reading is a problem for the semantic conception of free enrichment because it is mysterious how an entity like a bound variable could get itself inserted into the represented state of affairs if it isn’t articulated syntactically.

On the *syntactic* conception of free enrichment, however, it is not the state of affairs itself that is “enriched” (relative to the sentence uttered), but rather an enriched “syntactic” representation is derived, which *then* is assigned truth conditions.

On the syntactic conception, free enrichment determines aspects of the representation which is interpreted: it contributes further symbols, further representational elements, which are unarticulated in the sense that nothing corresponds to them in the natural language sentence that has been uttered. The output of this process of free enrichment in the syntactic sense is a mental representation which articulates what the speaker means by his utterance, including those aspects of the speaker’s message that are not articulated in the natural language sentence she uses.

This is supposed to provide a response to the bound-variable problem, since if what is produced by free enrichment is an enriched *logical form* (i.e., syntactically articulated), then there should be no problem with producing a logical form that contains an “extra” bound variable.

There are a number of problems with Recanati's response. First of all, it isn't clear who he has in mind as defending the semantic view of free enrichment. As he notes, the syntactic conception is held by an "overwhelming majority of TCP-theorists"(340). Indeed, the semantic conception seems barely coherent. For if we think of free enrichment as a *pragmatic* mechanism, then, presumably, it is supposed to be dependent on the speaker's communicative intentions. But in that case, we would expect there to be a level of representation, somewhere in the speaker's mind, where the unarticulation is resolved – for if the unarticulated constituent is unarticulated even with respect to the speaker's *thought*, it seems mysterious how the process could be under voluntary control, as suggested by the term *free* in *free enrichment*. (Indeed, Recanati describes the semantic conception as one in which the unarticulated constituent is *contextually supplied* – this can charitably be read as meaning *supplied in context*, rather than the pernicious *supplied by context*, but the ambiguity is potentially revealing.)

Thus, the semantic conception of free enrichment actually seems closest to what is known as the "implicit approach" to unarticulation/underspecification. On this view, the unarticulation is, effectively, *never* resolved – instead, the intuitively desired truth conditions are achieved by implicitly restricting the situation or domain relative to which the statement is evaluated. But it is counter-intuitive to think of this mechanism as a process of free enrichment, since there is nothing that is being "enriched" and, if the unarticulation exists at the level of thought, it is mysterious how it can be voluntarily controlled. Thus, it seems to be neither "free" nor an "enrichment," although one is cautious of putting too much stock in a label.

A further problem with Recanati's response to the bound-variable problem is that it does not really seem to address the main worry. Even if, per the syntactic approach, we grant that the process of free enrichment produces "linguistic" representations (in some representational medium), if we maintain that this is still determined by the speaker's intentions, the main issue arises again – how can a speaker's intentions provide a bound variable as the value of some expression, if the speaker does not have any concept of a bound variable?

Recanati writes: "Again, the variables in question will be found in the representations delivered by pragmatic processes as they apply to natural language structures; they will not be found in the natural language structures themselves, hence they will not be articulated in the relevant sense"(340). But an intentionalist will hold that the speaker's intentions constitutively determine what is meant on any given occasion of use, and this applies to the explicitly articulated meaning as well.

Thus, upon reflection, it appears that Stanley's argument against free enrichment for bound variables should also be a concern for the intentionalist for *articulated quantificational structures* as well. There is nothing specifically troublesome about the free enrichment part – if a speaker can refer to X, then it seems that she should be able to do so either implicitly or explicitly. For if speaker reference is, following Stephen Neale, a matter of speaker-meaning a singular proposition, then what matters is the thought the speaker is trying to convey – whether she conveys that thought explicitly or implicitly depends on her current interests and goals.

Thus, if Stanley is simply raising an instance of the meaning-intention problem, the core issue for the intentionalist is really how speakers manage to "refer" to bound variables (or, more properly, to express thoughts containing bound variables as parts) even if they lack the concept of a bound

variable. So, we can effectively push the issue back into the broader set of issues surrounding the meaning-intention problem. This is not to provide a resolution to the problem, but rather to note that there is nothing especially problematic about *implicitly* referring to a bound variable, and therefore, nothing especially troubling for the proponent of free enrichment for bound variables.

C. No Resolution

We have observed a few different degrees of unarticulation: (i) a sentence might be underarticulated relative to its logical form; (ii) the logical form might be underarticulated relative to the interpretation (thought) that is arrived at by the hearer. A more radical view of unarticulation, defended by John Perry, is that even *thoughts themselves* may be underarticulated relative to the state of affairs they represent. On this view, the unarticulation is *never* resolved – there remains a mismatch between vehicle and content at every level of representation.

But how could a thought have an unarticulated constituent? Perry motivates this possibility by way of a persuasive thought experiment. It concerns a community of speakers who inhabit a small, isolated island known as Z-land. They have no contact with the outside world and never leave Z-land. Furthermore, Z-land is small enough that whenever it rains, it rains everywhere on Z-land (i.e., it doesn't occur to them that it might be raining on the eastern coast of Z-land, but not on the western coast). Now suppose one of these Z-landers says:

(7) It's raining.

Intuitively, the speaker's utterance should be judged true or false depending on whether or not it's raining *in Z-land*. But the speaker didn't say "Z-land"! Thus, we seem to have a case where a speaker has expressed a thought whose content involves Z-land (in some sense), even though the speaker has not explicitly referred to Z-land.

What makes the Z-lander example especially problematic is that there seems to be good reason to doubt that the location of Z-land is specified in the thought the speaker expresses. As I have alluded to, by hypothesis, it simply never *occurs* to them that rain is something that takes place at different locations. For Z-landers, if it's raining, it's raining in Z-land, full stop. Thus, attaching a locative modifier to the verb *rain* would be as strange as an ordinary English speaker saying something like:

(8) Triangles have three sides in America.

Although (8) could certainly be used in a joking context, in ordinary situations, the modifier *in America* would be a pointless add-on that would violate the hearer's expectations of Relevance, or, in Gricean terms, the Maxim of Manner. Similarly, it would be very surprising to find that when people have the thought *triangles have three sides*, the Mentalese sentence contains a locative modifier. By parity of reasoning, it would be surprising to find that when Z-landers have the thought *It's raining*, the Mentalese sentence contains a modifier like *in Z-land*.

Thus, Perry's Z-land example, although a thought experiment, presents a strong challenge for Stanley's syntactic view. For when Stanley writes, "all truth-conditional effects of extra-linguistic context can be traced to logical form," this is not presented as a mere contingency, but is thought to express something like a law or a linguistic principle. Thus, even hypothetical counter-examples

should be taken seriously. And the Z-land case presents us with speakers who seem to be expressing thoughts *about* Z-land, even though we have reason to doubt that Z-land is articulated even at the level of thought (and certainly not in the syntax).

In Perry's terms, Z-land weather thoughts and utterances *concern* Z-land, although they are not *about* Z-land. Thus, Perry seems to grant the existence of two different forms of intentionality –*aboutness*, and *concerning*.

D. Summary

We have looked at three different approaches to the question of where in the communicative process unarticulation gets resolved: first, Stanley's syntactic view, which claims that it is resolved in the linguistic logical form; second, the view of Recanati and many other pragmatists, that unarticulation is resolved at the level of communicative intentions; and, finally, Perry's view, motivated by the Z-land example, that even thoughts can have unarticulated constituents. Is Perry's thought experiment strong enough to refute the syntactic approach? In the next section, I will argue that it is not a mere thought experiment, and that there are many cases that cast doubt on Stanley's bold hypothesis.

II. Totally Unarticulated Constituents

So as to free the concept from its origin in Z-land, I will use the term *totally unarticulated constituent* (TUC) to refer to an unarticulated constituent that is unarticulated relative to every level of

representation. I will argue that TUCs are (a) clearly attested in reality, and (b) fundamental to primitive representational systems.

A. TUCs in Animal Signals

Although it is rarely cited in relation to the current debate, if we accept that animal signals have representational content of some kind, it follows immediately from the observation that most or all animal signals lack meaningful syntax, that TUCs are rampant in animal signaling.

Consider, for example, the celebrated vervet monkey's alarm call. Upon seeing a predator, the vervet monkey typically produces a call, presumably with the function of alerting his conspecifics to the presence of a threat. Interestingly, the sound of the call varies in relation to what kind of threat the monkey perceives. Many of them have three or four calls that are differentiated depending on whether the threat perceived is aerial (hawks), on the ground (jaguars), or in the grass (snakes). Similarly, the monkeys who hear these calls react in a way that is appropriate to the kind of threat perceived – running up a tree to avoid a snake, or hiding in the bushes from the hawk.

Intuitively, one is inclined to say that when a vervet monkey produces the hawk-specific call, that, in some sense, it *means* that there is a hawk. But how can animal signals mean things? And what kinds of things do they mean?

i. Natural and Non-Natural Meaning

Grice began the contemporary discussion of such issues with his introduction of the distinction between *natural* and *non-natural* meaning. Natural meaning is the kind of meaning found in a sentence like,

(9) Those tree rings mean that the tree is 60 years old.

Natural meaning is information that is “out there” in the world. Also, it is factive, in the sense that (9) implies,

(10) That tree is 60 years old.

In contrast, non-natural meaning is *speaker's meaning* – it is constituted by the presence of a special kind of higher-order intention (communicative intention). Also, it is non-factive, in the sense that,

(11) By giving the thumbs up, Alice meant that she was having a good time.

does not imply,

(12) Alice was having a good time.

One major problem with Grice's distinction between natural and non-natural meaning is that the terms suggest exhaustiveness (since *non-X* plausibly expresses the dual of *X*), but animal signals do not seem to fit nicely into either category. First, they don't seem to count as non-natural meaning,

since it is implausible to suppose that most (or any) animal signals are produced by complex higher-order intentions. Thus, we might expect them to fit in better with the tree rings and such, on the natural meaning side (I mean, they're both in nature, after all). However, unlike natural meaning, animal signaling is, importantly, *non-factive*. Once we grant accuracy conditions, or some kind of representational content, to the vervet monkey's *hawk*-call, then if the call is produced in the absence of a hawk, we should conclude that the call *misrepresents* the world.

What becomes, then, of Grice's distinction? Although it was a useful initial description of the subject matter, it ultimately should be superceded. First, the notion of natural meaning should be identified with the much clearer and more precise notion of "Shannon information." Shannon information, or, "information" as defined by Claude Shannon, refers to a *reduction in uncertainty* provided by the occurrence of a given event regarding the occurrence of some other event. In probability terms, some event E carries information about some other event E' just in case the conditional probability of E' given E is greater than the probability of E' alone (i.e., $P(E' | E) > P(E')$). Thus, clouds carry information about the rain because the presence of clouds makes it significantly more likely that it will rain.

There is a debate in the literature about whether information can be inaccurate. For it seems that if an event E raises the probability of some other event E' even *slightly*, E can be said to carry *some degree* of information about E'. But we can't infer from E that E' actually occurs, for suppose $P(E') = 0.5$, and $P(E' | E) = 0.6$. Then, although E seems to carry some information about E', E' is still only 60% likely to occur, given E. What shall we say in the cases where E' fails to occur? Did E still carry information about E'? Does E still *mean* that E' will occur?

A simple way to resolve this largely verbal issue is to claim that Shannon information is in fact immune to error. However, the propositional content of the information is not simply the proposition that, say, E will occur. Rather, following the above example, we shall say that the information that is carried by E is *that E' is 20% more likely to occur*. In other words, if we assume that the content of the information itself contains a probability, then we can maintain both that (a) Shannon information is just “pure” information floating around “out there,” incapable of misrepresenting, and, that (b) Shannon information is not factive in the sense that *E carries information about E'* does not imply that E' occurs.

On this proposal, *natural meaning* should really be taken to be something like *natural suggesting*, as in:

(13) Those clouds suggest that it will rain.

(13) would be problematic for Grice, since (a) it appears to be a case of natural meaning, but (b) it would seem that (13) could be true even though it doesn't in fact rain. (Hence, the use of the weaker *suggest*.) But that means that it is non-factive. So we have apparently discovered another case of non-factive natural meaning.

Grice's initial claim that natural meaning is factive was based on the intuition that a sentence like the following is inappropriate:

(14) #Those clouds mean that it will rain, but it's not going to.

Many speakers, myself included, find this intuition compelling. What if we try to repair the sentence?

(15) #Those clouds mean that it will probably rain, but it's not going to.

(15) seems unacceptable as well. But why should this be? *It's not going to rain* is not inconsistent with *It will probably rain*, thus the sentence does not conflict with the assumption that natural meaning is factive. Shall we conclude, then, that it is *super-factive*, in the sense that not only must the dependent clause express a true proposition, but it must be a proposition with very high probability as well?

There is a much simpler solution. First, observe that a similar issue arises with *suggest*, which we have already seen is non-factive:

(16) ?Those clouds suggest that it will rain, but it's not going to.

The problem is really that these utterances conflict with the knowledge norm of assertion, namely the conversational principle that, in general, a speaker should only assert *p* if she *knows* that *p*.

Williamson defends the knowledge norm by way of modified Moore's Paradox examples, such as:

(17) #It will rain tomorrow, but I don't know that it will.

Whereas Moore's Paradox illustrates the infelicity of uttering *p, but I don't believe that p*, Williamson's examples persuasively show that the exact same phenomenon occurs with *p, but I don't know that p*.

If we assume that *S knows that p and q* entails *S knows that p* and *S knows that q*, then we can easily uncover a conflict with the knowledge norm. Let us suppose that S has uttered (16). Then:

- (i) If S is trying to obey the knowledge norm, then S believes that S knows that those clouds suggest that it will rain, and S believes that S knows that it is not going to.
- (ii) But (assuming S is rational) if S believes that S knows that it is not going to rain, then S must think that there is little to no chance that it will rain.
- (iii) However, if S believes that S knows that those clouds suggest that it will rain, then S has reason to believe that there *is* a non-negligible chance that it will rain.
- (iv) Therefore, if S is following the knowledge norm, then S believes that there is little to no chance that it will rain, and S believes that there is a non-negligible chance that it will rain.
- (v) Therefore, either S is not following the knowledge norm, or S believes a contradiction.

The same kind of reasoning can be used to explain why it seems plausible that natural meaning is factive (even if it is not). For we infer that natural meaning is factive on the basis of examples like:

(14) #Those clouds mean that it will rain, but it's not going to.

Those spots meant measles, but he hadn't got measles.

The recent budget means that we shall have a hard year, but we shan't have.

Those spots meant he probably had measles, but in fact he didn't.

But the most one can infer from this is that natural meaning is "probably-factive," since (15) is just as bad as (14):

(15) #Those clouds mean that it will probably rain, but it's not going to.

Another complication is that Grice's examples involve singular reference (*those clouds*), whereas other natural meaning statements, such as (17) and (18) do not:

(17) Clouds mean rain.

(18) Smoke means fire.

These sentences seem true, even though the presence of clouds does not perfectly predict the presence of rain. So, generic natural meaning statements are not factive in that universal sense. But this may be due to the fact that they are generics, and not universal statements.

We seem to be getting at a type/token distinction in natural meaning. *Clouds mean rain* is true in virtue of the physical correlation between cloud events and rain events. But *That cloud means rain* seems to be in virtue of the fact that *Clouds mean rain*. So, the token cloud event means that a token rain event will occur in virtue of the natural meaning properties of the type *cloud event* (specifically that cloud events naturally correlate with rain events). And, if we follow Grice, it must also be true that the rain event will in fact occur.

Grice notices that natural meaning statements can be relativized to an agent, as in:

(19) Those spots didn't mean anything to me, but to the doctor they meant measles.

This raises an interesting question: can a given (particular) event naturally mean that p to one person and naturally mean that q to someone else? If so, need both p and q be true?

(20) Those spots meant I had the plague to me, but to the doctor they meant measles.

(21) Those spots meant that I would be missing school to me, but to the doctor they meant measles.

I do not have strong intuitions about the acceptability of (20), or (21), but they seem relatively fine to me. Of course, (20) would present a direct counter-example to Grice's original characterization of natural meaning. I will leave it to the reader to adjudicate the matter according to his or her own judgment. However, there are other cases that, although they involve some non-literality, also seem to contradict the factivity of natural meaning:

(23) For Bob, it meant the end of the world, but we all knew he was overreacting.

(24) For Alice, her wedding meant that she would be happy forever, but she was just being naïve.

I claim that the reason natural meaning statements appear to be factive is that they are effectively equivalent to Moore's Paradox statements. This can be shown via a set of transformations that illustrate highly plausible inferences:

(22a) The recent budget means that we shall have a hard year, but we shan't have.

(22b) The recent budget means *to me* that we shall have a hard year, but we shan't have.

(22c) I infer from the recent budget that we shall have a hard year, but we shan't have.

(22d) I believe that we shall have a hard year, but we shan't have.

(22b) is derived from the assumption that natural meaning must mean something *to someone*, and the default agent, when not otherwise specified, is the speaker. (22c) is derived from the assumption that if some event E naturally means that *p* to S, then S knows, or believes, that there is a correlation between E and *p*, which provides good reason for S to infer that *p*. Therefore, if E naturally means that *p* to S, then S should infer that *p* upon perceiving E. Finally, (22d) is derived from the entailment from *inferring p* to *believing p*. And (22d) is a straightforward Moore's Paradox (with the asserting and doubting clauses being in the opposite positions compared to the canonical formulation).

So, I conclude that we can explain why sentences like (14) are infelicitous without assuming that natural meaning statements are factive. Namely, it is because they strongly imply a Moore's Paradox statement.

In any case, exegetical questions to the side, I do not consider natural meaning to be an important explanatory concept, and I recommend dropping talk of natural meaning, except as shorthand for statements about Shannon information.

ii. Animal Signals

To explain how animal signals can have representational content, and what that content is, we must look beyond the resources of Gricean theory and intention-based semantics. But first, why should we think that animal communication is representational at all?

A skeptic might argue that when we ascribe meaning to animal communication, we do so in a merely pretend way, or by analogy. For instance, if a dog is barking excitedly, the owner might say:

(25) That means that Fluffy is ready to go to the store and she wants me to put her leash on.

Uttering (25) does not imply that the speaker thinks that Fluffy really has a concept of the store, or a leash, etc. Rather, it is a kind of dramatic portrayal of Fluffy's state of mind, but it is not intended to match its intentional content. Similarly, we can easily adopt the "intentional stance" towards even inanimate objects, as in:

(26) My computer wants me to enter my password.

So, perhaps when we discuss vervet monkeys as making calls which "mean *hawk*," we are merely adopting the intentional stance towards the behavior, and it should not be taken as literally meaning anything.

This skeptical position does not appear justified. The suggestion is that one is *merely* adopting the intentional stance towards certain kinds of behavior, but one is doing so with knowledge that it is just a kind of pretense. However, this fails to hold up if the "pretense" figures in some explanatory theory. For, if adopting the intentional stance is explanatory and predictive, then it is not merely a pretense to do so, but is considered "real."

Teleosemantics provides us with a framework within which attributing meaning to animal signals has explanatory virtue. Thus, if teleosemantics is plausible, it is plausible that animal signals have representational content.

Teleosemantic theories are those that hold that the representational content of a representation is determined by its *proper function*. Within the class of *biosemantic* theories, proper function is characterized as follows: An organism's trait/behavior X has the proper function to do F just in case doing F explains (in accordance with evolutionary theory) the continued reproduction and proliferation of organisms that have X. For instance, the heart both circulates oxygen and produces a noise. However, it is highly plausible that the heart was *selected for* its ability to circulate oxygen, but not for its tendency to produce a noise. For circulating oxygen is an adaptive trait that would give its possessors an evolutionary advantage, thus proliferating the number of organisms with hearts, while the same cannot be said of the heartbeat.

Teleosemantic theories employ an *etiological* notion of function (Wright 1973), according to which the function of some organism's trait depends on the evolutionary history of the species. Some may question whether such a concept of function is explanatory. First, one might think that in order to understand something's function, you need only look at the organism that possesses it and how it interacts with other parts of the organism. This is *function* in the sense of "functionalism," or conceptual role semantics. Second, and related, one might question whether etiological explanations are causal. For one might think that the causal properties of an organism, and its traits, should supervene on the physical properties of the organism itself, and not its evolutionary ancestors.

Larry Wright, who introduced the etiological notion of function, responds to these worries by claiming that etiological explanations should be thought of as causal in an “extended sense.” “And this is indeed what I wish to argue: functional explanations, although plainly not causal in the usual, restricted sense, do concern how the thing with the function got there. Hence they are etiological, which is to say ‘causal’ in an extended sense”(1973, 156). In other words, etiological explanations do not explain the causal properties of any individual thing; rather, they explain why things of that type exist (in the frequency that they do). To say that a function of the heart is to circulate oxygen is to say that the fact that hearts circulate oxygen helps explain why there are so many hearts around in the world. It is hard to deny that such attributions have some explanatory value.

On the other hand, if, as Jerry Fodor claims, the theory of natural selection is itself non-explanatory, then it would follow that such attributions of biological function would be incorrect. However, I will not attempt to respond to Fodor’s criticisms of natural selection as a theory in this essay.

Let us now turn back to the vervet monkeys. Given the preceding discussion, to say that the function of a vervet’s “*hawk-call*” is to alert its conspecifics of the presence of hawk, is to say that alerting conspecifics of the presence of a hawk is an effect that that type of call has historically had, which promoted the fitness of that trait’s possessor, and which explains why so many monkeys act in that way.²⁹

In Millikan’s theory, as previously discussed, representations are always midway between a producer and a consumer. Something gains representational status when it acquires the function of being used

²⁹ “Alert” can have an intentional reading, but what I mean by “alerting its conspecifics of the presence of a hawk” is really, “prompting behavior that is adaptive in that situation, given the presence of a hawk.”

by the consumer to adapt its behavior in a way that is appropriate to some state of affairs, as “perceived” by the producer, in accordance with a biologically Normal explanation. The state of affairs that the consumer is supposed to adapt to is captured by the representation’s *semantic mapping function*.

Granting that animal signals have representational status in accordance with the above explanation, one might still doubt whether what animal signals represent is *propositional*. Similarly, one might have the intuition that although a vervet alarm call may be *accurate* or *inaccurate*, it cannot be *true* or *false*. Is there any theoretical grounding for such intuitions?

There are good reasons to want to reserve the term *true* and *false* for those representations that express propositions. (In particular, since we *do* apply those terms to propositional representations, such as natural language sentences, we may wish to restrict its usage to such cases, so as not to equivocate.) But what important difference could there be between my thought *that there is a hawk present*, and the vervet monkey’s *hawk-call*?

Well, for one thing, my thought that there is a hawk present is, in virtue of being a thought, part of an inferential network, and is composed of concepts that appear in other thoughts as well. Because of that, thoughts are able to achieve an incredibly fine level of grain. It is doubtful that primitive, atomic representations can do the same. For example, when we try to “translate” the vervet’s call into English, we might arrive at any of the following:

(27) There is a hawk present.

(28) Hide under the bushes!

(29) Look out, I see a hawk!

(30) There is a flying predator.

Each of these seems to be a reasonably fair “translation” of the vervet call. But in English, these sentences do not have the same content as each other. Indeed, some of them are in the indicative mood, thus capturing the indicative side of the *pushmi-pullyu* nature of the vervet call, while others are in the imperative mood.

It does not, in general, seem possible to give a uniquely correct translation of an animal signal into English. This is because the contents expressible with English sentences are extremely fine-grained, whereas I claim the contents of animal signals are generally quite coarse-grained (relative to English). This might not seem like a problem since, one might think, what you can do with a coarse-grained brush you can do with a fine-grained one, but that is not so, for the coarseness also introduces a certain amount of indeterminacy (relative to English). Ultimately, the representational content of an animal signal depends on its evolutionary history, in particular, the effect it has that explains (historically) the proliferation of the trait. But historical explanations admit of degrees and borderline cases. Being articulated in an inferential network can help reduce the indeterminacy because, in effect, the indeterminacy is reined in by the inferential role of the representation in the whole system. But vervet calls, and indeed most animal signals, are not articulated into an inferential network. Therefore, they are bound to be highly indeterminate in comparison to English, and this indeterminacy simply cannot be fully resolved.

Thus, I claim that while animal signals and human thoughts are both representational, the former represent non-propositional states of affairs, while the latter represent propositional states of affairs.

(If it helps, one can think of non-propositional states of affairs (that is, states of affairs individuated in a non-propositional way), as fuzzy sets of “real” (propositional) states of affairs, although I myself don’t find this helpful.) Similarly, I will reserve the terms *true* and *false* for propositional representations, and use *accurate* and *inaccurate* (*veridical* / *non-veridical*) for non-propositional representations.

iii. Totally Unarticulated Constituents

I have argued that we ought to attribute a coarse-grained kind of representational content to animal signals, in accordance with teleosemantics. Granting this claim, it is easily observed that animal signals exhibit a phenomenon very much like the unarticulated constituents that exist in English speech. For example, let us consider once more the *hawk*-call. It is reasonably translating as meaning:

(31) There is a hawk!

Note, however, that implicit in this translation is the fact that there is a hawk *here* and *now*.

(32) There is a hawk here, now!

In other words, the state of affairs represented by the *hawk*-call, when described in English, appears to pertain to the time and place of “utterance.” And, indeed, this is how it should be translated. For the *hawk*-call semantically maps onto a state of affairs in which there is a hawk present at the local time and place. It would not be judged accurate in virtue of there having been a hawk here yesterday, or a hawk on the other side of the mountain now. Similarly, the behavior that the call Normally

produces is adaptive given the presence of a hawk at the time and place of utterance, and not given that a hawk will be present two weeks from then.

Thus, there is a strong sense in which animal signals *concern* the local time and place. However, such entities are obviously not articulated in the structure of the signal. For almost all animal signals lack meaningful structure (they are, essentially, atomic). Thus, we have a representation that represents a state of affairs in which X has the property F, but there is no part of that representation that encodes reference to X itself. In other words, we have a case of TUCs.

Arguably, we have as strong of a case for claiming that animal signals display the phenomenon of TUCs as we do for thinking that Z-lander thoughts display it. In both cases, we conclude that the representation concerns a given entity by considering the state of affairs that the representation is *used* to orient its consumers towards. Although I have argued that animal signals' contents are more indeterminate and coarse-grained than those expressed in English, it is still sensible to say that they involve specific entities. Perhaps we cannot specify whether it is *here, this place, the presently visible surroundings, where I'm standing*, etc., as the "correct" English translation of the monkey's *hawk*-call. Still, these attempts have enough in common – they all concern *roughly* the current location – that we can say that the location (individuated somewhat vaguely) is a constituent of the state of affairs represented.

In brief, our inability to provide a uniquely correct English translation of the vervet call should not undermine the fact that the call concerns the time and location of utterance, in a partially vague sense. The fact is that *something* (which we can't precisely individuate) is playing a role in the state of affairs represented, but is not itself represented by anything.

Concerning the local time and place is not an idiosyncratic feature of vervet calls, either. Animal signals nearly universally concern the local time and place. First, carnival tricks notwithstanding, it would appear that no animal has ever produced an eternal sentence (such as a mathematical equality). The development of a meaningful signal through natural selection is a rather “costly” process (in the sense that it requires a number of historical conditions to be in place), and so animal signals will tend to be highly adaptive, “useful” behaviors. Furthermore, what is often of greatest concern is the local environment of the organisms. Most animals lack the capacity for future planning or reasoning about the distant future, and *ceteris paribus*, we would expect there to be more natural signs of a given state of affairs in the local environment than in very remote environments, and so it is natural that animal signals will almost invariably concern some relatively local state of affairs. Given this, in conjunction with the assumption that virtually all animal signals lack meaningful parts, we can infer that virtually all animal signals exhibit TUCs.

If this is correct, then we have already demonstrated that Perry’s Z-land case is not a mere thought experiment, for it exists in the animal kingdom. Upon reflection, we see that TUCs are not an esoteric feature of the English language, but are instead a fundamental aspect of primitive representations. This makes a great deal of sense when one considers how primitive representations tend to be used, and the fact that they tend to lack structure.

The foregoing does not immediately entail that Stanley’s syntactic thesis is incorrect. For it is conceivable that TUCs are pervasive in animal signals, but for some reason, totally absent in human speech and thought. However, it certainly casts doubt upon the thesis, for it seems to imply that human speech is somehow impoverished relative to animal signals.

B. TUCs in Natural Language

We may also clearly identify cases of TUCs in natural language. There are at least three kinds of example: historical; developmental; *prima facie*.

i. Historical Cases

There are clear cases where people have been in exactly analogous situations to the Z-landers.

Consider, first, the concept of *weight*. Many people now know that weight is in fact a relative notion.

Weight is a measurement of the force on an object due to gravity. On Earth, we are almost always interested in the gravitational force *from Earth*, so we tend to forget that it is a relative concept.

However, anyone who has seen footage of astronauts bouncing around on the moon probably knows that one's weight would be much lower on the moon.

However, although this fact is widely known today, there was a time when it was not even a hypothesis. Let us assume that that time was just before Newton.³⁰ People were able to think about weight, talk about weight, measure weight, etc., well before Newton proposed the laws of gravitational attraction, and well before anyone would have thought that *weight* really denoted *weight (on Earth)*. Thus, if a pre-Newtonian utters,

(33) This pumpkin weighs ten pounds

³⁰ Even if this is false, since we know definitively that there *has been* such a time (say, in ancient Babylon), it is a harmless pretense.

we can conclude that she has said something concerning weight on Earth, even though no part of the sentence encodes reference to Earth. Furthermore, no part of her *thoughts* about weight refers to Earth, for to a pre-Newtonian,

(34) This pumpkin weighs ten pounds on Earth

would sound as strange as

(8) Triangles have three sides in America.

I.e., it simply wouldn't occur to a pre-Newtonian that weight can be relativized, so it would be implausible to claim that their weight thoughts contain reference to Earth. Thus, we seem to have a very strong case for a genuine historical precedent for natural language TUCs.

Another similar case can be found in the concept of simultaneity. According to contemporary physics, whether two events occur at the same time, or simultaneously, is not an absolute notion, but depends on a temporal frame of reference. If you are traveling in a certain direction at a certain speed, two events may be simultaneous to you that are not simultaneous to me. Thus, statements that events X and Y “took place at the same time” or “were simultaneous” must be interpreted as meaning “were simultaneous relative to the default frame of reference.”

However, importantly, (a) many people simply *don't* know this – they have never encountered relativity physics; (b) many people know *about* this, but don't really have a grip on the concepts

involved (I would count myself among this group); and (c) there was a time in recent history when nobody knew this. Nevertheless, none of this stops people from making claims about simultaneity – claims which we must interpret as containing a temporal frame of reference as a TUC.

I conclude that undeniable historical cases of TUCs can be found. Thus, Stanley’s thesis is refuted. It is not true that “all truth-conditional effects of extra-linguistic context can be traced to logical form,” for we have found cases where extra-linguistic context affects truth conditions in a way that cannot be attributed to logical form (unless one makes the ridiculous assumption that there was a covert variable in *simultaneous* all along, even before relativity theory was discovered, thus implying that syntax is omniscient).

ii. Developmental Cases

What’s true of the development of humanity applies, in this case, to the individual as well. Namely, individuals acquire concepts that are relativized to some other entity X before they have acquired a concept of X. In such cases, application of the concept expresses an X-involving state of affairs, but no part of the thoughts containing the concept refers to X.

First, this is true of the weight and simultaneity cases as well: people learn to apply the concept *weight* before they have learned about the difference between weight and mass. And most people use *simultaneous* without ever mastering the concept of a frame of reference. Even those who do are able to apply the term meaningfully before studying Einstein.

From a developmental standpoint, it is reasonable to assume that relativized concepts are frequently acquired before the concept of the thing to which they are relativized. Consider, for instance, the predicate *delicious*. Although many people think that *delicious* might express a slightly different property to one person than it does to another, a child might acquire the word thinking simply that whatever is delicious *to me*, is delicious. It would take some time for the child to learn that Brussels sprouts might be delicious to Mommy but not to herself.

For another example, consider time reports:

(35) It's four o'clock in the afternoon.

Everybody knows that (35) must be relativized to a time zone. But it is doubtful that we know this as soon as we gain competence with sentences like (35). The understanding that clock time can differ from one region to the next is a relatively sophisticated intellectual achievement compared to the mere ability to "tell the time" from a clock. Thus, all else equal, we would expect the latter to precede the former from a developmental point of view. This provides further evidence that TUCs are a prevalent feature during conceptual development.

iii. *Prima Facie* Cases

Finally, I wish to raise a class of examples for which I have no strong proof that they exhibit TUCs, but for which there is a strong *prima facie* case to be made. Consider (35) once again:

(35) It's four o'clock in the afternoon.

Now, unlike relativity theory, the fact that there are different time zones is almost universally known amongst adults. Thus, it is not hard to realize that (35), uttered in New York City, must be interpreted as equivalent to:

(35') It's four o'clock in the afternoon (Eastern Standard Time).

However, although this fact is rather easily accessible, I claim that in many ordinary cases, *speakers simply ignore it*. In other words, the time zone is totally unarticulated, even though the speaker knows that the concept being employed is a relativized one and possesses a concept of its relatum.

Consider, for instance, a person who never travels, doesn't talk to people in other time zones, and generally keeps to himself in the house. Now suppose he looks at the clock and thinks, *It's four o'clock in the afternoon*. Why must we assume that the structure of his thought is really something like, *It's four o'clock in the afternoon in the current time zone*? Is it really plausible to hold that *in every case* where people think about what time it is, they also token a concept of the time zone? I think not, although I recognize that this is not a dispositive argument.

For another case, consider weight. People think a lot about weight, and relatively little about interplanetary travel. Thus, *almost* always, when people talk about weight (excluding scientists, astronauts, physics students, and the like), they are talking about weight on Earth. And, as mentioned, most people know that weight is a relativized notion. However, since the average person's thoughts rarely pertain to his or her weight on another planet, is it really plausible to

assume that *in every case* where people think about weight, they token a concept of the body or planet exerting the gravitational force? Again, *prima facie*, the answer is no.

I claim, then, that for many cases of unarticulated constituents / relativized concepts, although people know that the concept is relativized, if the fact that it is relativized is hardly ever cognitively or conversationally relevant, it is simply ignored, and subjects entertain thoughts as though the concepts were not relativized. In such cases, we would have TUCs.

iv. An Argument from Cognitive Efficiency

We can try to strengthen the force of these *prima facie* cases of TUCs with an argument regarding cognitive efficiency:

- (i) There is some non-trivial cost to tokening a concept.
- (ii) *Ceteris paribus*, the brain will reduce its tendency to perform unnecessary tasks.
- (iii) Therefore, *ceteris paribus*, the brain will reduce its tendency to token concepts unnecessarily.
- (iv) By “tacitly ignoring” the fact that a concept is relativized, a subject can gain all the usual benefits of tokening the concept without the cost of also tokening a concept of its relatum.
- (v) Therefore, *ceteris paribus*, subjects will treat relativized concepts as having TUCs and “tacitly ignore” the fact that they are relativized, when doing so is cognitively beneficial.

First, we must unpack “tacitly ignoring.” We can do so in an architectural way. We shall say that a subject S tacitly ignores the fact that a concept C is relativized to some relatum R just in case S knows that C is relativized to R, but S temporarily treats the concept as *not* being relativized, namely by granting thoughts about C the inferential role that she would grant thoughts about C *relative to R*. In other words, *X is C* is given the inferential status of *X is C relative to R*.

This assumes that such an operation is cognitively realistic, i.e., the mind can temporarily *alias* one thought for another. Although this remains speculation, it is not implausible.

The argument also assumes that the mind/brain tends to optimize its own activity (see premise (ii)). This is also partly speculative, and partly vague, but again not implausible. It also follows from the Cognitive Principle of Relevance, which states that “human cognition is geared to the maximisation of relevance.”³¹ Thus, this argument should be of interest to Relevance Theorists.

The topic of TUCs is not frequently associated with Relevance Theory. Relevance Theorists are more often grouped with the contextualists as holding the view that thoughts are generally articulated at the level of thought, although semantically and syntactically underspecified in communication. However, if one takes seriously the Cognitive Principle of Relevance, then the most natural position would be one in which unarticulated constituents remain unarticulated in thought (totally unarticulated) whenever doing so is cognitively “safe.” Since Relevance is *decreased* in relation to processing cost, a syntactically cheaper representation (one with TUCs) would often be more Relevant than a fully articulated one. Therefore, Relevance Theorists ought to expect that whenever doing so is “safe,” subjects will treat unarticulated constituents as totally unarticulated.

³¹ (Sperber & Wilson 2002, 14).

This, of course, raises the question of when such “tacit ignorance” is “cognitively safe.” Intuitively, the idea is that it is safe to implicitly ignore the time zone when you are in your house, talking about the time of day, and what you plan to do that day, and you don’t have any relatives in Japan, and you don’t plan to travel soon, etc. On the other hand, it is not safe to do so when just disembarking from an international flight, and you are trying to think about whether your child is still awake back home. In the latter situation, the time zone is a contextually relevant relatum, and tacitly ignoring it could lead to errors and bad inferences.

Thus, the argument assumes that the brain is somehow able to, on average, “know” or predict when tacitly ignoring a relatum is cognitively safe, and it is geared to maximize efficiency in those situations by tacitly ignoring the relatum.

Whether all of this is true cannot be convincingly argued here. However, I consider the argument sketched here a strong *prima facie* argument for the prevalence of TUCs. In conjunction with the *prima facie* plausible cases from the previous section, that is a double dose of *prima facie* plausibility for the prevalence of TUCs in natural language.

v. Conclusions

I have tried, through examples, to refute Stanley’s syntactic hypothesis. If even one TUC can be found in natural language, then the view must be rejected. First, I argued that animal signals universally exhibit TUCs, since they concern things like the present time and location, even though they generally have no meaningful parts which refer to those things. Second, I claimed that there are

undeniable historical precedents for TUCs, where we ourselves were in a similar situation with respect to the relativized nature of weight as Z-landers are regarding the nature of rain. Finally, I tried to argue that TUCs may in fact be pervasive. First, I appealed to some examples where it seems intuitively doubtful that the unarticulated relatum is articulated at the level of thought. Second, I proposed an argument from cognitive efficiency (or, the Cognitive Principle of Relevance), according to which there is a cognitive benefit in using TUCs (rather than their fully articulated counterparts), and so subjects will often implicitly ignore the fact that a certain concept is relativized and treat it as absolute. In such cases, the relatum is a TUC of the thought. The argument makes a number of strong assumptions about the nature of the mind/brain, but it too has a certain plausibility.

What I hope to have established is that: (a) it is undeniable that animal signals exhibit TUCs; (b) It is undeniable that there are historical cases of TUCs in natural language (*weight*, *simultaneity*, etc.); (c) it is *prima facie* plausible that TUCs are in fact quite prevalent in natural language. Furthermore, I take (b) to be a decisive refutation of Stanley's admittedly ambitious hypothesis that all context-sensitivity can be traced to logical form. Pre-Newtonian *weight* thoughts and statements were context-sensitive (they were relativized to Earth, even though this fact was not understood), but this context-sensitivity cannot be traced to logical form unless one assumes that logical form is omniscient, in the sense that it prefiguratively contains covert variables for all the forms of relativized concepts that humanity will discover in the future.

What about the pragmatists and contextualists, who claim that unarticulation is resolved at the level of thought? It is reasonable to think that this phenomenon does frequently occur as well. I hope to have established the reality, and perhaps prevalence, of TUCs, but I do not rule out the contextual

story as applicable in many contexts as well.

3. Is a Uniform Account Possible?

I. Introduction

Stanley (2007) defends the view that “all truth-conditional effects of extra-linguistic context can be traced to logical form.” I have argued that the existence of Totally Unarticulated Constituents (TUCs) refutes this position, and I claim that clear examples of TUCs have been attested. However, Stanley’s view has the theoretical virtue of providing a uniform account of context-sensitivity. Is it possible to retain a uniform position if we abandon Stanley’s strong syntactic hypothesis?

On a contextualist position, one would hold that the unarticulation is resolved at the level of what is meant – that is, although there is no linguistic part of my utterance of ‘It’s raining’ that encodes reference to New York City, there is a part of my *mental* representation that refers to a (vague) place. A TUC, on the other hand, is an entity which a thought concerns, but which is not represented by any part of the thought.

If my arguments in the previous chapter were sound, then there are at least some cases of TUCs. Therefore, the only way to maintain a uniform account would be to defend the claim that *all truth-conditional effects of extra-linguistic context can be traced to TUCs*. In this chapter, I will explore whether this thesis is tenable, and if not, how we can identify when we are confronted with a TUC. First, however, I must address in more detail Stanley’s primary argument in defense of the syntactic thesis.

II. The Argument From Binding

As discussed in the previous chapter, an argument made repeatedly by Stanley in defense of the existence of covert variables comes from the fact that unarticulated constituents may acquire a bound variable reading, as in the following type of case:

(1) Every time I light a cigarette, it rains.

What is significant about a sentence like (1) is that the location of the rain seems to co-vary with the location of the lighting of the cigarette, as in:

(1*) Every time t such that I light a cigarette at t at location l , it rains at t at l .

How can an unarticulated constituent acquire such a reading? According to Stanley, the best explanation is that the bound reading exists due to the existence of syntactic binding of a covert variable, and there are no acceptable alternative explanations. “On the third (and weakest) interpretation of the binding argument, it is an inference to the best explanation. By postulating a covert variable, one can account for the bound reading, and there is no other satisfactory way to account for it.”³² This quotation suggests that Stanley considers the covert variable explanation the “only game in town,” but at least it is clear that Stanley favors reading the binding argument as an inference to the best (or only) explanation.

³² (Stanley 2007, 214).

Stanley rejects the idea that a bound variable could be supplied via free enrichment. The main problem is that if the processes described by the defender of free enrichment existed, we would expect to find a number of readings that don't in fact exist. Therefore, since no one has ever offered explanatory constraints on free enrichment, appealing to it is unexplanatory.

However, we can demonstrate by example that the syntactic explanation cannot possibly be the only way that a sentence like (1) can acquire a bound variable reading. Let's consider a variation on the Z-land case.³³ Suppose we have a nomadic tribe that has never had any contact with other groups. They move from one location to another, and sometimes it rains where they are. When it does, they say things like 'It's raining.' It does not occur to them that it might be raining somewhere else at a given time, but not raining at their location. So, when they say things like 'It's raining,' their utterances concern the location *that they are in at the time*.

In other words, it seems that the unarticulated constituent in this case does not refer to a specific, constant location, but rather concerns the location of utterance – it is an indexical unarticulated constituent. However, we cannot assume that it is a pure indexical, such as 'I.' For if one of them utters 'It rained last week,' we can imagine that this utterance is judged true just in case it rained last week *at some place where they were at the time* (not necessarily the current location of utterance). Again, for the same reasons that motivate one to hold that the Z-land case is an instance of a TUC, this case seems to involve a TUC as well – since (by hypothesis) it never occurs to them that it might be raining at one location but not another, the idea that 'rains' would be relativized to a location would seem bizarre.

³³ Cf. (Perry & Blackburn 1986).

Extending the example, let's suppose that this tribe enjoys smoking tobacco, although they only partake once every week or two. As it just so happens, for the past two months, every time they have gathered to smoke tobacco together, it started to rain at their location. Thus, we can imagine such a speaker uttering:

(2) Every time we smoke tobacco, it rains.

What is the natural reading of this utterance? Precisely the bound variable reading:

(2*) Every time t such that we smoke tobacco at t at location l , it rains at t at l .

But if this is really a TUC, then there can't be any location variable either in the logical form of the sentence, or even of the corresponding thought. Thus, either the scenario sketched above does not constitute a possible TUC, or it is possible for TUCs to acquire a bound variable reading, despite the absence of syntactic binding.

I believe the case for (2) being an example of a TUC is just as strong as with Perry's original example. However, one might argue that this is not really a case of a bound variable reading. Since the imagined tribe just doesn't consider whether rain occurs at one location but not another, perhaps the correct reading is that 'rains' is for them something that happens *everywhere*. So, (2) should really be analyzed as meaning something like:

(2**) Every time t such that we smoke tobacco at t , it rains (*everywhere*) at t .

In this case, the location of the rain is not dependent on the location of the smoking. However, this reading does not accurately capture the function of ‘rains’-utterances for this tribe. For if we consider bare utterances like ‘It’s raining,’ then this utterance should be judged true or false simply on the basis of whether it’s raining at the general location of utterance – i.e., there is no reason to claim that the speakers are mistakenly asserting that it is raining everywhere whenever they utter ‘It’s raining.’ This would result in a massive error theory for utterances that are unproblematically accepted as true or false by the community based on the immediate meteorological circumstances. (Similar problems would arise for any attempt to read ‘rains’ in (2) as ‘rains *somewhere*.’)

Unfortunately, the historical examples of TUCs (such as weight and simultaneity) do not lend themselves easily to a similar thought experiment. However, a scenario involving weight can perhaps be contrived. Suppose that some pre-Newtonian civilization somehow acquires a vehicle for space travel. They start flying around from planet to planet, and one of the rituals they do at each planet is to step on a scale and weigh themselves. As it turns out, they happen to travel to increasingly dense, large bodies of mass – thus, their weight on each planet that they visit is greater than their weight on the previous planet. One of them notices this and exclaims:

(3) Every time I weigh myself, my weight is higher.

Here, the natural reading is:

(3*) Every time t such that I weigh myself at t on planet x , my weight at t on x is higher (than my weight at t' on y).

By hypothesis, the pre-Newtonian does not consider that weight is a relativized property (though they might soon figure this out if they carry on). Thus, they are puzzled as to why their weight appears to be increasing. However, there is at least one way of construing this scenario such that the utterance of (3) is true and would be accepted as true by the speech community. If this scenario is plausible, then we have another case of a bound variable reading of a TUC.

Thus, if either of the thought experiments sketched in this section are plausible, that would demonstrate that TUCs can acquire bound variable readings. Since there can't be any syntactic binding taking place with a TUC, this would demonstrate that syntactic binding is not the only possible source of a bound variable reading of an unarticulated constituent.

Stanley defends the uniform syntactic hypothesis by claiming, essentially, that it is the only game in town – there is no other “satisfactory” way to explain bound variable readings than by appealing to syntactic binding. But if bound variable readings can exist with TUCs, then this claim is refuted.

III. Uniform Unarticulation Thesis

I have argued that the argument from binding is refuted by counterexample. So, if we wish to maintain a uniform account of unarticulated constituents, we must consider the thesis that all cases of unarticulated constituents involve TUCs. Let us call this the “Uniform Unarticulation Thesis” (UUT).

We can find some support for UUT in the fact that TUCs appear to be a primitive feature of all, or almost all, representational systems. As previously argued, animal signals almost invariably pertain to occurrent facts relatively local to the agent's environment. Thus, as a general tendency, they concern the present time and place (e.g., *Aerial predator is present here now!*), but insofar as they lack meaningful structure, they do not represent or articulate the time and place. Thus, TUCs are a primitive feature of communication systems. Similarly, I have argued that they are prior with respect to development as well. Often, a child will become competent with a relativized word/concept (such as the time of day, weight, etc.) before they acquire the awareness that it is relativized. Thus, since TUCs are basic with respect to individual human development, and with respect to communication systems more generally, this lends some *prima facie* support for the view that they are pervasive in adult human speech as well.

Further motivation for UUT derives from the hypothesis that complete articulation is never, in fact, possible. This can be argued by invoking Searle's (1978) notion of "the background." This pertains to the fact that when I have a thought, such as that expressed by "The cat is on the mat," there are an indefinite number of background conditions that affect the literal meaning of that sentence, or which have to obtain in order for the sentence to have a literal meaning at all.³⁴ For instance, a somewhat normal gravitational situation has to be assumed – if the cat is floating in space and is touching the mat, then it doesn't really count as being *on* the mat, etc. But of course, I normally would never consider the significance of a gravitational field when I entertain such a thought. And there are an indefinite number of further conditions that bear on the meaning/meaningfulness of that thought but which are totally remote from my awareness – that is, they are in "the background."

³⁴ I find it somewhat difficult to parse Searle's overarching point in the essay 'Literal Meaning' (1978), although it is a valuable resource.

As Searle writes: “The truth conditions of the sentence will vary with variations in these background assumptions; and given the absence or presence of some background assumptions the sentence does not have determinate truth conditions.”³⁵

Searle presents “the background” as a unique form of indeterminacy (e.g., distinct from vagueness) that is irresolvable. This is due to the fact that (i) the background conditions for a sentence in context are of an indefinite number; and, (ii) every attempt to articulate a background condition will invoke notions that themselves have background conditions for their application, which leads to an infinite (or otherwise vicious) regress:

The general point is that representation, whether linguistic or otherwise in general goes on against a background of assumptions which are not and in most cases could not also be completely represented as part of or as presuppositions of the representation, for the two reasons we have already stated: the assumptions are indefinite in number and any attempt to represent them will tend to bring in other assumptions.³⁶

Although Searle’s framing of the issue pertains to the “literal meaning” of a sentence (a notion which I do not find particularly helpful), and he takes as “received opinion” some hypotheses that are rejected by contextualist theorists, there is a lesson to be gleaned from “the background.” It is a sort of Wittgensteinian lesson, which reminds one that precision and articulation are not needed for

³⁵ (Searle 1978, 214).

³⁶ (Searle 1978, 219).

language to fulfill its function, and that vagueness, imprecision, underspecification, and the like, are not imperfections, or deviations from an ideal language. As Wittgenstein writes in the *Investigations*:

If I tell someone “Stay roughly here” – may this explanation not work perfectly? And may not any other one fail too?

“But isn't it an inexact explanation?” - Yes; why shouldn't we call it “inexact”? Only let's understand what “inexact” means! For it does not mean “unusable”.³⁷

The entities and conditions that comprise “the background” of a thought/utterance are distinct from the notion of unarticulated constituents. However, they both provide contrast with a picture of representation according to which an ideal representation is one that is fully precise and articulated. Rather, language is intrinsically situational, vague, and underarticulated. Articulation is a cognitive achievement that is a result of reflection and scientific investigation, but there is nothing defective about underarticulation – it is inevitable. This picture resonates deeply with the motivation for situation semantics, as explained by Jon Barwise & John Perry (1981):

Situations are basic and ubiquitous. We are always in some situation or other. Human cognitive activity categorizes these situations in terms of objects having attributes and standing in relations to one another at locations – connected regions of space-time. Human languages reflect (and enhance) this cognitive activity by giving us a

³⁷ (Wittgenstein 1953, §88).

way of communicating information about situations, both those we find ourselves in and those removed from us in space and time.³⁸

So far, I have been attempting to marshal some *prima facie* considerations that support a picture according to which TUCs are pervasive. However, UUT faces some serious difficulties.

First of all, for all the examples of TUCs involving Z-land, weight, animal signals, etc., the unarticulated constituent is typically a local, static feature of the agent's environment, such as the planet that they are on. But the phenomenon of unarticulated constituents in natural language is not so constrained. An utterance of 'It's raining' could be used by a speaker in New York City to talk about the weather in New York, Los Angeles, London, Mars, or anywhere else. There is no systematic connection between the immediate context of utterance and the location that is unarticulated. So, does the natural language phenomenon really resemble the primitive cases of TUCs that I have argued for?

With primitive cases of TUCs, such as Z-land weather reports, or pre-Newtonian weight reports, the unarticulated constituent factors into the proposition expressed in virtue of how the thought/utterance is *used*. Pre-Newtonian weight reports concern Earth because they were used to orient behavior in a way that was appropriate for weight *on Earth*. And we can infer that Earth was unarticulated because they had no notion that weight was a relativized property. However, most speakers *are* aware that things like weight, time of day, temperature, etc., are all relativized in various ways. And the correct interpretation of a given application of such notions typically depends on what the speaker *has in mind*. How can TUCs be sensitive to the speaker's intentions in this manner?

³⁸ (Barwise & Perry 1981, 668).

One possibility is that whereas the Z-landers are *ignorant* of the fact that weather-reports are location-relative, a normal speaker can simply *ignore* that fact in a given context. For instance, my weight-reports are almost always about weight on Earth; so, why can I not “pretend,” or assume for the purposes of the conversation, that *weight* always pertains to *weight on Earth*? That is, I know that weight is a relativized property, but because the relativization is so infrequently relevant, I can somehow ignore that information, and it will be understood in the conversation that when I talk about “weight” I am talking about “weight on Earth.”

This presupposes that the common ground, or the situation that provides the background for the discourse, is subject to the speaker’s intentional control. In some cases, this may appear unproblematic. If we are in London discussing whether to go for a walk, then surely there is no harm in my pretending that *weather* always pertains to *weather-in-London*. But this is only a prototypical case. Speakers can intentionally shift the discourse situation (or so it would appear, on this view) and just expect the hearer to follow along. For instance, suppose that while discussing whether to take a walk in London, I am also reading the news, and I say,

(4) Oh dear, there’s going to be a massive hurricane,

and mean thereby that there’s going to be a massive hurricane in the Caribbean, then the attentive listener will be able to work out that I’ve tacitly shifted the discourse situation to pertain to hurricanes in the Caribbean (or, rather, elsewhere than London, where a hurricane would be highly unlikely), despite the overall discussion pertaining to the climate in London. Perhaps the changing

discourse situation can be explained through a process akin to what David Lewis calls “accommodation.”³⁹

As I sketched in the previous chapter, it is arguable that this kind of strategy, when available, could increase the Relevance of my utterances. First, Relevance is decreased in proportion to processing costs. Then, if we assume that, *ceteris paribus*, a representation with fewer constituent parts is easier to process, then it would follow that invoking a TUC, in the manner just described, could decrease the cost of processing the utterance, and hence increase its Relevance. To put it in other terms, it would certainly be a colossal waste of breath to constantly have to insert the qualifier ‘on Earth’ into sentences like ‘That book weighs five pounds,’ so maybe it is a waste of cognitive resources to even token the *concept* of Earth in entertaining that thought (under normal circumstances).

On the other hand, there is a concurrent risk involved in invoking TUCs. By leaving some entity unarticulated, you are assuming that the hearer will automatically recognize the correct interpretation. However, this cannot always be taken for granted. As a member of one of the few populations that still employs the imperial system of measurement, I frequently find myself getting into confusion when discussing the weather with my Canadian or British friends. For instance, an utterance like,

(5) It’s going to be 20 degrees tomorrow! We should go to the beach,

³⁹ Cf. (Lewis 1979). In a discussion of language as a matter of “scorekeeping,” Lewis (1979, 347) writes: “As I hope my examples will show, conversational score does tend to evolve in such a way as is required in order make whatever occurs count as correct play.” For instance, if a speaker presupposes a proposition that was not already in the common ground, it will “straightway come into existence.”

will likely strike me as ironic (or simply confusing) before I remember the existence of the metric system.

The possibility of confusion (reduction of clarity) arguably decreases the Relevance of a statement involving TUCs, since a fully attentive hearer might have to do some extra interpretive work in order to properly track the discourse situation that is being assumed. This would offset any potential gains stemming from processing a syntactically simpler representation.

It is important, at this point, to clearly separate two distinct issues. One has to do with (i) the circumstances in which a person will token a thought involving a TUC. The other has to do with (ii) the correct analysis of *utterances* involving unarticulated constituents.

Of course, these issues are not unrelated. In both instances, it seems like what we want is a system which allows for TUCs, but somehow manages to confine their use to cognitive/discourse contexts in which the unarticulation has no potentially harmful consequences (such as confusion or miscommunication). Perhaps, in keeping with the Cognitive Principle of Relevance,⁴⁰ we can speculate that the mind/brain has evolved in such a way as to *minimize articulation to the degree that is "safe,"* perhaps relying on heuristic mechanisms to ensure (relative) safety.

These are difficult questions for cognitive science. However, even if we were able to resolve the issue of under what circumstances subjects token thoughts involving TUCs, the question of the content of *utterances* would still be unresolved.

⁴⁰ (Sperber & Wilson 1995).

On the current hypothesis (UUT), unarticulated constituents remain unarticulated at the level of thought, and speakers manage to communicate thoughts involving unarticulated constituents by intentionally (or automatically) assuming a particular situation as the common ground of the discourse. On the competing free enrichment account, unarticulated constituents (in speech) *are* articulated in the thoughts that the speaker expresses, and the speaker intends the hearer to recover the articulated thought via pragmatic processing. So, the question of which account is correct depends on whether the speaker intends the hearer to entertain an *articulated thought*, or an *underarticulated thought*.

This seems like the correct way to frame the disagreement between the “implicit” and “explicit” approaches to underspecification. The difficulty, however, is that it is doubtful that speakers typically *have* intentions that would support either side of the debate. If I utter, ‘It’s raining,’ then for one account to be correct, I ought to have a distinct intention that the hearer entertain an (un)articulated thought. But why should speakers ever have such an intention? Why should one ever *care* whether the hearer’s thought be articulated or unarticulated? As long as the propositional content of the meaning recovered by the hearer corresponds to the speaker’s intentions, then communication has succeeded – the precise form of the hearer’s interpretation of the speaker’s meaning is irrelevant. So, since speakers lack the kinds of individuating intentions that would determine whether the content of the *utterance* should be analyzed as involving a TUC, we appear to be facing a deadlock. Not only do we lack clear support for one hypothesis over the other, we have found reason to doubt that the kinds of intentions that would be required to support either hypothesis generally exist.

IV. Compositionality

Jerry Fodor (2001) presents an argument which raises problems for the very notion of a TUC. Fodor is discussing the priority of language vs. thought with respect to having intentional content. He argues that thought must be prior to language because linguistic utterances are often underspecified with respect to the thoughts that they express. But if they are underspecified, then they cannot be compositional. And if only one of the two is compositional, then that one must be prior. Therefore, Fodor infers from the existence of unarticulated constituents in language that language is not compositional, and therefore cannot be prior, and so therefore thought must be prior to language, and thought had better be compositional:

The moral is that if language is compositional... then how a sentence is put together must be very explicit about how the corresponding thought is put together. But... sentences are remarkably *inexplicit* with respect to how the thoughts they express are put together. So either the content of the thought is different from the content of the sentence that expresses it, or the sentence isn't compositional. I take it that the first disjunct is preposterous; so I take it that the second disjunct must be true.⁴¹

Fodor is here referring to the fact that Gricean maxims compel speakers to leave information out of the explicit utterance when it is obvious, such as 'here' and 'now' in a typical utterance of 'It's three o'clock (here and now).' However, he holds that this underspecification must be resolved at the level of thought. He even writes that 'It's three o'clock' "is *syntactically* inexplicit; presumably what I've

⁴¹ (Fodor 2001, 12).

uttered is an abbreviated form of something like ‘it’s three o’clock here and now.’”⁴² Regarding the possibility of unarticulated constituents or underspecification in the thought itself, Fodor writes:

No such objections as I’ve been arguing against the compositionality of language can hold against the compositionality of thought. For, whereas the content of a sentence may be inexplicit with respect to the content of the thought it expresses, a thought can’t be inexplicit with respect to its own content; there can’t be more – or less – to a thought than there is to its content because a thought just *is* its content.⁴³

Unfortunately, Fodor does not elucidate the assumption that a compositional representational system cannot exhibit underspecification (perhaps because he sees it as self-evident). Indeed, if we classify thoughts according to their *content*, then it would appear to be impossible to explain how two thoughts could have the same content but different form. But this is precisely what is required by the defender of TUCs – that the thought *It’s raining* can have the same propositional content as *It’s raining here and now*.

Since it isn’t clear to me exactly what assumptions Fodor is relying on for this claim, it is difficult to engage with it. However, one apparent consequence of Fodor’s view is that there cannot be atomic representational contents, like animal signals. For an atomic representation has no parts, and therefore whatever content it has must be “inexplicit.” But a thought, or indeed any representation that has “original intentionality,” “can’t be inexplicit with respect to its own content.” So, this would

⁴² (Fodor 2001, 12).

⁴³ (Fodor 2001, 14).

entail that there cannot be atomic representations. I find this consequence very troubling. At least, it is firmly out of keeping with the overall aim of the current project.

Nevertheless, we might ask whether there is any reason to think that underspecification conflicts with compositionality. I see no reason to hold that it does. It depends, of course, on what one means by ‘compositionality.’ In the case of the nomadic Z-landers, it’s true that ‘It’s raining’ does not contribute the same content to the propositions it expresses relative to different contexts of utterance. However, its behavior is not unlike that of an indexical. And I see no reason to think that indexicals must be excluded from the language of thought, or from any compositional system. We can analyze Z-landers’ *rains*-thoughts as having the same content as *rains-here-now*-thoughts. (Of course, perspicuity/degree-of-articulation is not something that is preserved in analysis.) So, if we allow that two-dimensional semantic systems can be compositional, and we think of TUCs as a kind of two-dimensionality, then I see no reason to think that the existence of TUCs is incompatible with compositionality, *pace* Fodor.

V. Conclusion

This, then, is how I see the current state of play:

(i) Stanley’s Syntactic Thesis: Stanley’s main claims, including the argument from binding, are refuted by the actual and hypothetical examples of TUCs cited earlier in this work. Stanley offers his view, admittedly somewhat radical, as the only game in town, but the existence of TUCs demonstrates that there must be some way in which context can affect truth-conditional content besides fixing the

values of covert variables. Given this fact, it seems more promising to explore the alternatives to the syntactic approach than to make strong assumptions about logical form.

(ii) Free Enrichment vs. TUC: On the free enrichment account, the speaker intends the hearer to recover a fully articulated thought – unarticulation is resolved at the level of the thought. On the TUC account, the speaker intends the hearer to recover a partially unarticulated, or underspecified thought. The primary difficulty in adjudicating this issue is that speakers do not typically have either type of intention.

Unfortunately, I do not see any clear way out of this difficulty. The existence of historical cases of TUCs (weight, simultaneity, etc.) demonstrates that the TUC account is necessary *at least* in those cases (and arguably for many others as well). Meanwhile, Neale (2007) has presented some examples that he cites as highly problematic for a situation-based, or implicit, approach to underspecification:

(6) The Russian voted for the Russian.

Neale imagines a context in which (6) is used “to explain how one of the boxing judges voted in a boxing match between a Swede and a Russian.” The problem with this sort of example is that if situations are supposed to tacitly restrict the domain of quantification, so that apparently incomplete descriptions can determine a unique referent, then a sentence like (6) will require at least *two* situations, since the two occurrences of ‘the Russian’ are non-coreferential. Thus, Neale argues that a situation-based approach would have to endorse the view that *each nominal* can be associated with a different situation, and this can shift in the course of an utterance. This suggests that a free enrichment account may be preferable at least for this kind of case.

However, for the vast majority of speech situations, there seems to be no reason to prefer one account over the other. That is, we do not simply lack the necessary kind of evidence (e.g., about the precise structure of a subject's thought), we have reason to doubt that the requisite conditions necessary for either account to be correct obtain.

This may strike one as a disappointing conclusion, but if the foregoing arguments are correct, then it is indeterminate, at least for a vast majority of cases, whether the free enrichment account or the TUC account is correct. This is why it is important to separate the cognitive and linguistic issues. Even if we assume that there *is* a fact of the matter about whether any individual thought involves an unarticulated constituent (i.e., that thoughts have fully determinate structure), that does *not* appear to shed much light on the linguistic question – that is, of which analysis offers the correct account of speaker's meaning.

4. Is There a Meaning-Intention Problem?

I. What Is the Meaning-Intention Problem?

A. The Meaning-Intention Problem: Background

When confronted with some recalcitrant data, it is common practice in linguistic theorizing to invoke hidden representational structure. Hidden structures abound in syntactic theory, but they are common in semantics as well. If one is unable to generate the intuitively correct truth conditions for a sentence simply by assigning semantic values to the overt parts, one can often resolve this problem by appealing to covert variables or modifiers.

For example, it is notoriously hard to account for truth value judgments regarding belief attributions without invoking such hidden structure:

- (1) Alice believes that George Eliot is a man.
- (2) Alice believes that Mary Ann Evans is a man.

Let us suppose that Alice has heard the names ‘George Eliot’ and ‘Mary Ann Evans’ before, and she thinks they denote two distinct English authors. Then it is plausible to suppose that (1) and (2) might receive different truth values. But a naïve, extensional analysis of proper names would make it difficult to account for this difference. This is in part the motivation for Fregean theories of proper names.

On a Fregean analysis, expressions appearing in intensional contexts do not denote their reference, but instead denote their (customary) *sense*. Since ‘George Eliot’ and ‘Mary Ann Evans’ differ in sense, their occurrences in (1) and (2) differ in reference, and hence (1) and (2) may differ in truth value. Thus, although a classical Fregean analysis does not require positing any hidden *structure* in a belief report, it does imply that the proposition meant contains *senses*, entities of which speakers may lack a concept.

A competing account can be found in the writings of neo-Russellian theorists such as Nathan Salmon (1986). These theorists differ from the Fregeans in holding that the semantic value of a proper name is exhausted by its referent, however they explain the apparent failure of substitutivity by positing substantial hidden complexity in belief attributions. In particular, whereas a naïve view might hold that *believes* expresses a two-place relation between a believer and a proposition believed, Salmon proposes that belief attributions in fact express a three-place relation between a believer, a proposition believed, and a propositional *guise* under which the believer accepts the proposition:

- (3) *For some guise* x , Alice grasps that George Eliot is a man by means of x and BEL(Alice, that George Eliot is a man, x).

According to this view, the semantic analysis of both (1) and (2) is (3). Thus, both sentences are literally true. However, (2) is highly misleading, and pragmatically conveys something false (e.g., that Alice would assent to ‘Mary Ann Evans is a man’). Thus, we are simply judging the sentence according to what is pragmatically conveyed, rather than its semantic content. This view also allows us to account for sentences like,

- (4) Alice believes that Mary Ann Evans is not a man,

as follows:

- (5) *For some x , Alice grasps that George Eliot is not a man by means of x and BEL(Alice, that George Eliot is not a man, x).*

Alice may accept (BEL) that George Eliot/Mary Ann Evans is a man under some guises, and reject it (or accept its negation) under other guises. This elegant analysis allows one to preserve a simplified, Millian view of proper names, while accommodating the cases of substitution failure in a fairly systematic way.

However, the neo-Russellian analysis of belief attribution also requires the theorist to posit entities – guises – of which the speaker is ignorant, in the semantic analysis of ordinary belief reports. In addition, it implies that the logical structure of a belief attribution is more complex than an ordinary speaker might recognize. These facts are the basis of a persuasive argument against such theories, first developed by Stephen Schiffer (1992). This argument, known as the *meaning-intention problem*, claims that such analyses cannot be correct, because speakers do not, and cannot, mean what such analyses require them to mean.

Schiffer develops his argument in response to the *hidden-indexical theory* of belief attributions. This view, similar to Salmon's, analyzes belief reports as follows:

- (6) Ralph believes that Fido is a dog.
- (7) (For some m) (Π^*m & B(Ralph, <Fido, doghood>, m))

On this view, (6) is analyzed as the statement that Ralph believes that Fido is a dog under some mode of presentation (MOP) m , where m satisfies some contextually specified constraint on MOPs Π^* .

“The meaning-intention problem” in fact covers a number of related concerns raised by Schiffer for this type of analysis:

- (a) The Awareness Problem – If the hidden-indexical theory of belief reports is true, then speakers lack full, conscious awareness of what they mean, for most speakers have no idea that they are referring to a MOP-property: “Thus, if the hidden-indexical theory is correct, then [the speaker] has no conscious awareness of what she means, or of what she is saying... and this is a *prima facie* reason to deny that she means what the theory is committed to saying she means.”⁴⁴
- (b) The Cognitive-Resources Problem – If the hidden-indexical theory is true, then what a speaker means may involve entities about which the speaker is totally ignorant (such as MOP-properties).
- (c) The Specificity Problem – If the hidden indexical theory is true, then there must be some *particular* property Π^* of MOPs that the speaker is referring to. But it’s doubtful that the speaker’s intentions serve to pick out any such particular property (the speaker lacks “specifying intentions,” we may say).

⁴⁴ (Schiffer 1992, 514).

On the assumption that speakers *do*, generally, have full, conscious awareness of their own speech and thought contents, and that speaker-meaning requires a speaker to have a conceptual grasp of the particular entities and properties that comprise the contents of their speech acts, the meaning-intention problem raises serious doubts about the hidden-indexical theory of belief reports. Similar doubts would arise for any theory that posits, in the semantic analysis of some sentence, entities, properties, or structure about which the speaker (a) lacks conscious awareness, (b) is totally ignorant, or (c) lacks specifying intentions.

Schiffer (1996) deploys this argument against another prominent contextualist theory that invokes hidden indexicals: contextualist theories of knowledge claims. Such theories are in part an attempt to respond to the fact that a sentence like ‘I know that I have hands’ might be judged true in an ordinary context, but false in the context of a discussion about Descartes’s evil demon.

Contextualists claim that knowledge claims are tacitly relativized to a *standard of knowledge*, which affects the kind of evidence needed to support a knowledge claim in a given context. Thus, ‘I know that I have hands *relative to a low standard of knowledge*’ might be true, while ‘I know that I have hands *relative to a high standard of knowledge*’ might be false.

Schiffer sees two problems with this approach. First, it again implies that speakers are not fully aware of what they mean, think, or say: “But no ordinary person who utters ‘I know that *p*’, however articulate, would dream of telling you that what he meant and was implicitly stating was that he knew that *p* relative to such-and-such standard.”⁴⁵ Second, it implies that speakers’ conception of their own speech contents may not only be *incomplete*, but may be seriously mistaken. For the contextualist

⁴⁵ (Schiffer 1996, 326-327).

hopes to explain why the following kind of argument *appears* to be sound, even though the conclusion seems incorrect:

- (A) I don't know that I'm not a brain-in-a-vat.
- (B) I know that I have hands only if I know that I'm not a brain in a vat.
- (C) Therefore, I don't know that I have hands.

The response is that (C) strikes us as false because we are prone to interpret it as:

(C*) Therefore, I don't know that I have hands *relative to a low standard of knowledge*,

even when the context dictates that we ought to interpret it as:

(C**) Therefore, I don't know that I have hands *relative to a high standard of knowledge*.

The problem, according to Schiffer, is that it's highly implausible to suppose that speakers could be so confused about the content of their own speech acts: "It's as though a fluent, sane, and alert speaker, who knows where she is, were actually to assert the proposition that it's raining in London, when she mistakenly thinks she's asserting the proposition that it's raining in Oxford."⁴⁶

B. Assumptions Behind the Meaning-Intention Problem

⁴⁶ (Schiffer 1996, 326).

As we have seen, the meaning-intention problem attempts to cast doubt on contextualist semantic theories by arguing that we can't take speakers to mean what such theories require them to mean. However, the argument is deployed with different emphasis in different places. Thus, it is worthwhile to clearly articulate the assumptions upon which such an argument rests:

(Awareness) If a speaker, S , performs an utterance, and means thereby that p , and x is a constituent of the proposition p , then S must be consciously aware that x is part of what she means.

(Cognitive Resources) If x is a constituent of the proposition meant by an utterance of sentence S , then any competent speaker who utters S must have a concept of x , whether consciously accessible or not.

(Specificity) If x is a constituent of the proposition meant by an utterance of sentence S , then any competent speaker who utters S must have referential intentions that uniquely determine reference to x , rather than any other nearby candidates that might appear to serve just as well in the context.

Here, I adopt a structured-proposition view of propositional content for ease of exposition. Saying that x is a constituent of the proposition p is simply an attempt to capture the intuitive idea that p is, in some sense, "about" x . If a speaker utters 'It's raining' and means thereby that it's raining in London, then she has said something about London. The truth conditions of her utterance are sensitive to how things are in London.

Although these principles are distinct, they all attempt to use the fact that meaning is an intentional act performed by speakers as a way to constrain what counts as an acceptable semantic analysis. In

Section II, I argue that these principles are not sound, and that the meaning-intention problem, as stated, is far too powerful.

C. The Aphonic-Intention Problem

Stephen Neale (2016) defends Schiffer's meaning-intention problem, and claims there is an additional problem related to the use of aphonic referring expressions. The problem is based on the Gricean analysis of what is it to refer *with* an expression:

(RW) In uttering x , S referred to o *with* (or *using*) e , relative to its i -th occurrence in x , iff for some audience A and relation R , S intended A to recognize that $R(e, x, i, o)$ and, at least partly on the basis of this, that S referred to o in uttering x .

R , in this definition, is, effectively, the *inference-base* feature of the expression e , that is, the property that S believes e has (relative to its position i in the sentence x), such that the audience will recognize that S is referring to o partly on the basis of recognizing this feature. Often, this feature is simply the fact that e is conventionally used to refer to o .

This definition of *referring-with* invokes the more basic notion of *speaker-referring*, which Neale follows Schiffer in defining as:

(SR) In f-ing, S referred to o iff what S meant by f-ing is an o -dependent proposition (a singular proposition that has o as a constituent).

This definition, in turn, analyzes speaker-referring in terms of the more basic notion of *speaker's meaning*, and hence, together these definitions allow us to explain what it is for a speaker to refer *with* an expression in basic Gricean terms.

What's important, for the present discussion, about this Gricean analysis of referring-with is that it entails that if a speaker uses *e* to refer to *o*, then the speaker must have an intention that *has the expression e as part of its content*. It's not simply that, in referring to *o*, a speaker must employ the expression *e* in actualizing her intention, the way she must perform an alveolar stop in pronouncing the word *to*, something which most English speakers are capable of, whether or not they have a concept of *alveolar stop*. For in general, it isn't the case that if actualizing some intention requires a subject to utilize some capacity, then she must be aware of possessing that capacity. However, due to the nature of communicative intentions, if a speaker uses some expression *e* to refer to *o*, then she must intend for the audience to recognize *that e has some inference-base feature*, and thus, she must be capable of forming intentions whose content involves the expression *e*. This implies that referring with some expression *e* requires the subject to have some level of conceptual grasp of the expression *e* itself.

The problem that Neale raises is that this consequence is extremely dubious for cases of aphonic referring expressions, since most speakers have no idea that such expressions exist. Thus, there is a meaning-intention problem not only with respect to the entities that a speaker refers *to*, but also with respect to the entities that speakers refer *with*: "So an implicit reference theory according to which

speakers refer aphonically to mode of presentation types faces a compound problem: the theory has ordinary speakers referring *to things* they don't know about *with things* they don't know about."⁴⁷

This variation of the meaning-intention problem appears to rely on the following assumption, which is a fairly straightforward consequence of (RW):

(Syntactic Knowledge) If a speaker S refers to some entity o with some expression e , then S must be able to form an intention that has e as part of its content, and hence must "know about" e , in some sense.

I will argue that this assumption is incompatible with the data.

II. Extrinsic Parameters

There are a number of examples that appear to contradict the proposed principles supporting the meaning-intention problem. They seem to show quite clearly that a speaker S can express a proposition p that is about, or concerns, some entity e , even though S either lacks a concept of e , *contra* (Cognitive Resources), or simply lacks conscious awareness of expressing a proposition about e , *contra* (Awareness).

(8) This book weighs five pounds.

(9) The flash and the bang happened at the same time.

⁴⁷ (Neale 2016, 154).

(10) It's summer.

Each of these sentences expresses a proposition whose truth value depends on the state of some entity that is not explicitly mentioned in the sentence itself:

(8*) This book weighs five pounds (*on Earth*).

(9*) The flash and the bang happened at the same time (*relative to Earth as a frame of reference*).

(10*) It's summer (*in the Northern hemisphere*).

Most educated speakers know that weight is something that can actually vary depending on what large body of mass you happen to be standing on. And many (though perhaps not most) speakers know that simultaneity is not an absolute relation, but rather that two events are simultaneous only with respect to a reference frame.

What's important for our discussion is that despite the fact that many educated speakers are aware of such relativization, *many competent speakers are not*. Indeed, for large periods of human history, *no competent speaker* grasped that weight is relativized or had any idea what a reference frame was. This did not in any way prevent them from communicating complete thoughts in uttering sentences like (8)-(10). Therefore, *contra* (Cognitive Resources), it is not the case that if a speaker expresses some proposition *p* that contains some entity *o*, then she must have some conceptual grasp of *o*. *A fortiori*, she need not have conscious awareness of such a concept either, *contra* (Awareness).

Considering sentence (10), even if we assume that all, or almost all, competent speakers know that what season it is depends on what hemisphere you are in, it is entirely plausible that a speaker might

utter (10) and not be thinking about hemispheres at all, i.e., not have any awareness of saying something that is about a hemisphere. This again contradicts (Awareness). Even for those cases in which speakers have the requisite concepts, they may not be aware of invoking those concepts in performing their utterance.

Neale (2016, 160ff.) raises similar examples, referring to such cases as involving “extrinsic parameters,” but fails to draw any substantial conclusions from them. Instead, he focuses on the disanalogies between these cases and MOPs, in order to show that whatever solace the hidden indexical theorist hopes to find in such cases does not help the hidden indexical theory of belief reports. Examples (8)-(10) relate to

“factors *external* to us... about which we may be ignorant but about which we may acquire knowledge and thereby easily refine our linguistic behaviour. Mode of presentation types are not like this at all. They are supposed to be things *under* which beliefs are *had*, and learning about their existence and a great deal of information about their roles in theories of language and mind doesn’t even put theorists in a position to articulate the truth conditions of the propositions they actually express on given occasions using belief sentences if the hidden-indexical theory of belief reports is true.”⁴⁸

Thus, unlike with MOPs, “once speakers *learn* about time-zones, hemispheres and rest-frames, and learn a few additional words, they can easily describe the parameters relevant to the truth or falsity of what they are saying...”

⁴⁸ (Neale 2016, 163).

First of all, this seems doubtful. For instance, I know that whether or not two events are simultaneous depends on a reference frame. However, my knowledge of reference frames ends there. It is fair to say that I know far more about MOPs than about reference frames.

The main issue, though, is not whether MOPs are easier to grasp than reference frames, or *vice versa*. The problem is rather that examples like (8)-(10) seriously undermine the general assumptions that support the meaning-intention problem. It isn't clear what comfort the *proponent* of the meaning-intention problem is supposed to find in the fact that even though most speakers lack a proper concept of a reference frame, and many would never "dream of telling you that what he meant and was implicitly stating" was something to do with reference frames, there are *others* who do grasp the concept, and perhaps with sufficient training the rest of the population could do so as well.

The fact is that relatively few competent speakers grasp the concept of a reference frame. Some might be able to recall that simultaneity is a relative notion, if pressed, but even this is reserved for an educated segment of the population. Nevertheless, ordinary speakers are perfectly competent with phrases like "at the same time." This, again, shows that the principles adduced to support the meaning-intention problem cannot be sustained. There may be important differences between MOPs and reference frames, but these differences cannot be used to salvage the meaning-intention problem in its current form.

What about Neale's syntactic analogue of the original meaning-intention problem? Is there an aphonic-intention problem?

I claim, once again by *reductio ad absurdum*, that (Syntactic Knowledge) simply cannot be supported in light of the data. Consider the following sentence of Spanish:

(11) Quiere comer. [*He wants to eat.*]

According to standard assumptions of generative syntax, in order to comply with the Extended Projection Principle, it is argued that (11) must contain a phonologically null subject, which is typically expressed as *pro*:

(12) *pro* quiere comer.

Pro (distinct from *PRO*) is a phonologically null pronoun, which is the subject of the main clause, and whose semantic value is determined by the speaker's intentions. Importantly, although its existence is not supported by every syntactician, *pro* is an established posit in syntax with independent syntactic support – in other words, it is not something that is just posited by philosophers in order to generate their desired truth conditions.

Precisely the same concerns that Neale raises for aphonic indexicals in 'Silent Reference' would have to apply to *pro* (which is, effectively, an aphonic indexical). Most speakers lack any conceptual grasp of the expression *pro*, and the speaker of (11) might positively deny that she used a tacit referring expression at all. Therefore, if Neale's arguments are sound, we find that we must reject some basic posits of mainstream syntactic theory as well.

Regardless of whether *pro* exists or not, is doubtful that Neale would be sanguine about this consequence. He often takes philosophers to task for their tendency to make non-trivial claims about syntax purely on the basis of philosophical considerations. But that is precisely what we would have to do if we accept (Syntactic Knowledge).

On reflection, we can see that (Syntactic Knowledge) is a rather demanding principle. It effectively implies that linguistic ability and metalinguistic knowledge must proceed “in tandem” – I cannot *use* an expression to refer unless I have conceptual grasp of that expression. Even for overt expressions, however, it is not obvious that we should accept such a principle.

The problem for the Gricean is that (Syntactic Knowledge) is a fairly direct consequence of (RW). But it appears that (Syntactic Knowledge) must be rejected. Therefore, it seems that (RW) must be rejected as well.

III. Tacit States

One direct consequence of the preceding discussion is that (Awareness) should be abandoned. This is all for the good, for the picture of meaning that it presupposes is a strongly Cartesian one. Why should we assume that speakers *do* have privileged access to every aspect of the contents of their speech acts? This is certainly not required in order for them to have meaning-intentions, assuming that such intentions may fail to be fully conscious. And why shouldn't this be the case? Certainly, some strong arguments would be needed to establish that intentions must be conscious; or, at least, that meaning-intentions are special in that they must be conscious. But the Gricean ought to be very

cautious about the latter claim, for her entire program is based around positing a certain kind of complex intention as the basis for communication, where this intention does not simply reveal itself through introspection.

The question, then, is whether (RW) can be salvaged by interpreting the intentional verbs *intend* and *recognize* in terms of tacit states – tacit intentions and tacit recognition.

First, it is worth noting that (RW), as it stands, is inadequate – it fails to provide either necessary or sufficient conditions for referring-with:

(i) Consider the following sentence:

(13) I saw Alice_i and then she_i disappeared.

Here I have referred to Alice, and there are two expressions *with which* I referred to Alice. The problem is only the first instance can satisfy (RW). Recall (RW):

(RW) In uttering x , S referred to o with (or using) e , relative to its i -th occurrence in x , iff for some audience A and relation R , S intended A to recognize that $R(e, x, i, o)$ and, at least partly on the basis of this, that S referred to o in uttering x .

(RW) requires that if a speaker S uses e to refer to o , then S must intend her audience A to recognize that S referred to o in uttering the sentence x , and to do so at least partly on the basis of her utterance of e . But the problem is that A will recognize that S referred to Alice simply on the basis

of the utterance of ‘Alice.’ Since “referring to Alice” really means “expressing an Alice-dependent proposition,” then \mathcal{A} will know that S referred to Alice as soon as ‘Alice’ is uttered (and S will know this). Therefore, \mathcal{A} will recognize (and S will know that \mathcal{A} will recognize) that S referred to Alice purely on the basis of the first reference to Alice – and if this provides a sufficient reason for \mathcal{A} to believe that S referred to Alice (as indeed it does), then S cannot intend for \mathcal{A} to recognize this even *partly* on the basis of the second reference to Alice (‘she’). But, intuitively, S referred to Alice with ‘she.’ Therefore, (RW) does not provide a necessary condition for referring-with.

To resolve this problem, (RW) must therefore be modified along the following lines:

(RW*) In uttering x , S referred to o *with* (or *using*) e , relative to its i -th occurrence in x , iff for some audience \mathcal{A} and relation R , S intended \mathcal{A} to recognize that $R(e, x, i, o)$ and to recognize that $R(e, x, i, o)$ **provides a reason** to believe that S referred to o in uttering x .

This modification avoids the preceding worry because it does not require that S intend that \mathcal{A} ’s belief that S referred to o be derived *on the basis* of \mathcal{A} ’s recognition that $R(e, x, i, o)$, but simply that such recognition provide the hearer with *a basis* for arriving at such a belief (whether or not it is the basis that is in fact used). However, (RW*) does *not* require S to even be referring to o (or even to intend for \mathcal{A} to recognize that she is referring to o) – it only makes the weaker requirement that S intend for \mathcal{A} to recognize that S is doing something that *provides a reason* to believe that S referred to o . This seems too weak. Thus, perhaps the following will suffice:

(RW**) In uttering x , S referred to o *with* (or *using*) e , relative to its i -th occurrence in x , iff for some audience \mathcal{A} and relation R , S intended \mathcal{A} to recognize that $R(e, x, i, o)$ and to recognize that $R(e, x, i,$

o) provides a reason to believe that *S* referred to *o* in uttering *x*, **and to recognize that *S* referred to *o* in uttering *x*.**

(ii) The second problem with (RW) is that it does not only apply to referring expressions. Consider, e.g., a classroom in which there are two teachers, Alice and Bob, and a number of young students, one of whom is named Bob. Alice is talking to Student Bob's parents and says:

(14) Bob will be teaching the class today,

meaning thereby that Teacher Bob would be teaching the class. In uttering (14), Alice relies on the conventional meaning of *Bob* as a name for Bob, and thus, following (RW), refers to Bob with *Bob* insofar as she expects her audience to recognize that she is referring to Bob by uttering *Bob*.

However, she can't expect that uttering *Bob* is *sufficient* for her audience to recognize that she is referring to Teacher Bob, and not Student Bob. Rather, it is the context of the sentence as a whole – or, in particular, the verb 'teaching' – that makes it clear which Bob Alice is referring to. Thus, the right-hand side of (RW) appears to be satisfied by the utterance of 'teaching' as well. (RW) says that a speaker *S* refers to *x* with some expression *e* just in case for some audience *A* and relation *R*, *S* intended *A* to recognize that $R(e, x, i, o)$ and, at least partly on the basis of this, that *S* referred to *o* in uttering *x*. Instantiating the variables: Alice intended her audience to recognize that 'teaching' bears some relation to Teacher Bob and, at least partly on the basis of this, that *S* referred to Teacher Bob in uttering *x*. Thus, (RW) implies that Alice referred to Teacher Bob *with* the word 'teaching.' Since, intuitively, Alice did *not* refer to Bob *with* the word 'teaching,' this implies that (RW) does not provide a sufficient condition for referring-with either.

Unfortunately, no obvious solution to this problem presents itself. The simplicity of (RW) lies in the fact that it analyzes referring-with in terms of offering reasons to believe that one is expressing an *o*-dependent proposition. But, *prima facie*, there is no reason why this condition should be satisfiable by referring expressions only, since other information in the sentence might be intended to help convey what the speaker is referring to, as well.

These problems of definition notwithstanding, the question remains whether (RW) (or (RW**)) is acceptable if one reads the intentional verbs in terms of tacit states.

Two issues must be separated. In discussing tacit states, Neale suggests that ‘tacit’ amounts to ‘unconscious,’ glossing “The Tacit States Reply” as assuming that “S is not ‘consciously aware’ that she means a proposition of the form...”⁴⁹ This appears to be the sense of ‘tacit’ used by Brian Loar (1976) as well. However, Neale also describes the view as assuming “tacit knowledge” in Chomsky’s conception of the term. Chomskyan tacit knowledge is not merely unconscious, but is also functionally isolated – grammatical knowledge is not integrated into the web of belief, i.e., it is not accessible to central reasoning processes. One might assume that these two categories bear some logical relation to each other – for instance, that anything that is in the web of belief is accessible to consciousness – but this is by no means obvious.

Which of these senses is more appropriate for (RW)? (RW) must grant that speakers tacitly know, e.g., that the Spanish sentence *Quiere comer* has an unpronounced nominal expression in subject position. But it is doubtful that all competent Spanish speakers have a concept (even an unconscious one) of the expression *pro*. Rather, this “knowledge” is more akin to the kind of syntactic knowledge

⁴⁹ (Neale 2016, 163).

that comprises the language faculty – not fully conceptual, and not integrated into central reasoning. Therefore, the defender of (RW) must grant that the tacit knowledge mentioned in the definition is tacit in the *strong*, Chomskyan sense.

However, introducing tacit intentions in this manner presents a radical break from traditional Gricean thinking. For one of the apparent advantages of Grice's theory is that it accounts for the ways in which linguistic production interacts with global reasoning abilities. Speakers intentionally produce effects in the audience that are responsive to background knowledge, features of the context, and other factors that interact with central reasoning. By claiming that referential intentions can be tacit in the Chomskyan sense (functionally isolated), the defender of (RW-T) weakens the connection between speaker's meaning and rationality. If one pursues this approach to the end, one may find that the result is no longer distinctly Gricean.

IV. Specificity

I have argued that examples like,

(9) The flash and the bang happened at the same time,

show convincingly that the principles underlying the various aspects of “the meaning-intention problem,” in particular (Awareness) and (Cognitive Resources) are not tenable.

If we reject (Cognitive Resources), then we must reject (Specificity) as well. The problem raised by (Specificity) is that it seems like there ought to be something about the speaker's state of mind, or communicative intentions, that determines *which* of a number of competing alternatives is the actual semantic value of a hidden indexical. But if speakers can refer to *o* (express an *o*-dependent proposition) without even having a concept of *o*, as I claim is demonstrated by sentences like (9), then it cannot be the speaker's specifying intentions that resolve the (Specificity) problem.

However, it is worth noting that (Specificity) is not essentially a problem about language or communication, but rather is a puzzle about thought contents. For example, consider:

(15) The book is covered with paint.

Suppose a speaker utters (15) to communicate that the addressee's favorite book, *War and Peace*, which was left sitting on the dining room table, is covered with paint. The problem relating to (Specificity) is that *the book* does not serve to uniquely pick out the aforementioned book, and there does not seem to be any unique candidate for completing the nominal that is the one that the speaker intended to communicate. This is what is known as an *incomplete description*.

This presents a puzzle about the speaker's *thoughts* about the book, as well. In particular, what is the conceptual content of the speaker's thought, *the book is covered with paint*? One might assume that when a speaker tokens such a thought, there is some description or other (over and above *the book*) under which she thinks of the book. But why would she token one description or the other on any given occasion? Is doing so *necessary* in order for her to entertain thoughts about the book? But surely she needn't token a *uniquely* identifying description of the book, for coming up with such a

description is no simple task. But then it begins to look as though the descriptive information is not doing any cognitive work, and the thought is already complete as it is. But if that's the case, then why should an utterance like (15) not be complete as it is?

In sum, I am arguing that (a) examples like (9) show that (Specificity) must be rejected if (Cognitive Resources) is rejected, as I claim it must be; and, (b) the issue of incomplete definite descriptions is primarily an issue regarding mental contents, and only derivatively an issue pertaining to communication. Furthermore, brief consideration of the matter suggests that incomplete definite descriptions may very well exist in thought as well, which implies that the assumptions behind (Specificity) are misguided.

V. Constraints on Semantic Theories

If these arguments are correct, then the meaning-intention problem is based on implausible assumptions, and cannot be used as a way to refute the hidden indexical theory of belief attribution, or other hidden indexical theories. Speakers can express a proposition that is *o*-dependent even if they entirely lack the concept of *o*. However, there are significant differences between the Extrinsic Parameters Reply as applied to reference frames and as applied to MOPs and other more speculative philosophical entities.

According to Hofweber (1999), the relevant difference is that when it comes to reference frames (with respect to simultaneity) or bodies of mass (with respect to weight), the extrinsic parameter, or unarticulated constituent, is constant for all members of the speech community. Thus while people

failed to realize that weight is a relativized notion, there could never be any confusion or disagreement amongst the members of the community regarding which planet was relevant to their weight statements. This allowed discourse about weight to carry on unproblematically, without the body of mass being specified by the speakers' intentions. This is because the relativization did not affect "sameness and difference (or incompatibility) of contents" amongst speakers.

The same facts do not apply to MOPs. MOPs are *not* shared amongst all members of the community, and there *is* the tendency for confusion when making belief reports.

Hofweber concludes that the problem with the hidden indexical theory of belief reports is that it implies "not only... that speakers have no access (in the strong sense spelled out above) to the content of their utterances, but also no access to sameness, difference and incompatibility of the contents of their utterances."⁵⁰ However, he leaves it open why this fact is problematic.

I suspect that the problem is ill-formulated as once again pertaining to *access*. The reason why speakers must "share the referent" when it comes to genuinely unarticulated constituents (extrinsic parameters) is that testimony would break down otherwise. For two speakers might use the same sentence to express thoughts that concern distinct entities without their being aware of any difference. Thus, if one speaker asserts that sentence to another, the addressee might form a belief that has entirely different truth conditions. Over time, this discrepancy might cause problems and lead to the extinction of the use of the term.

⁵⁰ (Hofweber 1999, 102).

Alternatively, through the course of such discrepancies, speakers may learn to identify in what ways the truth conditions of each speaker's use diverge from each other, and thus may gain a conceptual understanding of the way that the sentence is relativized. This might be an important step in coming to grasp the true nature of the relation being discussed.

In sum, defenders of hidden indexical theories for belief attributions and knowledge claims cannot appeal to extrinsic parameters to supply the referents in such utterances, since the alleged referents are not constant amongst members of the community, or across contexts. The only alternative is to argue that speakers really do grasp concepts such as MOPs and their communicative intentions do serve to pick out properties of such MOPs in the required way. This would be to abandon the Extrinsic Parameters approach, but this leads back to the same problems initially raised by Schiffer and Neale.

VI. Conclusion

The meaning-intention problem has been used by theorists in the Gricean tradition to reject certain semantic hypotheses involving hidden indexical expressions. Hidden indexical theories imply that speakers lack full awareness of their own speech and thought contents, and in some cases lack the relevant concepts entirely. This appears incompatible with the Gricean approach to language, according to which speech act content depends on the speaker's communicative intentions. I have attempted to reverse this dialectic by showing that the assumptions that support the meaning-intention problem are incompatible with some basic data. Since these assumptions are natural ones for the Gricean to make, I have attempted to explore how this affects the project of intention-based

semantics. I have argued that the definition of referring-with that is offered by Schiffer and Neale needs to be reconsidered. It fails to provide either necessary or sufficient conditions for referring with an expression, and it is incompatible with facts about *pro*-drop languages, such as Spanish.

The meaning-intention problem, as traditionally conceived, thus appears to have little force. Nonetheless, I agree with Hofweber that appeals to extrinsic parameters must be tightly constrained. The clear cases of extrinsic parameters, or unarticulated constituents, that I have appealed to (such as weight and simultaneity) involve parameters that are constant amongst the members of the speech community. MOPs and standards of knowledge do not share this feature, and thus the same considerations do not apply.

5. A Gricean Twist on Teleosemantics

I. Teleosemantics and Intention-Based Semantics

Millikan's teleosemantic approach to natural language meaning and Grice's program of intention-based semantics (IBS) have traditionally been viewed as competing accounts of natural language meaning. In *Language, Thought, and Other Biological Categories*, Millikan offers a refutation of the Gricean view of speaker's meaning. But there has otherwise been minimal interaction between these programs. This is due in part to philosophical differences, but also to sociological factors – the literature on teleosemantics draws more heavily on biology and philosophy of science, with less emphasis on classical problems in linguistic semantics, while IBS draws largely from ordinary language philosophy and linguistics.

In one respect, there is no conflict at all between IBS as a theory of linguistic meaning and teleosemantics as a theory of *non*-linguistic representational content. For the Gricean program attempts to explain linguistic concepts like *reference*, *meaning*, etc., in terms of psychological notions like *intention*, *belief*, etc. But no precise theory of mental content is presupposed by the Gricean. Therefore, one might plausibly adopt a Millikanian view of the nature of mental content, and a Gricean view of linguistic content.

Nevertheless, as theories of natural language meaning, the two views have stark differences.

Although Millikan herself rejects such a characterization,⁵¹ Millikan's view is an instance of what Sperber and Wilson (1995) call the *code model* of communication.⁵² On this view, meanings are encoded into linguistic utterances, and interpretation consists in correctly decoding the utterance according to the rules of the language. On such a view, it is the meanings of the expressions themselves that constitutively determine the content of the speech act. However, for Grice, Sperber and Wilson, and other pragmatic theorists, linguistic meaning does not constitutively determine the content of the utterance. Rather, what is linguistically encoded provides the hearer with *evidence* about the speaker's communicative intentions, and it is these intentions alone that constitute the message. Sperber and Wilson describe this as an *inferential* model of communication, since the hearer must do more than mere decoding (however, it is worth noting that the inferential model still does *involve* decoding).

Theorists in the contextualist tradition have argued persuasively that the code model of communication is inadequate to explain a range of linguistic data, and that "Gricean reasoning" (abductive reasoning about the speaker's communicative intentions) is far more pervasive than the classical examples of conversational implicature. Nevertheless, Millikan has argued that the Gricean model is psychologically implausible. I will begin by discussing the ways in which Millikan's code model of communication is inadequate.

II. Millikan's Code Model

⁵¹ Cf. (Millikan 2005, 201).

⁵² Cf. (Sperber & Wilson 1995, 2ff.).

A code model of communication is one in which the propositional content of a speech act is constitutively determined by a set of rules that map sentences to meanings. On the production side, speakers have a thought or message that they want to communicate, so they encode that message into a sentence following the rules of the language. They then produce an utterance in accordance with the phonetic properties of the sentence. The language assigns semantic values to the primitive expressions, and the truth conditions of complex phrases are determined by the meanings of the parts and their syntactic combination. On the interpretation side, hearers identify the sentence uttered on the basis of the phonetic properties of the utterance, and then decode the message by computing the semantic value of the sentence.

The simplest kind of code would be a one-to-one function mapping forms onto propositional contents, but no natural language is a code in this sense. For natural languages involve ambiguity, indexicality, and other features that preclude such a strict mapping. Nevertheless, many theorists wish to maintain that the code model is essentially correct, and must simply be supplemented with a formal treatment of ambiguity and indexicality. It is reasonable to label such theories as falling under the Extended Code Model.

First, the introduction of ambiguity does not significantly alter the code model. Ambiguity may be an undesirable feature if it makes decoding more difficult, but it simply means that the meaning function will be one-to-many, rather than one-to-one. It does not alter the fact that the meaning is determined by an association between forms and meanings.

The introduction of indexicality means that propositional contents cannot always be computed simply on the basis of the rules of the language, but that features of the context of utterance must be factored in as well. However, the way in which context is invoked often appears to be quite rule-governed – e.g., *today* is typically used to refer to the day of utterance. Therefore, one might think that indexicality does not really disturb the code model either, since the language provides rules that determine how certain features of the context of utterance factor into the propositional content.

Millikan's account of natural language falls under the Extended Code Model. According to Millikan, the proper function of an indicative sentence is to produce a true belief in the hearer. The truth conditions are determined by the conventions of the language, in accordance with a *semantic-mapping function*.

What sets linguistic acts and, more generally, communicative acts apart from other acts with cooperative functions is that communicative forms work in part by mapping or, as Wittgenstein put it, 'picturing'. They correspond to states of affairs in accordance with semantic-mapping functions that have been determined by convention.⁵³

On this "picture model," sentences are mapped to states of affairs, and thus, in some sense, encode such states of affairs. The picture model and code model differ mainly in terms of metaphorical emphasis – on the code model, a speaker encodes her thought (which is about some state of affairs) into a signal, while on the picture model, the sentence provides a direct picture of the state of affairs.

⁵³ (Millikan 2005, 63).

In both cases, there is a mapping between sentences and propositional contents, and interpretation does not involve reasoning about the speaker's intentions.

Because of this, Millikan holds that acquiring information through communication is a form of perception.⁵⁴ Under Normal conditions, hearers will hear an utterance, decode that utterance into a proposition following the conventions of the language, and then accept that proposition.⁵⁵ This Normal process need not always occur, of course, for the hearer might be pre-disposed to distrust the speaker, e.g., and so will reject the proposition rather than accept it. But in the absence of overriding conditions, the hearer will accept the proposition that is encoded in the utterance, and since this process simply involves the automatic application of demonstrative inference rules, it is akin to direct perception.

In more recent work, Millikan has argued that her view does *not* involve a simple code model of communication: "But it is important to observe first that semantically conveyed information is *never* simply 'decoded.'"⁵⁶ However, her reasoning simply appeals to the fact that language conventions can overlap, resulting in the same signal having multiple proper functions, i.e., ambiguity.

The deep reason for this is that, on the present analysis, there can be no strictly dedicated forms in a language. Language is just a raggedy collection of reproduced speaker–hearer patterns having various origins and independent histories. There is nothing to prevent the same physical sign pattern from emerging from the

⁵⁴ Cf. (Millikan 2005, 117ff.).

⁵⁵ I follow Millikan in capitalizing 'Normal' when using it in the biological sense, which is not equivalent to 'typical' or 'common.' Cf. (Millikan 1984, 33ff.).

⁵⁶ (Millikan 2005, 201).

employment of separate linguistic conventions; indeed, it seems inevitable that linguistic conventions should sometimes accidentally cross.⁵⁷

Millikan appears to hold that the existence of ambiguity undermines the code model. This is due to the fact that where there is ambiguity, one must consider features of the context of utterance and apply non-demonstrative inference to disambiguate – there are no conventions telling you which convention applies in which circumstance (if there were, it would not be a genuine ambiguity).

“Wider context may always suddenly be needed for interpretation—needed in order to make it clear from which lineage of tokens this token has been copied, hence which is its true memetic function.”⁵⁸

I have argued that ambiguity does not undermine the essence of the code model – an ambiguous code is still a code. However, so as not to be get distracted with terminological issues, the key point is that Millikan’s view is not like the inferential model of communication that Sperber and Wilson propose as an alternative to the code model, for on Millikan’s account, (a) syntactic forms are mapped to complete propositional contents; (b) the proposition expressed is constituted by the conventions that are used (unlike in the inferential model, where the utterance merely provides *evidence* of the proposition expressed); and (c) the hearer makes inferences about *which convention is being used*, but not, typically, about the speaker’s communicative intentions.

Hence, Millikan’s account does fall under the Extended Code Model, and should be considered a code model of communication, despite Millikan’s remarks to the contrary. What distinguishes a code

⁵⁷ (Millikan 2005, 201).

⁵⁸ (Millikan 2005, 202).

model from an inferential model is the fact that the content is completely determined by an associative mapping (supplemented by the context in a suitably rule-governed way), and Millikan's theory retains this basic property.

III. The Perceptual View of Interpretation

How does Millikan reconcile the fact that interpretation is never simply decoding with the claim that interpretation is a form of perception? She attempts to defuse the tension by citing analogous situations in other forms of perception:

If you see one of Wittgenstein's duck-rabbits surrounded by drawings of ducks you will see it as a duck, but if surrounded by rabbits then as a rabbit. If you draw in water surrounding it, you will again see it as a duck. Consider how you immediately see a certain squiggly line in a cartoon strip to be shoes in the character's hand rather than, say, a bottle because the fellow is obviously tiptoeing up the stairs in stockinged feet, the clock on the wall saying 2.30.⁵⁹

As another example, Millikan cites the McGurk effect, in which subjects who are watching a video of a speaker making an utterance hear the same piece of audio as sounding like different phonemes depending on how the speaker's mouth is moving. This is an example of a cross-modal effect.

⁵⁹ (Millikan 2005, 211).

In these cases, “what is perceived is filled in by context.”⁶⁰ Although it is clearly not the same mechanism that explains the McGurk effect and the duck-rabbit example cited above, Millikan claims that a similar kind of process takes place with language comprehension: “Most linguistic forms that have multiple senses are effortlessly disambiguated in linguistic context. Other forms are disambiguated just as easily given the immediate external context.”⁶¹ Even metaphor is “often understood effortlessly in context.”⁶²

Thus, although she does not describe the various mechanisms by which this process of disambiguation is achieved, Millikan clearly holds that ambiguity, and indeed other forms of context-sensitivity, is resolved by the hearer by attending to contextual clues, in a way that is similar to other perceptual effects of context. Importantly, as she stresses, “...no thoughts of speaker’s minds need be involved in the process.”⁶³

Millikan does grant the existence of Gricean implicatures, along with the existence of hearers reasoning about speakers’ communicative intentions in a classically “Gricean way.” However, she holds that such cases are rare, and suggests that they are typically accessible to consciousness.⁶⁴ She claims that generalized conversational implicatures, and indeed many instances of particularized conversational implicatures,⁶⁵ can be recovered by the hearer without explicitly considering the speaker’s mind or intentions, appealing instead to contextual clues, what is obvious, or what “makes

⁶⁰ (Millikan 2005, 210).

⁶¹ (Millikan 2005, 212).

⁶² (Millikan 2005, 212). Stephen Neale (*p.c.*) correctly observes that saying something is done “effortlessly” or “easily” does not constitute much of an explanation.

⁶³ (Millikan 2005, 212).

⁶⁴ Millikan (1984, ch. 3) contrasts the use of conventional signs and improvised signs; her examples of improvised signs involve consciously planned communicative acts.

⁶⁵ Cf. (Millikan 2005, 218).

sense” in the situation. “But how far it is possible to get in understanding such implicatures prior to acquiring a representational theory of mind is not, I believe, obvious a priori.”⁶⁶

Millikan believes that the perceptual approach to interpretation is sufficient to deal with the whole range of cases of context-sensitivity that have been discussed in recent literature:

- (i) *Ambiguity*: Millikan describes ambiguity analogously to our perception of the duck-rabbit. Depending on the visual context, one may see the figure as a duck or a rabbit. “There are surely lots of ways of disambiguating ambiguous sentences and interpreting parasitic usages that do not depend in any way upon having thoughts of speaker intentions – like considering what makes sense or what would obviously be false.”⁶⁷
- (ii) *Sub-sentential utterances*: Sub-sentential utterances require completion in order to express full propositions. Millikan compares this process to the way that one perceives an object by directly seeing only part of it. “One perceives a chair or a cat, but what is responsible for one’s perception is only one side of a part of the chair or the cat.”⁶⁸
- (iii) *‘Ready’-cases*: Sentences like ‘Alice is ready’ have been used to motivate contextualist theories on the grounds that the sentence itself does not seem to encode a full proposition, and speakers might utter such a sentence to mean that Alice is ready for anything. Millikan responds to these cases along the same lines as her response to sub-sentential utterances: “Considering this, if language is a medium of direct perception, it certainly is not surprising that half a sentence, or even a single word uttered in the right context, often can convey the same as a full sentence would have. It is not surprising, for

⁶⁶ (Millikan 2005, 218).

⁶⁷ (Millikan 1984, 65).

⁶⁸ (Millikan 2005, 210).

example, that when some one calls that they are ready, one generally knows for what they are ready.”⁶⁹

- (iv) *Quantifier Restrictions*: Most uses of quantified phrases like ‘the book’ involve implicitly restricted domains of quantification (since otherwise they would fail to be uniquely identifying). As Millikan writes, “Similarly, the content of quantifiers and definite descriptions is nearly always determined by the domain the speaker is focusing on, not just by the immediate context of speaking.”⁷⁰ Thus, Millikan grants the importance of the speaker’s focus, however she claims that all that is required for successful interpretation of an implicitly restricted quantifier is that the hearer be able to track the speaker’s focus (which can be done without employing a representational theory of mind): “The trick is to show how it is possible to make use of natural information about where another person’s mind is focused in interpreting their words without employing a representational theory of mind.”⁷¹

“Somewhat similarly, if you are telling me about an event or a kind of event that I recognize, my mind will be focused where your mind is focused, and I will understand the proper names you use, your descriptions, and the domains of your quantifiers accordingly, without any concern for what’s inside your mind. We are focusing on the same scene or the same type of scene.”⁷²

⁶⁹ (Millikan 2005, 211).

⁷⁰ (Millikan 2005, 213).

⁷¹ (Millikan 2005, 213).

⁷² (Millikan 2005, 215)

- (v) *Vagueness and Comparison Classes*: "...when one hears that France is hexagonal, that a building is hexagonal, and that a machine nut is hexagonal, one interprets these hexagonals to have different degrees of precision."⁷³
- (vi) *Metaphor*: "As for metaphors, gross distortions of the right kind are sometimes more readily recognized in perception than the same items shown more accurately, as illustrated by good caricatures of people's faces. Metaphor is similar. It is often understood effortlessly in context, even sometimes striking a somewhat humorous note just as pictorial caricatures do."⁷⁴
- (vii) *Polysemy*: The possessive suffix appears to admit of a wide variety of relations between the object and its possessor – 'Carol's book' might mean the book that Carol wrote, the book that she loves, the book for which she designed the cover, etc. Thus the possessive suffix is a *polysemous* form – it admits a wide variety of meanings that are related in some way, but not distinct enough to merit individual senses. Here is Millikan on such expressions: "The semantics of a possessive allows it to refer to any pairing relation coupling possessors uniquely with things possessed, such as ownership, physical possession, current responsibility for, and so forth, so that 'John's book' may be the one he own or carries the one he wrote or bought or brought, and so forth. Still it is not surprising that, in context, it is usually easy to perceive immediately what specific kind of relation is being represented."⁷⁵
- (viii) *Generalized Conversational Implicatures*: Generalized conversational implicatures (GCIs) are exemplified by *scalar implicatures* ('Bob has two flowers' implicates that Bob does not have more than two flowers). Millikan claims that GCIs are in fact conventional uses:

⁷³ (Millikan 2005, 211-212).

⁷⁴ (Millikan 2005, 212).

⁷⁵ (Millikan 2005, 212).

“I have suggested that many interpretations of language forms that Grice described as generalized implicatures may actually be entirely conventional. ... Thus, if the form ‘Some As are B’ will be used by a cooperative speaker only when she does not possess information that all of the As are B, then in circumstances in which the speaker may be presumed to know whether all As are B, ‘Some As are B’ will actually entail that not all As are B. Nor does a hearer need to employ a representational theory of mind to be sensitive, in many cases, to whether the speaker has the relevant information.”⁷⁶

Interestingly, this passage reveals that Millikan accepts that ‘Some As are B’ may actually be used in such a way that it *entails* that not all As are B. However, what Grice calls a “conventional implicature” (e.g., ‘but’) does not appear to have the same properties as an entailment, despite being conventionalized. Unfortunately, Millikan does not comment on conventional implicature per se, or on the distinction between conventional implicatures and conventionalized (or standardized) conversational implicatures.

- (ix) *Particularized Conversational Implicatures*: Particularized conversational implicatures (PCIs) appear to present strong evidence that hearers perform inferences about speakers’ communicative intentions in the process of interpretation. For it is in the nature of a PCI that what the speaker implicates is not constrained by the conventional meaning of the sentence uttered. However, Millikan claims that the perceptual model of interpretation is sufficient to handle many (though perhaps not all) cases of PCI as well:

“If A says to be B ‘Let’s go play tennis’ and B replies ‘It’s raining’, the obviously intended

⁷⁶ (Millikan 2005, 217).

interpretation that B is unwilling to play tennis surely is not conveyed conventionally. On the other hand, without entertaining any thoughts of what's in B's mind, the mere information that it is raining would undoubtedly be enough to squelch A's hope that B would be disposed to play tennis, hence would serve as a reply to A's invitation. ... In general, when a speaker's purpose in what he says involves bringing something closely associated or implied to a hearer's mind, a hearer can understand that this effect on her is purposive without using a theory of mind."⁷⁷

As we can see from this overview of Millikan's perceptual response to context-sensitivity, when a hearer encounters some underspecification, ambiguity, or other context-sensitivity, he can arrive at an appropriate interpretation on the basis of contextual clues (linguistic and extra-linguistic), and this process has clear parallels in the way that context-sensitivity affects other forms of perception. Thus, according to Millikan, thinking about the speaker's intentions is rarely required to resolve context-sensitivity: "Employing a full-fledged representational theory of mind in order to interpret what another person is saying or meaning is always a possibility for an adult, but it is plausible, I believe, that this mechanism is seldom required."⁷⁸

IV. The Perceptual Model and the Metaphysics of Meaning

Millikan offers the perceptual model as a way of defusing the Gricean argument that utterances which exhibit underspecification or context-sensitivity can only be interpreted by reasoning about

⁷⁷ (Millikan 2005, 218)

⁷⁸ (Millikan 2005, 213).

the speaker's intentions; as an alternative, she suggests that interpretation may proceed in a manner akin to perception.

Importantly, in Millikan's view, language perception proceeds directly from hearing an utterance to forming a belief about the state of affairs pictured – this process need not (and usually does not) involve reasoning about the speaker's intentions. In Normal circumstances, the speaker will “purpose”⁷⁹ that the hearer accept what she is saying, but this purpose need not be conscious and need not be an explicit intention – the behavior can be purposive so long as it is adapted to produce that effect. Similarly, the hearer is not required to recognize the behavior as purposive in order for interpretation to proceed – the hearer simply perceives the world *through* language. The *speaker's purpose*, then, appears to play no special role in the theory, except that under Normal conditions, the speaker's purpose and the memetic purpose of the utterance will coincide. It certainly is given no special role with respect to the process of interpretation.

This raises a fundamental issue with the perceptual approach to language interpretation. Millikan grants that the speaker's purpose and the memetic purpose of an utterance need not coincide (e.g., with non-literal speech). In this case, the utterance will not serve its function in accordance with a Normal explanation, since one of the Normal conditions (that the speaker's purpose and the memetic purpose coincide) is absent. Therefore, whatever interpretation the hearer derives will not be generated in accordance with a Normal explanation.

⁷⁹ Millikan offers the verb ‘purpose’ to describe a behavior which has a proper function, but this function is not derived from the agent's intentions (the purpose is not represented by the agent).

Nevertheless, if the hearer's purpose and the memetic purpose of an utterance fail to coincide, this raises the question: *what constitutes the content of the speech act that is being performed?* There seem to be two possibilities: the speaker's purpose, and the memetic purpose of the utterance.

Millikan does not commit to either purpose as uniquely determining the content of an utterance. Rather, under Normal conditions, the two purposes coincide. When they diverge, neither one constitutes the "correct" meaning of the utterance; whatever interpretation the hearer derives, the sentence will not be able to serve its function in accordance with a Normal explanation, since Normal conditions do not hold. Therefore, it appears that teleosemantics does not stipulate either purpose as having metaphysical priority over the other with respect to the content of a speech act.

This is a fundamental problem with the perceptual approach to interpretation. For let us consider the perceptual model once more. Millikan appeals to examples like the duck-rabbit to show how perception frequently involves "filling in" from the surrounding context. But since a duck-rabbit is an illustration, Millikan's discussion focuses on perception as an epistemic process, noting the different contexts in which we perceive the image as a duck or a rabbit, and she ignores the question of veridicality – *is it a duck or a rabbit?* Since it's just an illustration, the question does not really arise. Now let us consider visual perception: if you see an animal and aren't sure whether it's a duck or a rabbit, then you may be inclined to see it one way rather than another depending on contextual factors. But ultimately, whether it *is* a duck or a rabbit is independent of its context – it is simply a biological fact. Thus, we have visual perception as an epistemic process, which is directed towards things in the world, and which causes us to have beliefs about such things, where the veridicality of those beliefs is determined not by the perceptual process itself but by its correspondence with an independent reality.

How does this work in the case of “linguistic perception”? Consider the following case:

- (1) Alice will have to take her check to the bank.
- (2) Bob is at the bank with his fishing pole.

Given the linguistic context provided by (8), one is strongly inclined to think that the meaning of (8) is that Alice will have to take her check to the financial bank, and not the river bank. Similarly, the most natural reading of (9) is that Bob is at the river bank with his fishing pole. Adopting Millikan’s model, we might say one perceives ‘bank’ *as* financial bank in (8), given the surrounding discourse context (cashing a check). But this is simply a description of the perceptual process. Perception can be veridical or non-veridical. Therefore, we must also say what makes it the case that the hearer has perceived the utterance correctly.

In the case of ambiguity, which of a pair of homonyms is being uttered depends, for Millikan, on the causal history of the utterance, and which linguistic convention it involves, for ambiguity is simply a matter of “criss-crossing” conventions (separate conventions that end up sharing the same form). Thus, regardless of whatever I perceive the utterance of ‘bank’ *as*, its actual semantic value is constituted independently of my grasp of it. This means that for ambiguous expressions, the perceptual model is acceptable since there is a clear account of what *constitutes* the difference between two uses of ‘bank’ beyond the fact that hearers may be disposed to *interpret* them differently.

However, this is not the case with all context-sensitive expressions. For consider ‘ready’-cases:

(3) Alice is ready [*to jump out of the airplane*].

(4) Alice is ready [*to begin the meeting*].

The context-sensitivity exhibited by ‘ready’ is not a form of ambiguity, and thus we cannot explain the different uses in terms of criss-crossing conventions. There is just the one convention, and the word ‘ready’ allows an open-ended set of complements. Therefore, the question one again arises – what determines the “correct” interpretation of a sentence like (3)/(4)?

As we have seen above, Millikan assimilates ‘ready’-cases to instances of perceptual completion – just as one can perceive an object by directly seeing only part of it, one can perceive the proposition that Alice is ready to φ by hearing the words ‘Alice is ready.’ As Millikan writes, “It is not surprising, for example, that when some one calls that they are ready, one generally knows for what they are ready.”⁸⁰ Millikan also illustrates this concept by considering the way one identifies objects in a photograph by taking into account the context of the scene: “If that’s a tiger, then that’s bars in the zoo and, now it’s clear, that’s another cage with a black panther in it, not a toy or a domestic cat.”⁸¹

Millikan’s account strikes me as a *prima facie* plausible description of how interpretation proceeds in some cases. However, the important point is that *it can only be an account of the mechanisms of interpretation* – the perceptual view does not explain what *constitutes* whether a given interpretation is correct or not. To extend Millikan’s zoo example, it is plausible to assert that reasonable subjects will infer from seeing a picture with bars and a tiger that there is a panther in the picture, and not a stuffed animal; but these contextual factors do not metaphysically *determine* that the object in

⁸⁰ (Millikan 2005, 211).

⁸¹ (Millikan 2005, 215)

question is a real animal – that is a matter that is independent of perception or context. Similarly, there must be something independent of perception that constitutes whether a given interpretation is correct or not.

I have argued that in the case of ambiguity, there *is* a clear answer that Millikan can give as to what determines the correct interpretation, independently of perception – namely, which convention is causally implicated in the utterance. But it is impossible to assimilate all types of context-sensitivity to ambiguity. When this fails (as I have argued it does for ‘ready’-cases), we are left without a metaphysical basis for ascribing content to a speech act.⁸²

If the linguistic conventions do not suffice to determine the propositional content of a speech act, there are two plausible candidates to fill in the gap: (a) what the speaker intends; (b) what a “reasonable hearer” would interpret the utterance to mean. Given the fact that Millikan regards the speaker’s intentions as typically irrelevant to the interpretive process, one might think that she would endorse (b). However, (b) is untenable.

⁸² Some theorists working on pragmatics, such as Michael Devitt (2013a), and Lepore & Stone (2014), have argued that the contextualists have overplayed their hand, and that many alleged cases of “free enrichment” or conversational implicature are in fact conventional. For instance, Devitt (2013a) defends “the traditional view that what is said is constituted by properties arising from linguistic conventions, disambiguations, and reference fixings...” and, he argues that “many of Pragmatism’s striking examples exemplify semantic rather than pragmatic properties.” Similarly, Lepore & Stone (2014, 3-4) claim that “speakers can SIGNAL apparently pragmatic meaning through linguistic devices that allow them to explicitly mark what they intend their utterances to do. Griceans have been wrong to categorize these cases as [Conversational Implicatures]...” These views are in keeping with Millikan’s approach to communication as almost entirely governed by conventional meaning, and my criticisms of such an account should be brought to bear against them as well.

First, (b) does not uniquely determine a correct interpretation in all cases because there can be reasonable disagreement amongst reasonable hearers regarding how to interpret an utterance. Second, (b) entails that if there is a choice between two interpretations, the correct interpretation is always the one that is more “reasonable” (likely to be true, consistent with the context, etc.). But meanings are not so constrained. A speaker might utter ‘Alice is taking her check to the bank’ and mean thereby that Alice is taking her check to the riverbank, even though interpreting ‘bank’ as ‘financial institution’ generally makes more sense in that kind of context. The same applies to ‘ready’-cases: it can’t be the case that what metaphysically determines the correct interpretation of an utterance of ‘ready’ depends on which reading a hearer would find more plausible.

Thus, we are led to the conclusion that the only thing that can plausibly complete the metaphysical gap between what is conventionally encoded and the propositional content of a speech act is *what the speaker intends*.

Millikan incorrectly offers the perceptual model as an alternative to the Gricean view of communication. However, at best, the perceptual model offers a description of the mechanisms of interpretation that does not make use of reasoning about the speaker’s intentions. It can in no way be construed as an account of the metaphysics of meaning, or of what determines whether an interpretation is *correct*.

For Millikan, the question of what constitutes a correct interpretation simply does not arise. This is because utterance comprehension is not really a matter of interpreting what someone else is trying to say, but rather of directly perceiving the world through their speech. Under Normal conditions, hearers arrive at a true belief about the world without considering the speaker’s thoughts or

intentions at all. However, this picture is shown to be inadequate in light of the fact that the conventions of language do not provide a metaphysical basis for speech act content in the face of radical underspecification. Millikan's attempt to alleviate this problem by appealing to perception misses the target, since perception epistemically reveals objects to the mind, but does not provide their metaphysical basis (ignoring forms of idealism, for the time being).

One final issue with the perceptual view of utterance interpretation is that the phenomenology of interpretation exhibits disanalogies with the phenomenology of perception. For one thing, perceptual illusions are persistent even when the subject has knowledge of the actual state of affairs. Thus, with the McGurk effect, even if I know that the audio is really a recording of a speaker saying "ba," I will still hear it as "fa" when I am watching the appropriate video. However, if I come across an utterance of "John brought his fishing pole to the bank," and I know the speaker was referring to a financial institution, there is nothing inhibiting me from accessing the appropriate reading of 'bank,' even though it is not the reading that is strongly suggested by the immediate linguistic context. This shows that the effect of context is less persistent in the case of linguistic interpretation than in other cases of perception.⁸³

V. The Proper Function of Linguistic Utterances

The failure of Millikan's program to provide a coherent metaphysics of meaning in the face of underspecification stems from the fact that she grants no special metaphysical status to speakers'

⁸³ Thanks to Stephen Neale for this observation.

intentions; she even claims that although speech is purposive (directed towards some end), it need not involve intentions at all (representations of the goal), and certainly not higher-order intentions.

I have argued that the metaphysical gap between what is conventionally encoded and the propositional content of a speech act can only be filled by the speaker's intentions. However, the question still remains whether hearers need to *recognize* this intention in the Normal course of interpretation. Millikan claims that because adult speakers do possess a representational theory of mind, it is always available to them to reflect on the purposiveness of linguistic behavior; however, for Grice, recognizing the speaker's intention is an essential part of speaker's meaning – the speaker produces an utterance *with the intention of the hearer recognizing that intention*.⁸⁴

Millikan ignores what is distinctive about human communicative exchanges – namely, that they involve *ostension*. Humans do not simply transmit information from one to another; rather, speakers intend that their informative intentions be “out in the open” – in Relevance-theoretic terms, they intend that their informative intentions be part of the mutual cognitive environment.

Millikan appears to dismiss the idea that there is something special about ostensive communication:

Suppose that you pass on the pickles to me as they circle around the table during dinner. You do this purposing that I take the pickle dish from your hand. If I reflected, I would know that you purpose this. And if you reflected, you would know that if I reflected I would know that you purpose this, etc. Are we to conclude that

⁸⁴ I do not mean to imply that such recognition must be conscious. Grice's views on the consciousness of communicative intentions are unclear.

there is something very special about acts of handing people things – a special “mechanism” – that puts these acts in a category deserving a large philosophical literature?⁸⁵

In fact, there *is* something special about acts that are performed in such a way that the actor intends her intention to be out in the open. There is an important difference between holding up the plate of pickles with the expectation that the person next to you will unthinkingly take them and pass them on, and holding up the plate of pickles with the intention that the person next to you *recognize* your intention to pass the pickles, and take them in part on the basis of that recognition. In other words, this is the distinction between *passing* the pickles to someone, and holding the pickles where the other person might take them. Of course, the passer might not explicitly intend for the recipient to recognize that the passer intends that the recipient recognize, etc. But it is crucial that the passer intend her act to be out in the open, that the recipient recognize this, and take the pickles partly on the basis of that recognition.

Thus, with ostensive communication, speakers intend for their informative intentions to be part of the mutual cognitive environment. And, in the case of meaning, they intend that recognition of that intention be part of the hearer’s reason for accepting what the speaker says. This means that speakers’ intentions play a much more central role in communication than they are given in Millikan’s theory.

Nevertheless, it is possible to accommodate Grice’s central insights in a teleosemantic framework. Recall that for Millikan, the proper function of an indicative sentence is to cause a true belief in the

⁸⁵ (Millikan 1984, 64).

hearer. However, this relies on a picture model of interpretation that was found to be inadequate. Since linguistic conventions cannot always determine the propositional content of a speech act, I argued that speakers' intentions must be given metaphysical priority. Accepting a statement as true is in fact a downstream effect of forming a hypothesis about what the speaker *means* in making that statement.

Thus, the goal of interpretation is to interpret *what the speaker means*, and the accuracy of an interpretation is judged not by comparing it to the world, but by comparing it to what the speaker intended. If this is the case, we can retain the framework of teleosemantics so long as we claim that the proper function of utterances is to cause a true belief in the hearer *about the speaker's intention*.

According to Millikan, the proper function of a linguistic form is determined by its *stabilizing function* – that is, whatever it does that promotes the interests of the producer and consumer mechanism and occurs often enough to stabilize the behavior.

Conventional language forms are selected for performing services satisfactory at once to both partners in communication. ... Because the conventional function of a linguistic form will remain stable only if it continues to serve the interests of both speakers and hearers often enough, I call it a 'stabilizing function'.⁸⁶

Millikan claims that the stabilizing function of an indicative sentence is to produce a true belief in the hearer, since this is generally in the interests of speaker and hearer, and if it did not happen often enough, people would stop producing indicative sentences. "And if hearers ceased ever using

⁸⁶ (Millikan 2005, 58).

indicative sentences as guides in forming beliefs, speakers would stop trying to use them for purposes that required imparting beliefs.”⁸⁷

Causing a true belief about the speaker’s communicative intentions is a stabilizing function of all linguistic forms. Whether or not what the speaker is saying is true, it is generally in the interests of both speaker and hearer for the hearer to form a correct hypothesis about the speaker’s meaning. It may be true (for a variety of reasons) that if speakers did not produce true statements frequently enough, then people would no longer make statements at all; but it is equally plausible that the survival of the indicative sentence form depends on the hearer correctly recovering speaker’s meaning. Therefore, the claim that the proper function of a linguistic utterance is to produce a true belief about the speaker’s intentions is consistent with Millikan’s conception of stabilizing function.

Millikan notes that some traits may have a variety of different proper functions. Some of these are “disjunctive,” in the sense that they are functionally independent – e.g., a bird’s feather’s might have the proper function of allowing the bird to fly and keeping it warm. Other functions may be “serial functions,” in the sense that performing one function depends on successfully performing another – e.g., the human heart has the proper function of pumping blood, which circulates oxygen, which allows humans to think and speak. However, we do not think of speaking as the proper function of the heart. Circulating oxygen is the last in the series of functions before they split into disjunctive functions (enabling the brain to think, enabling muscles to move, etc.). This last member in the series is known as the *focused function*: “[The focused function] will be the last function that such a

⁸⁷ (Millikan 2005, 58).

device performs before its serial functions finally diverge for good, being merely disjunctive and/or conjunctive after that point.”⁸⁸

With this terminology in place, we can see that the *focused function* of a linguistic utterance is to produce a true belief in the hearer about the speaker’s meaning. All other functions performed by uttering an indicative sentence (such as informing the hearer of something, telling a story, telling a joke, etc.) are disjunctive functions that depend for their successful performance on the hearer forming a correct hypothesis about the speaker’s meaning.⁸⁹

Thus, I claim that if we apply Millikan’s definitions of what counts as “the function” of a linguistic form, we are led to the conclusion that the (focused, stabilizing) proper function of a linguistic utterance is to cause a true belief in the hearer about the speaker’s communicative intentions. This Gricean “twist” on Millikan’s teleosemantic theory (or, equally, a Millikanian twist on intention-based semantics) retains the core definitions offered by Millikan, but also takes into account the failure of the picture/perception model of communication to provide an adequate metaphysics of meaning.

⁸⁸ (Millikan 1984, 36).

⁸⁹ McDowell (1998) objects that this kind of view falsifies the nature of communication, which is primarily about communicating information about the external world, and not about the speaker’s mind. McDowell (1998, 38) writes: “...the primary point of making assertions is not to instill into others beliefs about one’s own beliefs, but to inform others – to let them know – about the subject matter of one’s assertions...” I need not deny that this is an important proper function of communicative behavior – however, it is simply a more distal, disjunctive function. The more proximate, and indeed *focused*, function of communicative behavior is to cause true beliefs about the speaker’s mind, and it is *this* that we are interested in when we talk about “*the* proper function” of an utterance (even if utterances are often performed with the downstream goal of causing a true belief about the external world).

Furthermore, I claim that this view is closely aligned with Grice's own conception of linguistic meaning (timeless meaning), as determined by the notion of "having a certain procedure in one's repertoire."⁹⁰ Having a procedure in one's repertoire is a matter of having a disposition to utter an expression of a certain type when one has a certain kind of communicative intention. Grice does not offer a rigorous definition of the repertoire, but his comments on the matter are highly reminiscent of Millikan's discussion of conventional behaviors:

If U utters HW [Hand-Wave], his measure of expectation of success as regards effecting the intended response obviously depends (as has already been remarked) on A's knowledge of U's procedure; and normally, unless the signal is to be explained to each A, on A's repertoire containing the same procedure.⁹¹

A procedure in one's repertoire has a cooperative function, and it is reproduced because others have the same kind of procedure in their repertoire. There are obvious similarities to Millikan's notion of convention as a reproduced pattern of behavior that is reproduced largely by "weight of precedent": "To be thought of as conventional, a reproduced pattern must be perceived as proliferated due, in important part, to weight of precedent, not to its intrinsically superior capacity to produce a desired result, or due, say, to ignorance of any alternatives."⁹²

I claim that Millikan's notion of a "natural convention" can be used to elucidate the Gricean idea of "having a certain procedure in one's repertoire." Alternative accounts of convention, such as the Lewis-Schiffer account involving mutual knowledge, require extremely complex, and perhaps

⁹⁰ (Grice 1989, 126).

⁹¹ (Grice 1989, 126-127).

⁹² (Millikan 2005, 7).

unrealizable, cognitive states. Other writers in the Gricean tradition have spoken of “self-perpetuating regularities.”⁹³ However, as Millikan rightly observes, “The conventions of language are not regularities, either *de facto* or *de jure*.”⁹⁴ Millikan’s notion of convention is highly intuitive, robustly defined, and does not appeal to complex mental states or regularities. Thus, it avoids one of the major problems that has long plagued Gricean theorizing – how to understand mutual knowledge.

VI. Conclusion

Millikan presents a highly systematic account of the nature of linguistic meaning. However, despite her remarks to the contrary, I have argued that her view is an instance of the code model of communication, in which the propositional content of a speech act is constitutively determined by the conventions of the language. In more recent work, Millikan has attempted to respond to the pervasive context-sensitivity that has been identified in the literature on pragmatics. Her strategy is to assimilate the resolution of context-sensitivity in language to other, well-attested contextual effects and cross-modal effects in perception. This allows her to retain the view that language comprehension is a form of direct perception of the world, in the face of the problems raised by underspecification.

I have argued that the picture/perceptual model of linguistic meaning is fundamentally flawed, due to the fact that it fails to provide a coherent metaphysics of meaning. Perceptual completion can

⁹³ Cf. (Neale 2016, 259). Devitt (2013b) also speaks of conventions arising from (though not constituted by) “regular use.”

⁹⁴ (Millikan 2005, 26).

only ever be an account of the mechanisms of interpretation – it does not offer an account of what constitutes a *correct* interpretation, whenever this question cannot be resolved by appealing to the underlying conventions in use. Thus, sentences do not present pictures of states of affairs, and the gap between what is conventionally encoded and what is expressed cannot be filled (metaphysically) by perceptual processes.

The major flaw in Millikan's framework is that it grants no special role to speaker's intentions in the communicative process. Thus, it neglects what is special about ostensive communication. Speakers do not simply intend for certain outcomes to be achieved (such as getting the hearer to accept a piece of information); rather, they intend for such intentions to be out in the open (mutually known, or, part of the mutual cognitive environment).

I claim that we can retain the key insights of Grice's *and* Millikan's programs by adopting the view that the proper function of a linguistic utterance is to cause a true belief in the hearer *about the speaker's intentions*. This resolves the difficulties inherent in the picture/perceptual approach to language, and is consistent with Millikan's own definition of focused stabilizing function.

Furthermore, we can use Millikan's definition of a conventional pattern of behavior to clarify Grice's sketchy remarks regarding having a procedure in one's repertoire.

Works Cited

- Barwise, J., & Perry, J. 1981. Situations and Attitudes. *The Journal of Philosophy*, 78.11: 668-691.,
- Crimmins, M. 1992. *Talk About Beliefs*. Cambridge, MA: MIT Press.
- Devitt, M. 2013a. Three Methodological Flaws of Linguistic Pragmatism. In C. Penco & F. Domaneschi (eds.), *What is Said and What is Not: The Semantics/Pragmatics Interface*. Stanford: CSLI Publications: 285-300.
- Devitt, M. 2013b. What Makes a Property “Semantic”? In A. Capone et al. (eds.), *Perspectives on Pragmatics and Philosophy*, Cham: Springer: 87-112.
- Fodor, J. 1975. *The Language of Thought*. Harvard University Press: Cambridge, MA.
- Fodor, J. 2001. Language, Thought and Compositionality. *Mind & Language*, 16.1: 1-15.
- Grice, H. P. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Hofweber, T. 1999. Contextualism and the Meaning-Intention Problem. In K. Korta et al. (eds.), *Cognition, Agency and Rationality: Proceedings of the Fifth International Colloquium on Cognitive Science*, 93-104.
- Lepore, E., & Stone, M. 2014. *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford: Oxford University Press.
- Lewis, D. 1979. Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8: 339-359.
- Loar, B. 1976. The Semantics of Singular Terms. *Philosophical Studies*, 30.6: 353-377.
- McDowell, J. 1998. *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press.
- Millikan, R. G. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Millikan, R. G. 2005. *Language: A Biological Model*. New York: Oxford University Press.
- Neale, S. 2016. Silent Reference. In G. Ostertag (ed.), *Meanings and Other Things: Essays in Honor of Stephen Schiffer*. Oxford University Press: 229-344.
- Neale, S. 2007. On Location. In M. O’Rourke & C. Washington (eds.), *Situating Semantics: Essays on the Philosophy of John Perry*, Cambridge, MA: MIT Press (251-393).
- Perry, J. 1986. Thought Without Representation. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 60: 137-151.
- Perry, J. 1998. Indexicals, Contexts and Unarticulated Constituents. In A. Aliseda-Llera, et al. (eds.), *Proceedings of the 1995 CSLI-Armsterdam Logic, Language and Computation Conference*. CSLI Publications.
- Recanati, F. 2002. Unarticulated Constituents. *Linguistics & Philosophy*, 25: 299-345.
- Salmon, N. 1986, 1991. *Frege’s Puzzle*. Atascadero, California: Ridgeview Publishing Company.
- Schiffer, S. 1992. Belief Ascription. *The Journal of Philosophy*, 89.10: 499-521.
- Schiffer, S. 1996. Contextualist Solutions to Scepticism. *Proceedings of the Aristotelian Society*, 96.1: 317-333.
- Searle, J. 1978. Literal Meaning. *Erkenntnis*, 13: 207-224.
- Sennet, A. 2011. Unarticulated Constituents and Propositional Structure. *Mind & Language*, 26.4: 412-435.
- Soames, S. 1987. Direct Reference, Propositional Attitudes and Semantic Content. *Philosophical Topics*, 15: 47-87.
- Sperber, D., & Wilson, D. 1995, 1986. *Relevance: Communication and Cognition: Second Edition*. Cambridge, MA: Blackwell.
- Sperber, D., & Wilson, D. 2002. Pragmatics, Modularity, and Mind-reading. *Mind & Language*, 17.1/2: 3-23.
- Stanley, J. 2007. *Language in Context*. New York: Oxford University Press.

Wittgenstein, L. 1953. *Philosophical Investigations*. In Anscombe, Hacker, & Schulte (transl.), *Ludwig Wittgenstein: Philosophical Investigations* (2009), Malden, MA: Wiley-Blackwell.